Existence and uniqueness of positive-Q log-correlated metrics in dimension $d \geq 3$

Abstract

We consider exponential path metrics associated with the whole-space logarithmically correlated Gaussian field in every dimension $d \geq 3$. The parameter is defined initially by the upper point-distance exponent. Whenever $Q > 0$, the argument proves the genuine limit $a_\varepsilon = \varepsilon^{1-\xi Q+o(1)}$, identifies it with the crossing exponent, and proves a bounded comparison of their normalizers, constructs a separating lower-semicontinuous extended length metric, and identifies its normalization and its axiomatic uniqueness. The construction uses coordinate averages of path contacts and finite connected cores. The uniqueness argument uses nonseparation by finite unions of embedded interval subsets and a one-sided restoration argument.

Contents

1 Statement and conventions **3**

2 The field, the approximation, and the exponent **4**

2.1 Gaussian estimates used in the proof 8

2.2 Comparison of buffered point geometries 10

2.3 Finite-range counting and smooth changes of field 11

2.4 Many good annuli simultaneously at every point 12

3 A multiscale point connection estimate **13**

4 The crossing exponent and point moments **17**

4.1 An annular exponent 17

4.2 Small point moments from the upper point growth rate 21

5 Lower crossing bounds and averaged barriers **22**

6 One marked-network functional **25**

7 Uniform envelopes and simultaneous records **28**

8 A common core for arbitrary crossing arms **31**

9 Concentration of the marked-network functional **33**

10 The bounded point-to-crossing comparison **36**

11 Quantitative output for the metric construction **38**

12 From annular estimates to a metric **40**
12.1 Simultaneous annuli 41
12.2 A countable construction 41
12.3 Compact paths and the finite component 43
12.4 Recovery and limiting connectors 44
12.5 Actual intrinsic metrics 45
12.6 Boundary ports with one-sided access 46
12.7 Continuous weights 48
12.8 Conditional products before field measurability 49

13 Comparison, determination by the field, and rigidity **51**
13.1 Lower estimates for an arbitrary strong metric 52
13.2 A prefix comparison 54
13.3 Removal of residual randomness 60
13.4 Contact geometry and span estimates 63
13.5 A protected star 65
13.6 The signed local contract 67
13.7 Uniform certificates and two-sided rigidity 69
13.8 Exact covariance and identification of the normalization 85
13.9 Thin-shell restoration for a one-sided competitor 87
13.10 Protected trees with small unmarked pieces 89
13.11 Germ gates and a pure-discount contract 91
13.12 Strong witnesses and the final one-sided count 94

A Compactness, conditional laws, and measurable choices **98**

B Topological foundations **104**
B.1 A closed partition of an interval 104
B.2 Scalar extension and finite trees 105
B.3 Finite cochains and compact-cover limits 106
B.4 Separation computed by finite cochains 109
B.5 Finite unions of embedded closed interval subsets 111
B.6 Nonseparation by embedded interval subsets 112

References **113**

1 Statement and conventions

Throughout the paper, $d \geq 3$.

An *extended metric* is a symmetric function $D : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty]$ satisfying the triangle inequality and $D(x, y) = 0$ if and only if $x = y$. Its restriction to a finite-distance component is an ordinary metric. Lengths and continuity of curves are understood in that metric. For a Euclidean-open set U, the internal distance $D(x, y; U)$ is the infimum of the D-lengths of curves in U from x to y; an empty infimum is $+\infty$. A length metric equals its own global internal metric. For a continuous positive function w, $w \cdot D$ is the length metric obtained by integrating w against D-arclength.

The topology used for distance functions is convergence of lower-semicontinuous functions: $D_n \to D$ means

$$D(x, y) \leq \liminf_n D_n(x_n, y_n) \qquad \text{whenever } (x_n, y_n) \to (x, y), \tag{1.1}$$

$$D(x, y) = \lim_n D_n(x_n, y_n) \qquad \text{for some } (x_n, y_n) \to (x, y). \tag{1.2}$$

Limits and inequalities include $+\infty$. Equivalently, use the explicit countable epigraph model in Lemma A.4, which also proves metrizability. A *nonsingular* point is one having finite distance to another point.

The field h, its unit-sphere pinning, the approximants D_h^ε, and their internal point medians a_ε are defined in Section 2. Initially $Q(\xi)$ denotes the upper exponent

$$Q(\xi) = \frac{1}{\xi}\left(1 + \limsup_{\varepsilon \downarrow 0} \frac{\log a_\varepsilon}{\log(1/\varepsilon)}\right).$$

The positive-Q theorem below includes the assertion that this upper limit is a limit; no exponent-existence result is assumed.

The following axioms are the higher-dimensional form of the strong-metric axioms of Ding and Gwynne [5].

Definition 1.1 (Strong assignment)**.** For fixed $\xi, Q > 0$, a strong assignment is a Borel rule $g \mapsto D_g$ with values in separating Euclidean lower-semicontinuous extended metrics. Whenever g is the whole-space field plus a possibly dependent random continuous function, it has the following properties.

(i) Almost surely, D_g is a length metric.

(ii) For every deterministic open U, its actual internal metric $D_g(\cdot, \cdot; U)$ is almost surely determined by $g|_U$.

(iii) Almost surely, simultaneously for all continuous $f : \mathbb{R}^d \to \mathbb{R}$,

$$D_{g+f} = e^{\xi f} \cdot D_g. \tag{1.3}$$

(iv) For each fixed $r > 0$ and $z \in \mathbb{R}^d$, almost surely simultaneously for all x, y,

$$D_g(rx + z, ry + z) = D_{g(r\,\cdot + z) + Q \log r}(x, y). \tag{1.4}$$

(v) For every deterministic connected open U and deterministic disjoint compact connected nonsingleton sets $K_1, K_2 \subset U$,

$$D_g(K_1, K_2; U) := \inf_{x \in K_1,\, y \in K_2} D_g(x, y; U) < \infty \quad \text{almost surely.} \tag{1.5}$$

The last condition is an infimum between two sets. It asserts neither finite diameter of those sets nor finiteness at their every point.

Theorem 1.2 (Positive-Q existence, uniqueness, and convergence)**.** *Let $d \geq 3$ and $\xi > 0$, and suppose that $Q(\xi) > 0$. Then $a_\varepsilon = \varepsilon^{1-\xi Q(\xi)+o(1)}$, and a strong assignment with parameters $(\xi, Q(\xi))$ exists. It can be normalized uniquely by*

$$\operatorname{Med} D_h(0, e_1; B_2(0)) = 1. \tag{1.6}$$

With this normalization,

$$a_\varepsilon^{-1} D_h^\varepsilon \longrightarrow D_h \quad \text{in probability} \tag{1.7}$$

in the topology (1.1)–(1.2)*, jointly with h. The normalized internal distance in* (1.6) *also converges in probability. Its limit is positive, finite, atomless, and has unique median one. For each fixed $\lambda > 0$,*

$$\frac{a_{\lambda\varepsilon}}{a_\varepsilon} \longrightarrow \lambda^{1-\xi Q}. \tag{1.8}$$

If $\widetilde{D}$ is any other strong assignment with these parameters, then there is a deterministic $c \in (0, \infty)$ such that almost surely

$$\widetilde{D}_h(x, y) = c D_h(x, y) \qquad (x, y \in \mathbb{R}^d). \tag{1.9}$$

The equality holds on the same base-field event after every continuous Weyl perturbation. The nonsingular points form one complete geodesic component. Every fixed deterministic pair has finite internal distance in every fixed connected open domain containing it, and metric-bounded sets are Euclidean-bounded.

The hypothesis $Q > 0$ will be used to make endpoint chains summable and to make a coarse Gaussian influence decay. There is no hypersurface-diameter hypothesis. The proof first establishes the quantitative estimates needed for construction, then identifies actual internal distances and conditional local laws, and finally proves rigidity and fixes the scalar normalization.

Relation to earlier work

Ding, Gwynne, and Zhuang [6] develop the arbitrary-dimensional model and prove subcritical tightness for $Q > \sqrt{2d}$. Contreras Hip and Zhuang [3] establish the corresponding subcritical weak-metric framework and quantitative estimates. The distinction between point and set-to-set distance exponents is also central in [7]. The planar existence and uniqueness theory is established in [11, 4, 5]; this paper treats only $d \geq 3$.

The references identify earlier results and the sources of the methods used here. The positive-Q exponent comparison, bounded normalizer comparison, metric construction, and rigidity argument are proved below. The probability and topology appendices include the specialized representation, selection, and separation statements needed for the proof.

2 The field, the approximation, and the exponent

Throughout, $d \geq 3$, $\xi > 0$, $|\cdot|$ is the Euclidean norm, $B_r(z) = z + (-r, r)^d$, and $e_1 = (1, 0, \ldots, 0)$. All quantiles are lower quantiles. All distances constrained to an open set are actual internal length distances. A closed constraint will always be specified explicitly.

Choose a nonnegative radial $K \in C_c^\infty(\mathbb{R}^d)$ with $\|K\|_2 = 1$, and let W be white noise on $\mathbb{R}^d \times (0,\infty)$ with Lebesgue control measure. Set

$$H_t(x) = \int_{2^{-t}}^1 \int_{\mathbb{R}^d} u^{-(d+1)/2} K((x-y)/u)\, W(dy,du), \qquad t \geq 0. \tag{2.1}$$

Thus $\operatorname{Var} H_t(x) = t \log 2$, and disjoint scale intervals give independent fields. If K is supported in the Euclidean ball of radius R_K, noise below scale r determines the field in a set U using only noise in $(U + \overline{B}^{\mathrm{Eucl}}_{R_K r}) \times (0, r]$.

Write $H_{0,\infty}$ for the distribution obtained by integrating over $0 < u < 1$. This is a well-defined random distribution: for a smooth compactly supported test function φ, Young's inequality bounds the squared integrand norm by $u^{d-1}\|K\|_1^2\|\varphi\|_2^2$, which is integrable at zero. Define

$$g(x) = \int_1^\infty \int_{\mathbb{R}^d} u^{-(d+1)/2} \{K((x-y)/u) - K(-y/u)\}\, W(dy,du). \tag{2.2}$$

The field g has a smooth version. For example, the variance of its increment is at most $C|x - x'|^2 \int_1^\infty u^{-3}\, du$; the analogous estimates for all derivatives give local smoothness. Let σ be normalized surface measure on the Euclidean unit sphere, and put

$$h = H_{0,\infty} + g - (H_{0,\infty} + g)(\sigma). \tag{2.3}$$

Lemma 2.1 (Covariance, pinning, and scaling)**.** *For smooth compactly supported zero-integral tests φ, ψ,*

$$\operatorname{Cov}(h(\varphi), h(\psi)) = \iint \varphi(x)\psi(y) \log \frac{1}{|x-y|}\, dx\, dy. \tag{2.4}$$

The spherical pairing in (2.3) *is well defined. Its variance is finite, and* $\sup_V |g| + |(H_{0,\infty} + g)(\sigma)|$ *has a Gaussian upper tail on every fixed compact V. For $r > 0$ and $z \in \mathbb{R}^d$, the field $h(r\,\cdot + z) - h_r(z)$ has the same law as h, where $h_r(z) = h(\sigma_{r,z})$. On a fixed compact set of centers,*

$$\operatorname{Var} h_r(z) = \log(1/r) + O(1) \quad (r \downarrow 0), \qquad \operatorname{Var} h_R(0) = \log R + O(1) \quad (R \uparrow \infty). \tag{2.5}$$

Proof. Let $C_K(v) = \int K(w)K(w-v)\, dw$. This is radial, compactly supported, satisfies $C_K(0) = 1$, and $1 - C_K(v) = O(|v|^2)$ at zero. For $a, b > 0$, a change of variables gives

$$\lim_{T\to\infty} \int_0^T \{C_K(a/u) - C_K(b/u)\}\, \frac{du}{u} = \log(b/a).$$

Consequently the all-scale covariance is $-\log|x-y|$ up to terms killed by zero-integral tests. The anchoring term in (2.2) is killed in the same way. This proves (2.4).

For $0 < u < 1$, a spherical-cap estimate gives

$$\left\| u^{-(d+1)/2} K(\cdot/u) * \sigma \right\|_2^2 \leq C u^{d-2}.$$

Indeed the convolution is supported in an $O(u)$-neighborhood of the sphere, and is at most $C u^{(d-3)/2}$ there. The displayed bound is integrable because $d \geq 3$. It also justifies approximation of σ by smooth measures in the Gaussian L^2 space. The large-scale part is the average of the smooth field g. Gaussian supremum estimates give the stated tail.

Under $x \mapsto rx + z$, the covariance kernel changes by the constant $-\log r$. Centering by the unit-sphere average removes that constant, and the transformed centered field has covariance (2.4)

and zero unit-sphere average. These two properties determine its Gaussian law. Finally, pairing the logarithmic kernel with $\sigma_{r,z} - \sigma$ twice proves (2.5): the self-interaction of the radius-r sphere is $-\log r$ plus a constant, whereas its interaction with the fixed unit sphere stays bounded on compact sets of centers. For large concentric spheres the same calculation, after scaling, gives $\log R + O(1)$. □

Lemma 2.2 (Spectral variance from the scale integral)**.** *With the Fourier convention $\widehat{\varphi}(\zeta) = \int e^{-2\pi i x\cdot\zeta}\varphi(x)\,dx$, there is a positive constant κ_d such that every smooth compactly supported zero-integral test satisfies*

$$\operatorname{Var} h(\varphi) = \kappa_d \int_{\mathbb{R}^d} |\widehat{\varphi}(\zeta)|^2 |\zeta|^{-d}\,d\zeta.$$

The same formula, by polarization and Gaussian L^2 completion, gives the covariance inner product of all restriction coordinates used below.

Proof. The zero integral removes the anchoring and pinning terms, so $h(\varphi)$ is the all-scale white-noise integral of the convolution of φ with $u^{-(d+1)/2}K(\cdot/u)$. The white-noise isometry, Plancherel's identity, and Tonelli's theorem give

$$\operatorname{Var} h(\varphi) = \int_{\mathbb{R}^d} |\widehat{\varphi}(\zeta)|^2 \left(\int_0^\infty u^{d-1} |\widehat{K}(u\zeta)|^2\,du \right) d\zeta.$$

Radiality and the substitution $v = u|\zeta|$ identify the inner integral, for $\zeta \neq 0$, as

$$|\zeta|^{-d}\kappa_d, \qquad \kappa_d = \int_0^\infty v^{d-1} |\widehat{K}(ve_1)|^2\,dv.$$

This constant is finite and positive: $\widehat{K}$ decays faster than every power and $\widehat{K}(0) = \int K > 0$. In fact polar coordinates and $\|K\|_2 = 1$ show that it is the reciprocal of the unit-sphere area, so it depends only on d. At the origin $\widehat{\varphi}(\zeta) = O(|\zeta|)$, and at infinity it decays rapidly. These bounds also directly verify integrability of the displayed variance formula. Polarization and completion prove the remaining assertion. □

The next lemma combines decay of annular covariance with the Gaussian Hölder inequality of Chen, Dafnis, and Paouris [2]. We include the finite-dimensional semigroup argument.

Lemma 2.3 (Gaussian block inequality for annular restrictions)**.** *Fix a compact annular template with positive inner radius and fixed buffers. For every $p > 1$, there is an integer L such that the following holds. Let $r_i = 2^{-m_i}$ with successive m_i separated by at least L, and let their centers satisfy $|z_i - x| \le \kappa r_i$ for a common x and a sufficiently small fixed κ. Let $\mathcal{A}_i$ be the Gaussian sigma field generated by the restriction of h to the ith annulus modulo additive constants. For nonnegative $\mathcal{A}_i$-measurable random variables F_i,*

$$\mathbb{E} \prod_{i=1}^N F_i \le \prod_{i=1}^N (\mathbb{E} F_i^p)^{1/p}. \tag{2.6}$$

In particular, if local bad events have marginal probability at most q, then the probability that at least δN of them occur is at most $2^N q^{\delta N/p}$.

Proof. By Lemma 2.2, the squared Gaussian variance norm of a zero-mass test measure μ is a constant times $\int |\widehat{\mu}(\zeta)|^2 |\zeta|^{-d}\,d\zeta$. On tests supported in a fixed compact template this norm controls

the usual inhomogeneous $H^{-d/2}$ norm. Consider two rescaled zero-mass tests at radii $r \le s$, and put $\epsilon = r/s$. In unit coordinates their cross-covariance kernel, after using the zero mass of the smaller test, can be written

$$-\log|v + \delta_s - \epsilon(u + \delta_r)| + \log|v + \delta_s|.$$

Here the shifts δ_r, δ_s range in fixed small compact sets. Since the larger annulus stays away from the common center, this kernel and all its derivatives on the two fixed templates are $O(\epsilon)$, uniformly when the scales are sufficiently separated. Multiply by fixed smooth cutoffs around the templates. Sobolev duality in each variable bounds the bilinear form by $C\epsilon$ times the product of the two Gaussian variance norms. For clarity, the required product Sobolev norm of the kernel is bounded by a finite sum of its L^2 derivative norms of sufficiently large integer order, all of which are $O(\epsilon)$. This argument first applies to smooth tests; Gaussian L^2 closure gives every restriction variable.

Thus the cross-covariance operator of blocks i, j, after their marginal covariance norms are made the identity, has norm at most $C2^{-|m_i - m_j|}$. The sum of the off-diagonal operator norms in each row is at most $2C\sum_{k\ge 1} 2^{-kL}$. Choose L so that this is at most $p - 1$. The covariance matrix Σ of any finite collection of orthonormal coordinates in the blocks then satisfies $\Sigma \le pI$, and each diagonal block is the identity.

Here is the Gaussian inequality for such a matrix. For smooth bounded positive functions f_i, let $g_i(t, \cdot) = (P_t(f_i^p))^{1/p}$, where P_t is the standard Ornstein–Uhlenbeck semigroup on the ith marginal space, and put $w_i = \log g_i$. Its generator is $\Delta - x \cdot \nabla$, so

$$\partial_t w_i = (\Delta - x \cdot \nabla) w_i + p|\nabla w_i|^2.$$

Gaussian integration by parts for the joint vector X gives

$$\frac{d}{dt}\mathbb{E}\prod_i g_i(t, X_i) = \mathbb{E}\left[\prod_i g_i(t, X_i)\left(p\sum_i |\nabla w_i|^2 - \sum_{i,j} \nabla w_i \cdot \Sigma_{ij} \nabla w_j\right)\right] \ge 0.$$

Let $t \to \infty$ to obtain (2.6). Truncation, regularization, and approximation by finitely many restriction coordinates extend it to the asserted random variables. For indicators of a specified set of k bad events the bound is $q^{k/p}$. Sum over subsets of size at least δN. □

Fix a nonnegative radial $\rho \in C_c^\infty(\mathbb{R}^d)$ with integral one, and write $\rho_\varepsilon(x) = \varepsilon^{-d}\rho(x/\varepsilon)$. The approximation under consideration is

$$D_h^\varepsilon(x, y; U) = \inf_{P: x \to y,\, P \subset U} \int_0^1 e^{\xi(h * \rho_\varepsilon)(P(t))} |P'(t)|\, dt. \tag{2.7}$$

Here the infimum is over piecewise C^1 paths. Define

$$a_\varepsilon = \operatorname{Med} D_h^\varepsilon(0, e_1; B_2(0)). \tag{2.8}$$

Define the upper distance exponent by

$$Q = Q(\xi) := \frac{1}{\xi}\left(1 + \limsup_{\varepsilon \downarrow 0} \frac{\log a_\varepsilon}{\log(1/\varepsilon)}\right). \tag{2.9}$$

Elementary Gaussian supremum bounds make this a finite real number. When $Q > 0$, we will prove that the upper limit is a limit, so that $a_\varepsilon = \varepsilon^{1 - \xi Q + o(1)}$. Thus existence of a point-distance

exponent, equality with a crossing exponent, and bounded comparison of their normalizers will all be conclusions, rather than hypotheses.

The stationary convolution field and the corresponding path metric are

$$X_t = H_{0,\infty} * \rho_{2^{-t}}, \qquad D_t(x,y;U) = \inf_{P:x\to y,\, P\subset U} \int_P e^{\xi X_t}\,|dz|.$$

For the sharp approximation, the same notation with a superscript sh uses H_t instead of X_t. When an argument applies to either approximation we suppress the superscript. At $0 \le k \le t$ there is the exact independent decomposition

$$X_t = H_k * \rho_{2^{-t}} + X^f_{k,t}, \qquad X^f_{k,t}(2^{-k}\,\cdot + z) \overset{\text{law}}{=} X_{t-k}. \tag{2.10}$$

For H_t, replace the coarse convolution by H_k and the fine convolution by $H_{k,t}$. Every use of a bounded-factor estimate for convolution metrics below uses (2.10) directly.

2.1 Gaussian estimates used in the proof

The two inequalities below are the Gaussian concentration consequence of Gross's logarithmic Sobolev inequality [9] and the resampling variance bound of Efron and Stein [8]. Their required forms are proved here.

Lemma 2.4 (Gaussian concentration and resampling)**.** *Let F be a function of an isonormal Gaussian process, with Cameron–Martin Lipschitz constant L. Whenever F is integrable,*

$$\mathbb{P}(|F - \mathbb{E}F| > u) \le 2e^{-u^2/(2L^2)}. \tag{2.11}$$

For independent coordinates (Z_i), let $Z^{(i)}$ replace Z_i by an independent copy. Then

$$\operatorname{Var} f(Z) \le \sum_i \mathbb{E}\big[(f(Z^{(i)}) - f(Z))_+^2\big]. \tag{2.12}$$

Both statements hold conditionally on independent additional data.

Proof. For a finite standard Gaussian vector, the Gaussian logarithmic Sobolev inequality is $\operatorname{Ent}(e^v) \le \frac{1}{2}\mathbb{E}(e^v|\nabla v|^2)$. One way to obtain it is to differentiate $\mathbb{E}[P_t(e^v)\log P_t(e^v)]$ along the Ornstein–Uhlenbeck semigroup, integrate from 0 to ∞, and use $|\nabla P_t f| \le e^{-t}P_t|\nabla f|$ and Cauchy–Schwarz. Apply it to $v = sF$. Integration of the resulting differential inequality for $s^{-1}\log \mathbb{E}e^{sF}$ gives $\log \mathbb{E}e^{s(F-\mathbb{E}F)} \le s^2L^2/2$. Chernoff's bound, for positive and negative s, proves (2.11). Conditional expectations on finitely many Gaussian coordinates preserve the Lipschitz constant; martingale approximation proves the isonormal version.

For (2.12), reveal the independent coordinates successively and sum the orthogonal squared martingale differences. Conditional Jensen bounds their sum by $\frac{1}{2}\sum_i \mathbb{E}(f(Z^{(i)}) - f(Z))^2$. The pair $(Z, Z^{(i)})$ has an exchangeable law, so the two squared one-sided differences have equal expectations. Truncation and L^2 approximation give the countable-coordinate statement. □

Lemma 2.5 (Local Gaussian estimates)**.** *On every fixed bounded working region, the following bounds hold. Constants may depend on that region, K, ρ, and fixed exponents.*

1. *For $m \ge 0$ and $u > 0$,*

$$\mathbb{E}\exp\{u\sup|H_m|\} \le e^{C_u(m+1)}, \qquad \mathbb{E}\exp\{u\sup|X_m|\} \le e^{C_u(m+1)}.$$

2. *If $T_m(x)$ is a ball of radius $C(1+m)^A 2^{-m}$ about a deterministic x, then*

$$\mathbb{E}\exp\{u \sup_{T_m(x)} H_m\} \le C_u(1+m)^{C_u} 2^{u^2 m/2}. \tag{2.13}$$

The assertion is uniform in deterministic centers in a compact set.

3. *Let O_m be the maximum oscillation of H_m on all the polynomially enlarged scale-2^{-m} balls in a fixed working region. For every $\delta > 0$ and large m,*

$$\mathbb{P}(O_m > \delta m) \le C_\delta \exp\{-c_\delta m^2/\log(m+2)\}. \tag{2.14}$$

The same local oscillation bounds hold for the terminal convolution field at its cutoff scale. With a fixed enlargement instead of a polynomial one, the maximum oscillation over all scale-2^{-m} cells satisfies $\mathbb{E}e^{uO_m} \le e^{C_u\sqrt{m+1}+C_u}$.

Proof. The kernel formula gives

$$\mathbb{E}(H_m(x) - H_m(y))^2 \le C\int_{2^{-m}}^1 \min\{1, |x-y|^2/u^2\}\,\frac{du}{u}.$$

In particular the increment variance on a fixed cutoff-sized ball is bounded, and on a polynomially enlarged ball it is $O(\log(m+2))$. Smoothness of K gives uniformly bounded rescaled derivative variances on a cutoff-sized ball. A dyadic grid refinement and the elementary bound $\mathbb{P}(|N(0,v)| > t) \le 2e^{-t^2/(2v)}$ bound the expected supremum of its increment process by a constant. Covering the polynomial enlargement by polynomially many such balls gives increment supremum exponential moments bounded by a power of $1+m$. Covering a fixed region requires $O(2^{dm})$ balls, and gives expected absolute field supremum $O(m+1)$. Apply Lemma 2.4 to these suprema and oscillations. The maximal increment variance, together with the number of balls, gives (2.14) and the stated exponential moment for fixed enlargements.

For the sharper estimate (2.13), tilt the Gaussian law by $e^{uH_m(x)-u^2 m\log 2/2}$. Under this tilt the mean of each increment $H_m(y) - H_m(x)$ changes by

$$u\,\mathrm{Cov}(H_m(y) - H_m(x), H_m(x)) = -\tfrac{u}{2}\,\mathrm{Var}(H_m(y) - H_m(x)) \le 0.$$

The centered increment covariance is unchanged. Its supremum exponential moment is therefore at most the polynomial bound just proved. Restoring the Gaussian tilt proves (2.13). The convolution assertions follow from the same kernel estimates and Young's inequality; positivity and unit mass of ρ can also be used for the coarse part. □

Lemma 2.6 (Leading-order regularization comparison)**.** *On each fixed bounded region,*

$$\sup|H_t - X_t| = O_{\mathbb{P}}(\sqrt{t+1}). \tag{2.15}$$

Replacing t by $t+v$, with v in a fixed bounded interval, has the same bound. Pinning and adding the large-scale correction change path lengths by factors bounded by the exponential of a random variable with a fixed Gaussian tail. Consequently the sharp, stationary convolution, and pinned convolution point exponents agree. The same assertion holds for any two nonnegative radial compactly supported kernels of unit L^2 norm in the smooth class specified above, and also for the normalized unit-ball indicator.

Proof. Put $\varepsilon = 2^{-t}$. In the coarse integral $u \geq \varepsilon$, smoothing changes the kernel by a squared L^2 amount at most $C(\varepsilon/u)^2$ times the original scale variance. Integration against du/u is bounded. The smoothed fine integral $u < \varepsilon$ has variance at most

$$C \int_0^{\varepsilon} u^{d-1} \varepsilon^{-d} \, du \leq C.$$

Thus $H_t - X_t$ has bounded point variance. On an ε-sized ball its rescaled increment supremum has bounded mean and Gaussian tails. A union bound over $O(\varepsilon^{-d})$ balls proves (2.15). A bounded scale band is treated identically. The correction in (2.3) is controlled by Lemma 2.1.

For two kernels, construct the all-scale pinned fields separately. The covariance calculation in Lemma 2.1 shows that their laws coincide. Disintegrate the two white-noise realizations over one common pinned field. Compare each sharp cutoff with the mollification of that common field, as above. For the indicator kernel the squared L^2 translation increment is $O(|v|)$ rather than $O(|v|^2)$; the coarse error is then bounded by $C \int_{\varepsilon}^{1} (\varepsilon/u) \, du/u$. The local increment processes have a uniform positive Hölder exponent in their Gaussian canonical metric, which suffices for the same local supremum bound. Its large-scale correction is continuous, which is enough here; no smoothness of that correction is asserted. Finally $e^{-\xi \sup |f-g|} D_g \leq D_f \leq e^{\xi \sup |f-g|} D_g$ for every constrained path problem. The logarithmic error is $o_{\mathbb{P}}(t)$, proving the claim about exponents. This argument does not give bounded-factor comparisons of quantiles. □

2.2 Comparison of buffered point geometries

Lemma 2.7 (Buffered point geometries)**.** *Let $(x, y; U)$ range over a fixed compact family of buffered point geometries: the endpoints have separation bounded above and below, all domains lie in a fixed bounded region, and they contain connecting polygonal routes with uniformly bounded numbers of segments and a uniform positive clearance. For every fixed $p \in (0, 1)$,*

$$\left| \log q_p(D_n(x, y; U)) - \log a_{2^{-n}} \right| \leq C\sqrt{n+1}, \tag{2.16}$$

uniformly over the family. The assertion holds for sharp and convolution cutoffs, and uniformly over bounded fractional parts of the cutoff. In particular, the same bound holds with any fixed reference point quantile in place of $a_{2^{-n}}$. No exponent limit is used in this comparison.

Proof. For a fixed path problem in a bounded domain, the logarithm of the distance has Cameron–Martin Lipschitz constant at most $\xi\sqrt{n \log 2 + C}$. Indeed a Cameron–Martin shift of norm one changes the field at each point by at most that point's standard deviation; infima of positive path lengths preserve the resulting multiplicative bound. Lemma 2.4 shows that any two fixed quantiles differ by a factor $e^{O(\sqrt{n})}$, uniformly in the geometry. A bounded change of cutoff also changes the quantiles by $e^{O(\sqrt{n})}$, by Lemma 2.6 and a union bound; the quantile levels can first be moved to fixed slightly larger or smaller ones.

To obtain the upper comparison, cover a buffered polygonal route by finitely many translated, rotated, and scaled copies of the reference point geometry. All scale factors are bounded above and below. The exact split (2.10) says that the corresponding fine distances have the reference laws at cutoffs shifted by bounded amounts. The coarse field on each copy has a fixed supremum tail. Choose a high fixed quantile in each copy and take a union bound. The triangle inequality gives $q_p(D_n(x, y; U)) \leq e^{C\sqrt{n+1}} a_{2^{-n}}$, uniformly.

For the reverse comparison, rotate $y - x$ to the first coordinate axis. Shrink a copy of U by a fixed factor so that its endpoint displacement is e_1/M, where the integer M is sufficiently large.

Place M translated copies consecutively from 0 to e_1. Their domains all lie in $B_2(0)$ if M was large enough. The scale factor and the required cutoff shift stay in fixed compact intervals for the entire family. Applying the same exact split, fixed-quantile comparison, and union bound to these copies gives $a_{2^{-n}} \le e^{C\sqrt{n+1}} q_p(D_n(x,y;U))$. The copies need not be independent. This proves (2.16); bounded fractional parts were already included in the cutoff comparison. □

2.3 Finite-range counting and smooth changes of field

Lemma 2.8 (Sparse dependent bad sites)**.** *Suppose labels on $\mathbb{Z}^d$ have dependence range at most R, and every site is bad with probability at most p. If p is small enough in terms of d, R, then the probability that a bad $*$-connected set of diameter at least m meets a finite set A is at most $C|A|e^{-cm}$. The same bound holds for the reflected and folded labels used in a fixed orthant.*

Proof. A self-avoiding $*$-path with m vertices contains at least m/J vertices whose mutual distances exceed R, where J depends only on d, R. The probability that it is all bad is at most $p^{m/J}$. There are at most $(3^d)^m$ such paths from a given site. Sum this geometric series and then sum over A. A reflected connected set folds to a connected original set; the finite number of reflected copies changes diameter and counting constants only by a fixed factor. □

For later Gaussian changes of measure, we use the Cameron–Martin formula [1]. Smooth compactly supported functions are Cameron–Martin shifts of the field modulo constants; the required norm and density formula are derived next. In Fourier coordinates the squared Cameron–Martin norm is a fixed positive multiple of

$$\|f\|_{\mathcal{H}}^2 = c_d \int_{\mathbb{R}^d} |\widehat{f}(\zeta)|^2 |\zeta|^d \, d\zeta. \tag{2.17}$$

To verify this identification, Lemma 2.2 identifies the Gaussian covariance Hilbert space with the completion in the norm $\kappa_d \int |\widehat{\varphi}|^2 |\zeta|^{-d}$. The linear shift functional $\varphi \mapsto \int f\varphi$ has squared dual norm $\kappa_d^{-1} \int |\widehat{f}|^2 |\zeta|^d$, by Cauchy–Schwarz. Equality follows by approximating its weighted Fourier representative with smooth functions supported away from zero and infinity, and then approximating their inverse transforms by compactly supported zero-integral tests. The last approximation uses a smooth spatial cutoff and subtracts its vanishing integral error times one fixed test of integral one; rapid decay makes the error tend to zero in the displayed norm. The Hilbert-space Riesz representative gives a Gaussian linear functional G_f with this variance and with $\mathrm{Cov}(G_f, h(\varphi)) = \int f\varphi$. Tilting by its Gaussian exponential therefore shifts every test coordinate by $\int f\varphi$, as follows by the finite-dimensional Gaussian density formula. This proves the asserted shift and norm. In particular the norm is finite for $f \in C_c^\infty$, and invariant under translation and dilation of f's argument. In the pinned representative the shift is $f - f(\sigma)$; the subtracted constant is accounted for by constant Weyl scaling. If G_f is the centered Gaussian linear functional of variance $\|f\|_{\mathcal{H}}^2$, then the Radon–Nikodym derivative for the shift is

$$\exp\{G_f - \tfrac{1}{2}\|f\|_{\mathcal{H}}^2\}. \tag{2.18}$$

The formula is the elementary translated Gaussian density in finite dimensions and follows in the Gaussian Hilbert space by projection and L^1 convergence of the exponential martingale.

2.4 Many good annuli simultaneously at every point

We record the precise independence argument used for local annular events. It also applies when local metrics carry residual randomness, provided their local conditional kernels have the product property proved later. It does not assume such metrics are already functions of the field.

Lemma 2.9 (Annular density from independent fine blocks)**.** *Let U_0 be a bounded annular region contained in $\{a < |x|_\infty < b\}$, where $0 < a < b < \infty$. Fix a sufficiently small $\gamma > 0$, and at a location (z, r) split the field at scale γr. A fine local label is permitted to use only noise in*

$$(z + rU_0 + \overline{B}^{\mathrm{Eucl}}_{R_K\gamma r}) \times (0, \gamma r],$$

or an independent local conditional kernel over that noise. Suppose the following two properties hold uniformly in location, scale, and allowed cutoff.

1. *For every $p > 0$ there are deterministic quantitative parameters for which a fine label is bad with probability at most p.*
2. *If the fine label is good and the oscillation of the deleted coarse field on its working buffer is at most M, then the desired full-field conclusion holds, with deterministic parameters depending only on M and the parameters of the fine label.*

For every prescribed $B < \infty$ there are choices of the parameters, an integer scale separation L, and $\kappa > 0$ with the following property. For any N admissible labels at radii $r_i = 2^{-m_i}$ with $m_{i+1} - m_i \geq L$ and $|z_i - x|_\infty \leq \kappa r_i$ for a common x, the probability that fewer than half of them have the full-field conclusion is at most $2e^{-BN}$. Consequently, on a common probability-one event, every point of a fixed compact region is surrounded by a positive density of successful annuli in every sufficiently large dyadic logarithmic block. Their working buffers can be taken disjoint, and all have one common deterministic set of quantitative parameters.

Proof. Take $\kappa < a/8$ and $R_K\gamma < a/8$. The spatial footprint of the ith fine block lies in

$$\{(a - R_K\gamma - \kappa)r_i < |y - x|_\infty < (b + R_K\gamma + \kappa)r_i\}.$$

Choose L large enough that these shells are disjoint for successive admissible radii. Their whole fine-noise regions are then disjoint, even though their scale intervals overlap. The fine labels are independent. A union bound over subsets gives

$$\mathbb{P}(\text{at least } N/4 \text{ fine labels fail}) \leq 2^N p^{N/4}.$$

Choose p sufficiently small for this to be at most e^{-BN}, with a harmless increase of B to absorb constants.

It remains to control the coarse oscillations without asserting their independence. Write the deleted coarse field at scale γr_i as G_i. Its pinning scalar disappears from $G_i(u) - G_i(z_i)$. For u_i in the corresponding fixed-factor working buffer, let $v_i(u_i)$ be the white-noise kernel of that increment. Kernel translation estimates at scale s give

$$\big|\langle v_i(u_i), v_j(u_j)\rangle\big| \leq C \int_{\max(\gamma r_i, \gamma r_j)}^{\infty} \min\{1, r_i/s\} \min\{1, r_j/s\} \, \frac{ds}{s} \leq C_\gamma \frac{\min(r_i, r_j)}{\max(r_i, r_j)}.$$

The same estimate holds when the coarse field is mollified at a smaller cutoff, by convolution and positivity. It follows that $S = \sum_i \sup |G_i - G_i(z_i)|$ has squared Cameron–Martin Lipschitz constant

at most $C_\gamma N$: express S as a supremum of signed sums of these kernels and sum the geometrically decaying covariance bound. Each summand has a uniformly bounded expectation, by scaling and the smooth Gaussian increment estimates. Hence $\mathbb{E}S \le C_\gamma N$. If at least $N/4$ coarse oscillations exceed M, then $S \ge MN/4$. Lemma 2.4 makes this event have probability at most e^{-BN} once M is sufficiently large. The union of the two exceptional events proves the first assertion. No independence between a fine label and any coarse field was used.

For the last assertion, in a fixed compact region take the grid of mesh $2^{-2a}/C_0$ and, about each grid point, candidate radii 2^{-m} for $a \le m \le 2a$ in one residue class modulo L. There are $O(2^{2da})$ grid points and at least $c_L a$ candidate radii. Choose B so that $2^{2da}e^{-Bc_L a}$ is summable. The preceding estimate and Borel–Cantelli give success at a positive fraction of the candidate radii, simultaneously at all grid points, for all sufficiently large a. Every point is within $2^{-2a}/C_0$ of a grid point. Choose C_0 large enough that it lies in the inner region of every candidate annulus about that grid point. Fixed strict buffers and the scale separation keep the enlarged working annuli disjoint. Exhausting compact regions and taking countable intersections gives the stated common event. For a prelimit cutoff only stages wholly above that cutoff are used; each fixed finite stage is admissible for all sufficiently fine approximations. □

3 A multiscale point connection estimate

This section uses the stationary finite-range approximation

$$X_n = H_{0,\infty} * \rho_{2^{-n}}, \qquad X^{m,n} = H_{m,\infty} * \rho_{2^{-n}}, \qquad 0 \le m \le n.$$

The scale decomposition and its exact scaling law are

$$X_n = H_m * \rho_{2^{-n}} + X^{m,n}, \qquad \left(X^{m,n}(2^{-m}x)\right)_{x\in\mathbb{R}^d} \overset{\text{law}}{=} \left(X_{n-m}(x)\right)_{x\in\mathbb{R}^d}. \tag{3.1}$$

Write D_n for the length metric with density $e^{\xi X_n}$. All cubes in this section use the maximum norm. Let b_n be the lower q-quantile of $D_n(0, e_1; B_2(0))$, where $q < 1$ will be chosen sufficiently close to one. In particular b_0 denotes this quantile at cutoff zero; it is a positive finite constant. The point-geometry comparison identifies the upper exponent of these quantiles with the parameter in (2.9):

$$Q_p = \frac{1}{\xi}\left(1 + \limsup_{n\to\infty} \frac{\log_2 b_n}{n}\right) = Q. \tag{3.2}$$

The chaining construction itself uses neither an exponent limit nor a ratio estimate for b_{n-m}/b_n.

Lemma 3.1 (Finite-range bad clusters)**.** *Let labels on $\mathbb{Z}^d$ have a dependence range at most R_0: labels at sets of vertices separated by more than R_0 are independent. Suppose each vertex is bad with probability at most p. There are $p_0, c, C > 0$, depending only on d, R_0, such that, for $p \le p_0$, a box containing at most N vertices meets a bad star-connected component of diameter at least L with probability at most CNe^{-cL}.*

Proof. A star path of length k contains at least k/C_0 vertices at pairwise distance greater than R_0: select one vertex and discard the at most $C_0 = C_0(d, R_0)$ vertices within distance R_0, then repeat. For a fixed self-avoiding star path, independence bounds the probability that all its vertices are bad by p^{k/C_0}. There are at most $(3^d)^k$ star paths of length k starting at a specified vertex. Choose p_0 so that $3^d p_0^{1/C_0} < 1$. A component of diameter at least L contains a self-avoiding star path of length at least $L/2$ starting in, or within one lattice step of, the box. Summing the geometric bound over starting vertices proves the assertion, with an immaterial enlargement of the box. □

Lemma 3.2 (Deterministic lattice routing)**.** *Suppose only finitely many vertices of $\mathbb{Z}^d$ are bad, $d \geq 3$, and every bad star-connected component has diameter at most L. The good vertices have a unique infinite nearest-neighbor component $\mathcal{O}$. Every vertex lies within distance $C_d(L+1)$ of $\mathcal{O}$. If $u, v \in \mathcal{O}$ can be joined by a Euclidean path in an open set V, they can be joined by a good nearest-neighbor path in the $C_d(L+1)$-neighborhood of V. In particular, if u, v belong to a box of diameter $C_0(L+1)$, the joining path can be chosen in a box of diameter $C_1(L+1)$ and with at most $C_2(L+1)^d$ edges.*

Proof. We first prove the discrete boundary fact used below. If S is a finite nearest-neighbor connected subset of $\mathbb{Z}^d$, its *inner exterior boundary* and *outer exterior boundary* are star-connected. The outer boundary consists of vertices in the unbounded component of $\mathbb{Z}^d \setminus S$ adjacent to S; the inner boundary consists of their neighbors in S. There is only one unbounded complementary component: outside a box containing S all vertices are connected by coordinate paths around the box, using $d \geq 3$.

Here is a finite combinatorial proof of the boundary assertion. Fill every bounded complementary component of S, obtaining $\widehat{S}$. Each filled component has a neighbor in S, so $\widehat{S}$ is nearest-neighbor connected; its complement is the unbounded component and is connected as well. The finite set $\mathcal{E}$ of nearest-neighbor edges between these two shores is therefore a minimal nonempty edge cut: deleting a proper subset leaves both shores connected and at least one edge joining them. Join two edges of $\mathcal{E}$ whenever they belong to the same unit coordinate square. We claim this edge adjacency graph is connected. Otherwise choose a nonempty proper component $\mathcal{E}_0$. Every unit square contains an even number of edges of $\mathcal{E}$, by following the two-valued indicator of $\widehat{S}$ around that square. All cut edges in any one square belong to the same adjacency component. Thus every unit square contains an even number of edges of $\mathcal{E}_0$.

Every closed nearest-neighbor lattice walk has even intersection count with $\mathcal{E}_0$, counting repeated traversals. Indeed, interchanging consecutive steps in different coordinate directions changes the walk by the boundary of one unit square and preserves this parity. Remove consecutive opposite steps, which also preserves parity. By these operations move all steps in the first coordinate together and cancel them, then do the same for each remaining coordinate. A closed walk has zero total displacement in each coordinate, so the resulting walk is empty. Consequently the parity of $\mathcal{E}_0$ edges on a path from one fixed vertex to x is a well-defined function $\psi(x) \in \{0, 1\}$. Across an edge it changes exactly when that edge is in $\mathcal{E}_0$. Since $\mathcal{E}_0$ is nonempty, ψ takes both values, so deleting $\mathcal{E}_0$ disconnects the lattice. This contradicts minimality of $\mathcal{E}$. Hence its edge adjacency graph is connected. The inside endpoints, and separately the outside endpoints, of two cut edges in the same unit square are at maximum-norm distance at most one. They give the asserted star-connected boundaries. Filling holes does not change these exterior boundaries: a bounded complementary component cannot have an edge to the unbounded one.

Outside a sufficiently large box all vertices are good and are nearest-neighbor connected, so the good vertices have exactly one infinite component $\mathcal{O}$. Its complement is finite. Let H be a nearest-neighbor component of that complement. Every vertex outside H adjacent to H is in $\mathcal{O}$; otherwise it would belong to the same complementary component. In particular all vertices of the inner exterior boundary of H are bad, since a good vertex adjacent to $\mathcal{O}$ would be in $\mathcal{O}$. That inner boundary is star-connected by the preceding argument, so it is contained in one bad component and has diameter at most L. Each coordinate maximum and minimum of H is attained on this boundary: from an extremal vertex the outward coordinate ray stays outside H after its first step. Thus H itself has maximum-norm diameter at most L. A neighboring exterior vertex in $\mathcal{O}$ is within $L+1$ of every point of H. This proves the claimed density of $\mathcal{O}$.

We next prove a local joining statement. If $u, v \in \mathcal{O}$ lie in a lattice cube $B_R(z)$, then they have a good path inside $B_{R+L+4}(z)$, with integer radii enlarged harmlessly if needed. Suppose not, and let S be the vertices reachable from u by good paths inside that larger cube. It is finite and nearest-neighbor connected. The vertex v belongs to the unbounded component of its complement: follow a path in $\mathcal{O}$ from v until it first exits the larger cube. Its initial part cannot meet S, and after the exit a path outside the cube goes to infinity without meeting S. A coordinate path from u to v inside $B_R(z)$ therefore meets the outer exterior boundary of S inside $B_R(z)$. That boundary also has a vertex outside the larger cube: follow a good path from u to infinity until its first exit. The outer boundary is star-connected. Every one of its vertices inside the larger cube is bad, since a good neighbor of S there would be reachable from u. A star path in this boundary from the inner cube to the exterior of the larger cube consequently contains a bad star-connected segment of diameter greater than L. This is a contradiction and proves the local statement.

Finally take a compact Euclidean path in V from u to v and partition it into finitely many pieces of diameter at most $L+1$. For each partition point choose a vertex of $\mathcal{O}$ within $C_d(L+1)$, using the density just proved and nearest-lattice rounding; keep u, v as the end representatives. Consecutive representatives lie in a cube of diameter $C_d(L+1)$ about the corresponding path piece. The local joining statement connects them in a fixed larger cube of that order. Concatenation stays in the $C'_d(L+1)$-neighborhood of V. If the original endpoints lie in a cube of diameter $C_0(L+1)$, the local statement directly gives a route in a cube of diameter $C_1(L+1)$. Erase loops from this finite lattice route; its number of edges is at most the number of vertices in the containing cube, which is $C_2(L+1)^d$. This proves all the assertions. □

The next construction is needed to connect different mesh levels. A good component at one mesh need not meet a good component at the next mesh. The additional labels below enforce a common infinite component.

Fix a bounded working region and a larger deterministic cube containing all the paths and buffers under consideration. At mesh $r_m = 2^{-m}$, declare $v \in r_m\mathbb{Z}^d$ good if, for every signed coordinate neighbor $v + r_m e$, its fine-field internal distance satisfies

$$D_{X^{m,n}}(v, v + r_m e; B_{2r_m}(v)) \le r_m b_{n-m}. \tag{3.3}$$

Only labels in the larger working cube are needed; declare vertices sufficiently far outside it good. The large cube may be increased by a fixed factor whenever a routing neighborhood is used. By stationarity, radiality, and (3.1), each label fails with probability at most $2d(1-q)$. The white-noise support of a label is within a fixed multiple of r_m of v, because its fine scales are at most r_m and its convolution radius is at most r_m. Thus the dependence range in mesh units is bounded independently of m, n.

For $m < n$, call $v \in r_m\mathbb{Z}^d$ *nice* if it is good at level m and every level-$(m+1)$ vertex within maximum-norm distance $3r_m$ of v is good at level $m+1$. These labels have uniformly bounded dependence range and failure probability at most $C_d(1-q)$. The infinite nearest-neighbor nice component is contained in the level-m infinite good component. After subdividing each of its edges into two edges, it is also contained in the level-$(m+1)$ infinite good component. Indeed, all the subdivision vertices and all the signed fine edges used in this subdivision are required to be good by the nice label. The same statement holds for the harmless modification outside the large working cube if the nice labels there are defined from the modified good labels.

Proposition 3.3 (Raw successive-mesh chaining)**.** *There are $q_0 < 1$ and constants C, c, A, depending only on the fixed geometries, kernels, and dimension, with the following property. Choose $q > q_0$*

and put

$$T_m(S) = \{x : \mathrm{dist}_\infty(x, S) \le A(1+m)^2 2^{-m}\}.$$

For every cutoff n and integer $M \le n$, there is an event $G_{M,n}$ such that

$$\mathbb{P}(G_{M,n}^c) \le C \sum_{m \ge M-1} 2^{dm} e^{-c(1+m)^2}. \tag{3.4}$$

On this event, simultaneously for $n \ge r \ge M$ and for points $x, y \in S$ with $|x-y|_\infty \le 2^{-r}$,

$$\begin{aligned} D_n(x, y; U) \le C \sum_{m=r}^{n-1} (1+m)^{2d} 2^{-m} b_{n-m} \exp\left(\xi \sup_{T_m(S)} H_m\right) \\ + C(1+n)^{2d} 2^{-n} \exp\left(\xi \sup_{T_n(S)} X_n\right), \end{aligned} \tag{3.5}$$

provided the displayed tubes and their fixed edge buffers lie in U and in the working region. If S is a cube boundary of diameter at most $C_0 2^{-r}$, the same right-hand side, with a changed constant, bounds $\sup_{x,y \in S} D_n(x, y; U)$.

Proof. Choose q so that Lemma 3.1 applies to both the good and the nice labels. At level m, there are $O(2^{dm})$ relevant vertices. Except with probability $C 2^{dm} e^{-c(1+m)^2}$, all bad star-components for both label systems have diameter at most $C(1+m)^2$ in mesh units. Intersect these events for $M - 1 \le m < n$; this proves (3.4).

Fix x. At each level $r \le m < n$, choose a vertex $v_m(x)$ in the infinite nice component within distance $C(1+m)^2 2^{-m}$ of x. Such vertices exist by Lemma 3.2. The two vertices $v_m(x)$ and $v_{m+1}(x)$ belong to the same level-$(m+1)$ infinite good component. The deterministic routing lemma joins them by at most $C(1+m)^{2d}$ fine edges inside a ball about x of radius $C(1+m)^2 2^{-m}$. This is the point at which the nice labels, rather than only the separate good components, are essential.

The level-r vertices associated with x and y lie in the same level-r infinite good component and at mutual distance at most $C(1+r)^2 2^{-r}$. Join them in the same way. At the finest level, join x directly to $v_{n-1}(x)$, and similarly at y. The direct segments have length at most $C(1+n)^2 2^{-n}$ and lie in $T_n(S)$ after increasing A by a fixed factor.

Each good edge at level m has fine-field cost at most $2^{-m} b_{n-m}$. Its chosen approximating path lies in its defining edge cube. On that cube the coarse contribution $H_m * \rho_{2^{-n}}$ is bounded by $\sup_{T_m(S)} H_m$, after enlarging A to include the support of the mollifier. Hence its full-field cost is at most the corresponding summand in (3.5), divided by the edge-count factor. Choose the finitely many edge paths with total excess as small as desired. Summing their costs, adding the two direct segments, and letting this excess decrease to zero proves the estimate. When $n = r$, use the direct segment from x to y.

Finally, any two points of a cube boundary of diameter $C_0 2^{-r}$ can be joined along that boundary by a bounded number of links of length at most 2^{-r}, with the bound depending only on C_0, d. Apply the preceding estimate to these links. Nothing in this argument requires a limiting metric to be finite on the whole boundary. □

Remark 3.4. The estimate is a prelimit statement. To turn its surface supremum into a summable bound by a direct Gaussian maximum argument would require $Q_p > \sqrt{2(d-1)}$. For smaller positive Q_p, it will be used only at finitely many deterministic vertices and on selected finite connections. The subsequent averaging argument, rather than a supremum over a hypersurface, supplies annular connectors.

4 The crossing exponent and point moments

Throughout this part, $d \geq 3$ and $\xi > 0$. Write $B_r(z) = z + (-r, r)^d$ and $B_r = B_r(0)$. The stationary scale field is

$$H_t(x) = \int_{2^{-t}}^{1} \int_{\mathbb{R}^d} u^{-(d+1)/2} K((x-y)/u)\, W(dy, du), \qquad t \geq 0,$$

where K is nonnegative, radial, smooth, compactly supported, and $\|K\|_2 = 1$. Thus $\operatorname{Var} H_t(x) = t \log 2$. The symbol X_t means either H_t or $H_{0,\infty} * \rho_{2^{-t}}$, with the choice kept fixed in an argument. Let D_t be the path metric with density $e^{\xi X_t}$. Distances restricted to a fixed cube can equivalently be computed in its closure when the endpoints are interior: inward polygonal approximation and short endpoint segments prove this directly at every fixed cutoff. We use this convention when compactness of the admissible domain is convenient.

Deleting the first k scales gives an exact independent decomposition

$$X_n = G_{k,n} + F_{k,n}, \qquad F_{k,n}(2^{-k}\,\cdot) \overset{\text{law}}{=} X_{n-k}, \quad 0 \leq k \leq n, \tag{4.1}$$

where $G_{k,n} = H_k$ in the sharp model and $G_{k,n} = H_k * \rho_{2^{-n}}$ in the convolution model. The point variance of $G_{k,n}$ is at most $k \log 2$. Fine events in sets separated by a fixed multiple of 2^{-k} depend on disjoint white noise. All constants below may depend on the fixed kernel, mollifier, dimension, parameter, geometric buffers, and explicitly fixed probability or moment orders. They never depend on the cutoff.

Fix a sufficiently high quantile level $q < 1$ and put

$$b_n = \inf\{t : \mathbb{P}[D_n(0, e_1; B_2) \leq t] \geq q\}. \tag{4.2}$$

We use the raw lattice chaining statement of Proposition 3.3, which requires no point-exponent limit. Consistently with (3.2), define

$$\gamma_p = \limsup_{n\to\infty} \frac{\log_2 b_n}{n}, \qquad Q_p = \frac{1+\gamma_p}{\xi}, \qquad \chi_p = 1 - \xi Q_p = -\gamma_p. \tag{4.3}$$

The elementary compact field bounds below show that γ_p is finite. Throughout we assume $Q_p > 0$. The nonasymptotic point-geometry comparison in Lemma 2.7 says that fixed point quantiles in a compact family of separated buffered geometries have logarithms within $O(\sqrt{n+1})$ of $\log b_n$. It assumes no exponent limit. We will prove both existence of that limit and bounded comparison with crossing normalizers.

4.1 An annular exponent

We first work with the sharp field. Let

$$T_n = D_n(\overline{B}_1, \partial B_2; \overline{B}_2), \qquad \widehat{c}_n = \operatorname{Med} T_n, \qquad a_n = \log \widehat{c}_n.$$

Here T_n is also the unrestricted distance from $\overline{B}_1$ to ∂B_2: a path can be stopped at its first outer-boundary hit.

Lemma 4.1 (Elementary field and distance bounds)**.** *For a fixed bounded set V, integers $0 \leq a < b$, and fixed $\varepsilon > 0$,*

$$\mathbb{P}\Big[\sup_V |H_b - H_a| > C_{V,\varepsilon}\sqrt{(b-a)(b+1)}\Big] < \varepsilon.$$

For $0 < \delta \le 2^{-m}$, the oscillation of H_m between points in V at distance at most δ is, with probability at least $1-\varepsilon$, at most $C_{V,\varepsilon} 2^m \delta \sqrt{\log(2+\delta^{-1})}$. The weaker bound with $\sqrt{2^m \delta \log(2+\delta^{-1})}$ is also valid and suffices below. If L_n is a distance between fixed disjoint compact sets, in a fixed connected buffered domain or without a domain restriction, then

$$\mathbb{P}\big[|\log L_n - \log \operatorname{Med} L_n| > t\big] \le C \exp\left(-\frac{t^2}{C(n+1)}\right). \tag{4.4}$$

In particular, fixed quantiles of the same distance have logarithms within $O(\sqrt{n})$ of each other. Moreover,

$$|a_n| \le Cn, \qquad |a_{j+k} - a_j| \le C\sqrt{k(j+k+1)} \quad (k \ge 1). \tag{4.5}$$

Proof. White-noise isometry and smoothness of K give

$$\mathbb{E}|(H_b - H_a)(x) - (H_b - H_a)(y)|^2 \le C \min\{2^{2b}|x-y|^2, b-a\}.$$

On a cube of side 2^{-b}, subtracting the center value gives a Gaussian process with uniformly bounded entropy integral and variance. Its absolute supremum has bounded mean and a uniform Gaussian tail. A cover by $O(2^{db})$ such cubes, followed by a union bound for their center values of variance $(b-a)\log 2$, proves the first assertion. The same argument on cubes of side δ gives the oscillation estimate. For the normalized ball-indicator kernel the corresponding increment bound is $C2^b|x-y|$; the weaker displayed oscillation estimate and all subsequent arguments remain valid.

A Cameron–Martin perturbation v changes $H_n(x)$, uniformly in x, by at most $\sqrt{n \log 2}\|v\|_2$. In the convolution model $\operatorname{Var} X_n(x) = n \log 2 + O(1)$, so the corresponding bound is $C\sqrt{n+1}\|v\|_2$. Every path length therefore changes by a factor between $e^{-C\sqrt{n+1}\|v\|_2}$ and $e^{C\sqrt{n+1}\|v\|_2}$. Taking infima preserves these bounds. Gaussian concentration about the mean, and the resulting bound of order $\sqrt{n+1}$ on the mean-to-median difference, give (4.4) after enlarging C. The logarithm is integrable: a fixed polygonal connection gives an upper bound using a compact field supremum, and a bounded neighborhood of the first endpoint set, disjoint from the second, gives a positive width times a compact field infimum as a lower bound. A countable polygonal approximation proves measurability.

Since a reference crossing has Euclidean length at least one and a straight crossing of length one exists, the first field bound gives $|a_n| \le Cn$. On a probability-0.99 event, adding the k scale band changes log crossing lengths by at most $C\sqrt{k(j+k+1)}$. Intersecting with the probability-0.99 concentration event for T_j proves the second bound in (4.5). □

Lemma 4.2 (Slow change under a short cutoff shift)**.** *For either cutoff convention and integers $0 \le j < n$,*

$$|\log b_n - \log b_j| \le C\sqrt{(n-j)(n+1)}. \tag{4.6}$$

The same estimate holds for any fixed quantile in a fixed buffered point geometry. Writing b_t for the analogous quantile at real depth t, we also have

$$\sup_{0 \le u \le 1} |\log b_{n+u} - \log b_n| \le C\sqrt{n+1}. \tag{4.7}$$

In particular, if $j_n \le n$ and $n - j_n = o(n)$, then $\log b_n - \log b_{j_n} = o(n)$.

Proof. Put $k = n - j \geq 1$ and couple the two cutoffs using the same noise. For the sharp field the first estimate of Lemma 4.1 gives, on every fixed compact V,

$$\mathbb{P}\left[\sup_V |X_n - X_j| > C_{V,\varepsilon}\sqrt{k(n+1)}\right] < \varepsilon.$$

Here is the same bound for convolution cutoffs. Write $\epsilon_i = 2^{-i}$. In the white-noise isometry for $X_n - X_j$, scales $u \geq \epsilon_j$ contribute at most $C\int_{\epsilon_j}^1 (\epsilon_j/u)^2\, du/u \leq C$: smoothness of the kernel bounds the difference of its two averages. Scales $\epsilon_n \leq u \leq \epsilon_j$ contribute at most Ck. For $u \leq \epsilon_n$, the convolution estimate

$$\|K_u * \rho_{\epsilon_i}\|_2^2 \leq C(u/\epsilon_i)^d, \qquad K_u(x) = u^{-d/2}K(x/u),$$

bounds the remaining integral by a constant. Thus the point variance of the difference is at most Ck. Differentiating the kernel above its cutoff and the mollifier below it also gives

$$\mathbb{E}|(X_n - X_j)(x) - (X_n - X_j)(y)|^2 \leq C\min\{2^{2n}|x-y|^2, k\}.$$

On each cutoff cube, the centered increment process therefore has uniformly bounded Gaussian supremum moments. Cover V by $O(2^{dn})$ such cubes, and union-bound their center values and increment suprema. This proves the displayed probability bound for convolution cutoffs as well.

On this event every path in the reference buffer has its length changed by a factor between $e^{-C\sqrt{k(n+1)}}$ and $e^{C\sqrt{k(n+1)}}$. Choose $2\varepsilon < \min\{q, 1-q\}$. By (4.4), the old log distance lies within $C\sqrt{j+1}$ of $\log b_j$ with probability at least $1-\varepsilon$. On the intersection, of probability at least $1-2\varepsilon$, the new log distance lies within $C\sqrt{k(n+1)} + C\sqrt{j+1}$ of $\log b_j$. Its lower q quantile lies in this same interval: its probabilities below and above the interval are each at most 2ε. Since $k \geq 1$, the last term is absorbed in $C\sqrt{k(n+1)}$, proving (4.6). The case $j_n = n$ is immediate, and all other cases of the last assertion follow by division by n. For cutoffs n and $n+u$, $0 \leq u \leq 1$, the same kernel estimates give uniformly bounded point variance for their difference and increments at scale $2^{-(n+1)}$. The identical covering and quantile argument proves (4.7), uniformly in u. □

Lemma 4.3 (Erasing bad clusters)**.** *Let w, v be positive continuous densities in a convex compact domain Ω. On a mesh of size r, declare vertices good or bad, using adjacency at sup-norm distance at most $3r$. Suppose each bad component has at most L vertices, every good vertex z has*

$$D_v(\overline{B}_r(z), \partial B_{2r}(z)) \geq ra,$$

and the oscillation of $\log w$ between z and points within $8(L+1)r$ is at most η. Then, for $|x-y|_\infty > 6(L+1)r$,

$$D_{wv}(x, y; \Omega) \geq \frac{ae^{-2\eta}}{12\sqrt{d}(L+1)} D_w(x, y; \Omega). \tag{4.8}$$

The same conclusion, with changed fixed constants, holds for any fixed shell template and any fixed mesh-to-shell size ratio.

Proof. For a rectifiable path, round its current position to a nearest mesh vertex and stop the next segment at its first exit from the vertex's $2r$ cube. Consecutive vertices are adjacent, and every completed segment has Euclidean length at least $3r/2$. A completed segment at a good vertex costs at least $rae^{-\eta}w(z)$. A consecutive run of bad vertices stays in one bad component, irrespective of repeated visits. Erase these runs and join the remaining good segment starts, together with the original endpoints, by straight chords. Consecutive retained points have distance at most $6(L+1)r$;

at least one good segment exists by the endpoint separation. Convexity keeps the chords in Ω. Charge each outgoing chord to its good starting vertex, and the initial chord also to the first good vertex. The resulting w length is at most $12\sqrt{d}(L+1)re^{\eta}\sum_z w(z)$, whereas the original disjoint time segments cost at least $rae^{-\eta}\sum_z w(z)$. Their spatial images need not be disjoint. Taking infima gives the result. Changing the fixed shell template changes only the constants in this stopping construction. □

Proposition 4.4 (Existence of the annular exponent)**.** *There is a finite deterministic Q_a such that every fixed quantile of a fixed-ratio annular crossing has the form $2^{-(1-\xi Q_a)n+o(n)}$. The same exponent applies to the unrestricted distance between any two disjoint compact sets with nonempty interiors. It is unchanged by the choice between the sharp and convolution cutoffs. Furthermore, $Q_a \le Q_p$.*

Proof. We prove almost superadditivity of a_n. Let $N = m + n$, with $m \le n$, and choose

$$k = k_0 + 8\lceil \log_2(N+2)\rceil, \qquad j = m + k, \qquad \ell = n - k, \qquad r = 2^{-j}.$$

For large N, $\ell \ge 1$. Split $H_N = H_m + (H_j - H_m) + (H_N - H_j)$. A fine mesh vertex is good when its fine shell crossing is at least $r\widehat{c}_\ell e^{-C\sqrt{\ell}}$. By (4.4), C makes the bad probability as small as desired. There is a fixed finite partition of the mesh into classes within which the labels are independent. A connected vertex set of size s contains at least a fixed fraction of mutually independent vertices; the number of such sets rooted at one vertex is at most C_d^s. Thus, for sufficiently small bad probability,

$$\mathbb{P}[\text{some bad component in the working region has size} > L] \le C2^{dj}e^{-cL}.$$

Take $L = C_0(j+1)$ with C_0 large. With probability at least 0.99 all bad components satisfy this bound. The oscillation of H_m on distances $8(L+1)r$ is at most one with probability at least 0.99, because $2^m(L+1)r \le C(N+1)2^{-k}$. The intermediate band has supremum at most $C\sqrt{k(j+1)}$, and T_m is at least $\widehat{c}_m e^{-C\sqrt{m}}$, each with probability at least 0.99.

Apply Lemma 4.3 to every reference crossing, with $w = e^{\xi H_m}$ and $v = e^{\xi(H_N - H_j)}$, then restore the intermediate band. On the intersection of these four events,

$$T_N \ge \widehat{c}_m\widehat{c}_\ell \exp[-C(\sqrt{m} + \sqrt{\ell} + \sqrt{k(j+1)} + \log(L+1) + 1)].$$

Its probability exceeds $1/2$, so the same deterministic lower bound holds for the median. Equation (4.5) replaces a_ℓ by a_n at a cost $O(\sqrt{N\log(N+2)})$. Hence

$$a_{m+n} \ge a_m + a_n - C\sqrt{(m+n)\log(m+n+2)} \ge a_m + a_n - C'(m+n)^{2/3}. \tag{4.9}$$

The finitely many remaining indices are absorbed into C'.

For completeness, if $|u_n| \le Cn$ and $u_{m+n} \ge u_m + u_n - C(m+n)^\alpha$ with $\alpha < 1$, repeated doubling gives $u_{2^j r} \ge 2^j u_r - C_1 2^j r^\alpha$. Writing an integer q in binary and merging its dyadic blocks yields

$$u_{qr} \ge qu_r - C_1 qr^\alpha - C_2(qr)^\alpha \log(q+2).$$

Adding a remainder less than r, dividing by the total length, and letting the total tend to infinity shows $\liminf_n u_n/n \ge u_r/r - C_1 r^{\alpha-1}$. Let r run through a subsequence realizing the limsup. This proves existence of the limit. Apply it to (4.9) and write the limit as $-(1-\xi Q_a)\log 2$.

A fixed similarity changes logarithms of bounded-domain distances by $O_{\mathbb{P}}(\sqrt{n})$. Indeed, scaling by 2^c changes the white-noise interval $[2^{-n}, 1]$ to $[2^{-(n+c)}, 2^{-c}]$. Couple overlapping bands identically.

The difference consists of a fixed-width coarse band and a fixed-width band at the cutoff. The latter has bounded point variance; covering a fixed compact region by $O(2^{dn})$ cutoff cubes bounds its supremum by $O_{\mathbb{P}}(\sqrt{n})$. Fixed real c is handled by the same calculation. Rotations and translations preserve the law. Gaussian log-distance concentration then transfers medians and fixed quantiles. This proves the assertion for similar annuli.

Let K_1, K_2 be disjoint compact sets with nonempty interiors. Cover K_1 by finitely many small cubes whose doubled cubes avoid K_2. Every connecting path crosses one of their shells. Similarity comparison and a union bound give a lower median bound $\widehat{c}_n e^{-C\sqrt{n}}$. For the upper bound, cover the two boundaries of the reference annulus by finitely many small balls. Some pair of these balls has an internal connection of cost at most $\widehat{c}_n$ with a fixed positive probability. Map its centers by a similarity into fixed interior balls of K_1, K_2, choosing the cover small enough that the whole endpoint balls map inside them. Only finitely many similarities occur. Similarity comparison and (4.4) therefore give the upper median bound $\widehat{c}_n e^{C\sqrt{n}}$. The unrestricted target distance is no larger than the mapped bounded-domain distance, as required.

The common-field coupling in the field preliminaries has compact supremum error $O_{\mathbb{P}}(\sqrt{n})$ between a sharp approximation and a mollified approximation. Applying it only to bounded crossing geometries transfers their leading exponents. The preceding shape argument then applies to the convolution model as well. Finally, an annular crossing is bounded above by a fixed finite concatenation of reference point paths on a fixed positive-probability event. Concentration therefore bounds its median above by $Jb_n e^{C\sqrt{n+1}}$. Divide its logarithm by n and take the limsup to obtain $Q_a \le Q_p$. □

4.2 Small point moments from the upper point growth rate

The following deduction is useful because it does not assume any uniform ratio estimate for b_n.

Lemma 4.5 (Point moments before exponent identification)**.** *For each buffered deterministic point geometry and $0 < p < 2Q_p/\xi$, for every $\delta > 0$ there is $C_{p,\delta} < \infty$ such that*

$$\|D_n(x, y; U)\|_p \le C_{p,\delta} 2^{-(\chi_p - \delta)n}. \tag{4.10}$$

The bound is uniform over compact families of such geometries. Let A_n be the largest p norm in a fixed finite collection of local point geometries containing the reference one. For $\eta_p > 0$, set

$$\mathcal{B}_n = \max_{0 \le j \le n} A_j 2^{-(\chi_p + \eta_p)(n-j)}. \tag{4.11}$$

Then

$$\limsup_n \frac{\log_2 A_n}{n} = \limsup_n \frac{\log_2 \mathcal{B}_n}{n} = -\chi_p, \qquad \mathcal{B}_n \ge A_n \ge (1-q)^{1/p} b_n,$$

and

$$\mathcal{B}_{n-k}/\mathcal{B}_n \le 2^{(1-\xi Q_p + \eta_p)k}.$$

For every $\varepsilon > 0$, if η_p is sufficiently small, there are events whose probabilities tend to one as $m \to \infty$, uniformly in $n \ge m$, on which simultaneously for all nearby x, y in a fixed working region,

$$\frac{D_n(x, y; U)}{\mathcal{B}_n} \le C_\varepsilon \sum_{k=m}^{n} 2^{-(\xi Q_p - \varepsilon)k} \exp\left(\xi \sup_{T_k(x) \cup T_k(y)} H_k \right). \tag{4.12}$$

Here $|x - y| \le C2^{-m}$, and $T_k(x)$ has radius $C(1+k)^C 2^{-k}$. In the convolution model the terminal summand includes its bounded-variance cutoff residual, with the same estimates.

Proof. Fix $\delta > 0$. The definition by a limsup gives $b_j \le C_\delta 2^{-(1-\xi Q_p-\delta)j}$ for every j. Insert this into the raw successive-mesh chain and divide by $2^{-(1-\xi Q_p-\delta)n}$. The endpoint contribution is bounded by

$$C_\delta \sum_{k\ge 0} (1+k)^C 2^{-(\xi Q_p+\delta)k} e^{\xi \sup_{T_k(x)} H_k}.$$

For a deterministic x and fixed $u > 0$,

$$\mathbb{E} e^{u \sup_{T_k(x)} H_k} \le C_u (1+k)^{C_u} 2^{u^2 k/2}. \tag{4.13}$$

To see this, tilt by $uH_k(x)$. The mean added to each increment $H_k(y) - H_k(x)$ is nonpositive, because stationarity gives

$$\operatorname{Cov}(H_k(y) - H_k(x), H_k(x)) = -\tfrac{1}{2} \operatorname{Var}(H_k(y) - H_k(x)).$$

The exponential moments of the increment supremum on a polynomial number of cutoff cells are polynomial in k. This proves (4.13). Subadditivity for $p \le 1$, or Minkowski's inequality for $p > 1$, now proves summability whenever $p\xi Q_p - p^2\xi^2/2 > 0$.

The first depth M from which all the required good/nice mesh and oscillation events hold has a tail $\mathbb{P}[M > j] \le C \exp[-cj^2/\log(j+2)]$. The connected level-$M$ core has at most $C2^{dM}$ vertices and total normalized mass at most $e^{CM} e^{\xi \sup_U H_M}$. Its fixed moments grow at most exponentially in M; Cauchy–Schwarz and the displayed tail sum them without an independence assumption. If $M > n$, a fixed polygonal path has normalized moments growing only exponentially in n, again absorbed by the exceptional-event tail. The terminal cutoff segment is bounded by the same tilted Gaussian estimate, using $\operatorname{Var} X_n(x) = n \log 2 + O(1)$ directly. This proves (4.10). Markov's inequality for the reference distance gives $A_n \ge (1-q)^{1/p} b_n$, so the upper bound and (4.3) identify the limsup of $\log_2 A_n/n$.

For every $\delta > 0$, the same upper bound and the envelope definition give $\mathcal{B}_n \le C_{p,\delta} 2^{-(\chi_p-\delta)n}$, whereas $\mathcal{B}_n \ge A_n$ gives the matching limsup. The ratio assertion follows directly from the maximum definition. Markov's inequality bounds the needed high point quantiles by fixed multiples of A_j. Use these thresholds in the raw chain, apply the envelope ratio, and absorb polynomial factors into $2^{\varepsilon k}$. This gives (4.12). The simultaneous event is the same good/nice event as in raw chaining, whose failure probabilities are summable over $k \ge m$. □

5 Lower crossing bounds and averaged barriers

Fix a sufficiently large integer k_0, put $R = 4 \cdot 2^{k_0}$, and let c_n be a sufficiently low fixed p_0 quantile of the crossing of $B_{2R} \setminus \overline{B}_R$. The choices of k_0, p_0, q are made once: p_0 is small enough for every finite-range bad-site estimate below, and q is high enough for raw point chaining and for a finite point-path concatenation across this annulus. In particular,

$$c_n \le J b_n \tag{5.1}$$

for a fixed J: require the J point paths of a fixed concatenation to cost at most b_n, an event of probability at least $1 - J(1-q) > p_0$. For $\eta > 0$, define

$$C_n = \min_{0\le j\le n} c_j 2^{-(1-\xi Q_a-\eta)(n-j)}. \tag{5.2}$$

It has exponent Q_a, satisfies

$$C_{n-k}/C_n \ge 2^{(1-\xi Q_a-\eta)k}, \tag{5.3}$$

and has infinitely many contact indices $C_n = c_n$. Indeed, $c_n 2^{(1-\xi Q_a-\eta)n} = 2^{-\eta n + o(n)} \to 0$.

Lemma 5.1 (Contacts realizing the upper point growth)**.** *For each n, choose any index $j(n)$ attaining the minimum in* (5.2)*. Then $C_{j(n)} = c_{j(n)}$ and $n - j(n) = o(n)$. There is an unbounded sequence of contact indices, along which*

$$\log_2 b_n = -\chi_p n + o(n), \qquad \log_2 \mathcal{B}_n = -\chi_p n + o(n), \qquad C_n = c_n. \tag{5.4}$$

Proof. Put $\chi_a = 1 - \xi Q_a$ and $e_i = \log_2 c_i + \chi_a i = o(i)$. If $C_j < c_j$ for $j = j(n)$, an earlier minimizer for C_j, extended to index n, would strictly improve the minimum defining C_n. Hence j is a contact. Comparing its value with the candidate n gives

$$\eta(n - j) \le e_n - e_j.$$

The estimate $e_i = o(i)$ implies $\max_{0\le i\le n} |e_i| = o(n)$, by separating finitely many initial indices from the remaining ones. Thus $n - j = o(n)$.

Choose a sequence n_k realizing the limsup in (4.3), and put $j_k = j(n_k)$. Lemma 4.2 gives $\log b_{n_k} - \log b_{j_k} = o(n_k)$. Since $j_k/n_k \to 1$,

$$\log_2 b_{j_k} = -\chi_p j_k + o(j_k).$$

The lower bound $\mathcal{B}_{j_k} \ge (1 - q)^{1/p} b_{j_k}$ and the global upper-growth bound of Lemma 4.5 give the same asymptotic for $\mathcal{B}_{j_k}$. Extracting a strictly increasing subsequence of these contact indices proves (5.4). □

Lemma 5.2 (Inverse crossing tails)**.** *For every fixed nondegenerate annulus A,*

$$\mathbb{P}[D_n(\textit{across } A) < e^{-t} C_n] \le C_A \exp\left[-\frac{c_A t^2}{\log(t + 2)}\right] \tag{5.5}$$

for sufficiently large t, uniformly in n. Consequently all fixed inverse crossing moments in C_n units are uniformly bounded. For every $\delta > 0$, choosing η small enough gives events $\mathcal{G}_{m,n}$ with $\inf_{n\ge m} \mathbb{P}[\mathcal{G}_{m,n}] \to 1$ as $m \to \infty$, on which, simultaneously at all locations in a fixed compact region and all $m \le k \le n$, every completed scale-2^{-k} shell costs at least

$$C_n 2^{-(\xi Q_a + \delta)k} e^{\xi H_k(x)}. \tag{5.6}$$

Fixed geometric multiplicative constants may be absorbed into a slightly larger δ. On the same events the oscillation of H_k on every polynomially enlarged cutoff ball is at most δk.

Proof. Choose $\ell = \lceil 2 \log_2(t + 2) \rceil + C_A$. If $\ell \le n$, partition at scale $2^{-\ell}$ and declare a fine shell bad if its crossing is less than $2^{-\ell} c_{n-\ell}$. Its probability is at most p_0, and the dependence range is fixed. A self-avoiding bad-path count bounds the probability of a macroscopic spanning bad chain by $C 2^{d\ell} e^{-c2^{\ell}}$. Otherwise every crossing pays for a good shell and therefore has cost at least

$$2^{-\ell} c_{n-\ell} e^{\xi \inf_V G_{\ell,n}} \ge C_n 2^{-(\xi Q_a + \eta)\ell} e^{\xi \inf_V G_{\ell,n}}.$$

The coarse infimum has mean at least $-C(\ell + 1)$ and Gaussian variance parameter $C(\ell + 1)$. Since $\ell = O(\log(t + 2))$, its lower tail proves (5.5). If $\ell > n$, use the deterministic crossing width times $e^{\xi \inf_V X_n}$, $C_n \le c_0 2^{-(1-\xi Q_a - \eta)n}$, and $n = O(\log(t + 2))$.

For the simultaneous statement, round shell centers to a scale-2^{-k} mesh and require every relevant fine crossing to exceed $2^{-k} C_{n-k} 2^{-\delta' k}$. The preceding tail and a union bound over $O(2^{dk})$ centers give failure probability at most $C 2^{dk} e^{-ck^2/\log(k+2)}$. Coarse oscillations on a polynomial

cutoff neighborhood have increment variance $O(\log(k+2))$ and entropy giving the same summable tail for exceeding $\delta' k$. Restore the coarse factor and apply (5.3). Choosing η, δ' small in terms of δ gives (5.6). At the terminal scale use the density minimum and the cutoff oscillation estimate; the same envelope inequality with $n - k = 0$ applies. Summing over $k \geq m$ proves the assertion, including moving centers and field-selected subpaths. □

Lemma 5.3 (Weighted short arms)**.** *Fix two separated coordinate slabs traversed by a macroscopic crossing. Write $r = 2^{-m}$. Given $\zeta > 0$, with probability tending to one as $m \to \infty$, uniformly in $n \geq m$, every crossing path of cost at most LC_n has, in each slab, a subarm $A \subset B_{Cr}(z)$ with coordinate displacement at least cr and*

$$\mathrm{len}_{D_n}(A) \leq LC_n r^{\xi Q_a} e^{\xi H_m(z)} e^{\zeta m}. \tag{5.7}$$

The centers belong to a deterministic mesh, with $O(r^{-d})$ possibilities in each slab.

Proof. Use the last entrance to the lower slab face before the first hit of the upper one. Divide this slab path into first exits from balls of radius r, discarding its final incomplete segment. The polygonal chords of the full segments still traverse a smaller fixed slab. Round their starting points to mesh centers z_i and set $w_i = r^{\xi Q_a} e^{\xi H_m(z_i)}$. If O_m is the largest coarse oscillation on the relevant Cr balls, each chord has H_m length at most $Cre^{\xi H_m(z_i) + \xi O_m}$. Hence

$$\sum_i w_i \geq c r^{\xi Q_a - 1} e^{-\xi O_m} D_m^{\mathrm{sharp}}(\text{across the smaller slab}).$$

The slab crossing is bounded below by the minimum of finitely many fixed annular crossings. Lemma 5.2 and the annular exponent show that this last display is at least $e^{-\zeta m}$ with probability tending to one. We used also $O_m = o_{\mathbb{P}}(m)$, simultaneously over all centers. Since the old arm costs have total at most LC_n, one arm has cost at most $LC_n w_i / \sum_j w_j$. Its Euclidean displacement has a coordinate of magnitude at least $r/\sqrt{d}$. This proves the result in each slab separately. □

Fix from now on

$$\max\{0, Q_a - 1/\xi\} < \alpha < Q_p. \tag{5.8}$$

The interval is nonempty because $Q_p > 0$ and $Q_a \leq Q_p$. For a target of radius comparable to $r = 2^{-m}$ about z, define

$$Z_{m,n,z}(x) = \max_{0 \leq j \leq n-m} (H_{m+j}(x) - H_m(z) - \alpha j \log 2)_+. \tag{5.9}$$

Lemma 5.4 (Averaged barrier on a cheap arm)**.** *For every $\varepsilon > 0$, the slacks in the preceding lemmas can be chosen so that the arms of Lemma 5.3 satisfy*

$$\frac{1}{|I|} \int_I Z_{m,n,z}(x(t))\, dt \leq \varepsilon m \tag{5.10}$$

with probability tending to one, uniformly in $n \geq m + C_0$. Here I is a deterministic-grid coordinate interval of length comparable to r crossed by the arm, with a fixed buffer from its coordinate extremes, and $x(t)$ is any measurable section of the arm.

Proof. At relative depth $j \geq C_0$, an arm completes a shell of radius comparable to $2^{-(m+j)}$ in at least one direction from each of its section points. Partition I into intervals of this length. In each interval containing a level with

$$H_{m+j}(x(t)) - H_m(z) > \alpha j \log 2 + u$$

select one such level. A bounded coloring separates their supporting shells spatially, so their old path costs can be added. The shell charge (5.6) gives the following upper bound on the fraction of these levels:

$$C\frac{\mathrm{len}_{D_n}(A)}{C_n r^{\xi Q_a} e^{\xi H_m(z)}} 2^{\delta m} 2^{-(1+\xi\alpha-\xi Q_a-\delta)j} e^{-\xi u}. \tag{5.11}$$

The factor 2^{-j} is the shell width divided by $|I|$. Choose δ so that $1+\xi\alpha-\xi Q_a-\delta>0$. Substitute (5.7) and sum over j. The fraction of levels where the supremum over these j exceeds u is at most

$$\min\{1, CL\exp[(\zeta+\delta\log 2)m-\xi u]\}.$$

Integrating this tail bounds the average by $(\zeta+\delta\log 2)m/\xi+O_L(1)$. The finitely many omitted initial relative scales have maximal absolute increment $O_{\mathbb{P}}(\sqrt{m})$, by a union bound over the $O(2^{dm})$ target cells. Make ζ,δ small and then m large. No assumption about independence or typicality of the selected center value is used: its coarse weight cancels in (5.11). □

Lemma 5.5 (Tame points attach to deterministic centers)**.** *For every sufficiently small $\varepsilon_1>0$, choose the point envelope slack small enough. For every deterministic center z in a compact working region, with probability tending to one as $m\to\infty$, uniformly in $n\ge m$, simultaneously for every $x\in B_{C2^{-m}}(z)$ and $u\ge 0$,*

$$Z_{m,n,z}(x)\le u \quad\Longrightarrow\quad D_n(x,z;U)\le \mathcal{B}_n\exp[-(\xi Q_p\log 2-\varepsilon_1)m+\xi u]. \tag{5.12}$$

The probability bound is uniform over deterministic, including cutoff-dependent, choices of z. It is not a simultaneous typical-center assertion over all possible centers.

Proof. For a deterministic center, $H_m(z)=o_{\mathbb{P}}(m)$. Its subsequent scale process is a Brownian motion sampled at times $j\log 2$, so, since $\alpha>0$,

$$\sup_{j\ge 0}(H_{m+j}(z)-H_m(z)-\alpha j\log 2)=o_{\mathbb{P}}(m).$$

Impose these events, the simultaneous raw chain event, and the polynomial-neighborhood oscillation events before choosing x. Insert its barrier and the analogous center barrier into (4.12). The relative-scale series is geometric with positive decay exponent $\xi(Q_p-\alpha)$ after all slacks are made small. Its initial factor is $\mathcal{B}_n 2^{-\xi Q_p m}e^{\xi H_m(z)+\xi u+o(m)}$. Absorb the center and tube errors into $e^{\varepsilon_1 m}$. □

6 One marked-network functional

Take two fixed target cubes W_i of diameter comparable to 2^{-m}, with deterministic centers z_i, in a fixed bounded convex ambient cube U. In each target fix a coordinate and an interval I_i of length comparable to 2^{-m} in the middle quarter of that coordinate projection. A marked network consists of a finite connected polygonal graph $G\subset\overline{U}$ and connected subgraphs $G_i\subset G\cap W_i$ whose projections contain I_i. Its mass is weighted Hausdorff length, counting overlapping edges only once. Put

$$U_i(G_i)=\frac{1}{|I_i|}\int_{I_i}\min_{x\in G_i:\,x_{j_i}=t} Z_{m,n,z_i}(x)\,dt, \tag{6.1}$$

$$F_{n,m,\lambda}=\inf_{(G,G_0,G_1)}\{\,\mathrm{mass}_{D_n}(G)+C_n e^{\lambda U_0(G_0)}+C_n e^{\lambda U_1(G_1)}\},\qquad \lambda\ge 1. \tag{6.2}$$

The infimum can be restricted to graphs in a fixed countable polygonal family. In particular $F_{n,m,\lambda}\ge 2C_n$.

Lemma 6.1 (Measurability and bounded penalty energy)**.** *The functional* (6.2) *is measurable. Finite rectifiable marked networks are competitors by approximation, with no increase in the sum of their constituent edge masses or section penalties in the limit. For each fixed marked graph, the Cameron–Martin Lipschitz constant of U_i is bounded independently of m, n. Consequently* $\log F$ *has squared full-field Lipschitz constant at most $C(n + 1 + \lambda^2)$. Conditionally on the fine noise below 2^{-K}, $K \geq m$, its squared coarse Lipschitz constant is at most $C(K + 1 + \lambda^2)$.*

Proof. For a fixed finite polygonal graph, away from finitely many vertex levels its sections are finite collections of affine functions of the level. Thus their barrier minimum is Borel, and the countable infimum is measurable. We give the approximation argument without an area formula. If $q : [0, 1] \to \mathbb{R}$ is continuous, its local maximum and minimum *values* form a countable set: every such value is the maximum or minimum on some closed interval with rational endpoints contained in its defining neighborhood. Include the two endpoint values in this exceptional set. For every other value t and every s with $q(s) = t$, each neighborhood of s contains parameters where $q < t$ and where $q > t$. If $q_k \to q$ uniformly, the intermediate value theorem therefore supplies $s_k \to s$ with $q_k(s_k) = t$.

Parameterize each of the finitely many constituent edges of a rectifiable marked network continuously and approximate it by polygonal interpolation. Apply the preceding observation to its marked coordinate function. Outside one countable set of levels, every old section point is a limit of points at that same level on the approximating edge. Thus the upper limit of each new section minimum is no greater than the corresponding old minimum. The barrier is continuous and bounded on the fixed compact working region. Reverse Fatou, applied after subtracting from a common bound, gives the upper limit inequality for its section integral.

For clarity, polygonal interpolation also controls weighted length. On a parameter subinterval whose curve image has diameter at most δ, the chord has Euclidean length at most the old edge length there, and the continuous density differs from its value at the first endpoint by at most its modulus of continuity at 2δ. Summing and then taking $\delta \downarrow 0$ bounds the upper limit of new weighted length by the old weighted length. Apply this to each edge and add the bounds. Shared vertices are approximated identically; overlapping new edges may be merged, which only reduces mass. The networks used here are supplied with a finite edge presentation, and the cost estimates charge the sum of the presented edge lengths, counting a shared edge only once in that presentation. Consequently this sum is the mass bound needed both for original path competitors and for the surgery networks below.

An arbitrarily small inward perturbation and short endpoint segments preserve the buffered projection intervals. Subsequent rational perturbations preserve connectivity and the same bounds. This proves that a fixed countable polygonal family suffices.

Let $g_{k,x}$ be the white-noise kernel of $H_k(x)$. Uniformly in $k, \ell \geq m$ and $x, y \in W_i$,

$$|\langle g_{k,x} - g_{m,z_i}, g_{\ell,y} - g_{m,z_i} \rangle| \leq C + C \log_+ \frac{2^{-m}}{|x - y|}. \tag{6.3}$$

Indeed, above scale 2^{-m} the kernel differences have bounded norm by smoothness of K. The remaining cross-covariance is nonnegative, and compact support makes its integrand zero below a constant multiple of $|x - y|$; integrating du/u proves the bound. For sections $x(t)$ and $x(s)$, their coordinate separation implies $|x(t) - x(s)| \geq |t - s|$. Averaging arbitrary signed maximizing kernel choices therefore has squared norm at most

$$\frac{C}{|I_i|^2} \int_{I_i} \int_{I_i} \left(1 + \log_+ \frac{2^{-m}}{|t - s|} \right) dt\, ds \leq C.$$

The difference of two section minima is bounded using a minimizing section for one of them; the difference of the maxima in the barrier is bounded using maximizing scales. The displayed uniform bound proves both directions of the Lipschitz assertion, without differentiability or unique minimizers.

A fixed network mass changes under a coarse perturbation by a factor between $e^{-\xi\sqrt{K\log 2}\|v\|}$ and $e^{\xi\sqrt{K\log 2}\|v\|}$. A penalty changes by a factor between $e^{-C\lambda\|v\|}$ and $e^{C\lambda\|v\|}$. Positive sums and infima preserve the common multiplicative bounds. Restriction of a kernel to coarse noise cannot increase its norm. The full-field argument uses $\operatorname{Var} X_n = n\log 2 + O(1)$. □

Proposition 6.2 (Existence and equality of point and crossing exponents)**.** *If the upper point exponent Q_p in* (4.3) *is positive, then $Q_a = Q_p$ and the full point-exponent limit exists:*

$$b_n = 2^{-(1-\xi Q_p)n+o(n)}. \tag{6.4}$$

Proof. Suppose $g = \xi(Q_p - Q_a)\log 2 > 0$. Let n tend to infinity through the contact sequence of Lemma 5.1. A reference crossing of cost at most $2C_n$ exists with a fixed positive probability. A fixed finite choice of slabs contains two separated subarms. Apply Lemmas 5.3 and 5.4. There are at most $C2^{2dm}$ pairs of target centers, and only finitely many coordinate and interval choices for each pair. Pigeonholing gives deterministic target data, depending on m, n, with centers uniformly separated, for which

$$\mathbb{P}[F_{n,m,\lambda} \le 4C_n e^{\lambda\varepsilon m}] \ge e^{-c_0 m} \tag{6.5}$$

for all sufficiently large m. The chosen path and its arms are legitimate competitors by Lemma 6.1.

Gaussian concentration and (6.5) imply

$$\mathbb{E}\log(F/C_n) \le \lambda\varepsilon m + C\sqrt{(n+1+\lambda^2)(m+1)} + C.$$

Choose a small $\tau > 0$ with $2\xi\tau < \xi Q_p \log 2/4$, then $\varepsilon_1 < \xi Q_p \log 2/4$ and ε so small that $\varepsilon < \min\{\tau/4, g/8\}$. Fix all previous slacks with these choices. Finally choose a fixed λ sufficiently large and set $m = \lfloor n/\lambda \rfloor$. An upper Gaussian deviation of size $n^{3/4}$ gives, with probability tending to one,

$$\log(F/C_n) \le (\varepsilon + C/\sqrt{\lambda} + o(1))n.$$

The fixed choice of λ makes this coefficient smaller than both $g/2$ and $\tau/2$. Since $\log(\mathcal{B}_n/C_n) = gn + o(n)$,

$$F/\mathcal{B}_n \le e^{-gn/2}, \qquad \lambda^{-1}\log(F/C_n) \le \tau m$$

after slightly tightening the margins. Choose a measurable near-minimizer with objective at most $F + C_n$. Each marked section has average barrier at most $\tau m + o(m)$, so each has a point x_i with barrier at most $2\tau m$. By Lemma 5.5, these points attach to the two deterministic centers at costs at most $\mathcal{B}_n e^{-cm}$. The connected network contains a path between them of mass at most $F + C_n$. Hence

$$D_n(z_0, z_1; U)/\mathcal{B}_n \le e^{-c'n}$$

with probability tending to one, for a fixed $c' > 0$. The deterministic centers stay in a fixed compact family and have uniform positive separation. By the nonasymptotic point-geometry comparison, their log medians differ from $\log b_n$ by $O(\sqrt{n+1})$, uniformly in the chosen centers. Equation (5.4) makes those medians $\mathcal{B}_n e^{o(n)}$, contradicting the displayed high-probability exponential improvement. Thus $g = 0$.

Finally (5.1) gives

$$\liminf_{n\to\infty} \frac{\log_2 b_n}{n} \geq \lim_{n\to\infty} \frac{\log_2 c_n}{n} = -(1-\xi Q_a) = -(1-\xi Q_p).$$

The reverse limsup equality is the definition of Q_p. This proves (6.4). Equation (4.7) extends the limit to real cutoffs, uniformly over their fractional parts. □

Corollary 6.3 (Full point moment exponents)**.** *Write $Q = Q_a = Q_p > 0$. For every fixed buffered deterministic point geometry and $0 < p < 2Q/\xi$,*

$$\|D_n(x, y; U)\|_p = 2^{-(1-\xi Q)n+o(n)}. \tag{6.6}$$

The upper exponent is uniform over compact families of these geometries. The point-only envelope in (4.11) *also satisfies $\mathcal{B}_n = 2^{-(1-\xi Q)n+o(n)}$.*

Proof. The upper bound is Lemma 4.5. A fixed point quantile bounds its p norm below by a positive deterministic multiple of that quantile. The nonasymptotic geometry comparison and (6.4) give the matching lower exponent. For $\mathcal{B}_n$, its global upper-growth bound and $\mathcal{B}_n \geq (1-q)^{1/p} b_n$ give the two bounds directly. □

7 Uniform envelopes and simultaneous records

In the remainder of this part write $Q = Q_p = Q_a > 0$ and $\chi = 1 - \xi Q$. Choose $p > 0$ sufficiently small that $p < 1/8$ and $8p\xi < Q$. Further decreases will be harmless. Take $\eta > 0$ small in (5.2), and now set

$$A_n = \|D_n(0, e_1; B_2)\|_p, \qquad R_n = \max\{1, \max_{j\leq n} A_j/C_j\}, \qquad B_n = R_n C_n. \tag{7.1}$$

This B_n is different from the point-only envelope $\mathcal{B}_n$ used in the exponent comparison.

Lemma 7.1 (Large-gap crossing comparison)**.** *For every $\delta > 0$ there is $M_\delta < \infty$ such that*

$$c_{s+m} \geq c_s 2^{-(\chi+\delta)m}, \qquad s \geq 0, \quad m \geq M_\delta. \tag{7.2}$$

Here integer cutoffs suffice; the same proof is uniform over a fixed fractional part of the cutoff.

Proof. Split at depth m. On a mesh comparable to 2^{-m}, use translated and rescaled copies of the fixed wide reference shell. Call a fine shell good when its cost is at least $2^{-m}c_s$. Bad labels have probability at most p_0 and a fixed dependence range. Except with probability $C2^{dm}e^{-c(1+m)^2}$, their components have diameter at most $(1+m)^2$ mesh steps, and hence have at most $L_m = C_d(1+m)^{2d}$ vertices. This is the size parameter in Lemma 4.3. The largest oscillation of the coarse field on distances $CL_m 2^{-m}$ is $o_{\mathbb{P}}(m)$, uniformly in s. Indeed, its increment variance is $O(\log(m+2))$; covering the working region and applying Gaussian concentration gives a tail $C\exp[-c_\varepsilon m^2/\log(m+2)]$ for exceeding εm. Nonnegativity of the mollifier gives the same control for $G_{m,s+m} - H_m$.

Lemma 4.3, with the fixed mesh-to-shell ratio, converts every full crossing into a coarse polygonal crossing. It gives simultaneously for all paths

$$D_{s+m}(\text{across } B_{2R} \setminus \overline{B}_R) \geq c_s e^{-o_{\mathbb{P}}(m)} D_m^{\text{sharp}}(\text{across } B_{2R} \setminus \overline{B}_R),$$

where the polynomial loss in L_m is included in $e^{o(m)}$. The sharp annular exponent and Gaussian concentration make the last crossing at least $2^{-(\chi+\delta/2)m}$ with probability tending to one. The whole lower bound is uniform in s and has probability greater than $1 - p_0/2$ for large m. Its lower p_0 quantile is therefore at least the displayed threshold, proving (7.2). □

Lemma 7.2 (Envelope estimates)**.** *For every fixed $\delta > 0$ there is $K_{\eta,\delta} < \infty$ such that, for $0 \le k \le n$,*

$$2^{(\chi-\eta)k} \le C_{n-k}/C_n \le K_{\eta,\delta} 2^{(\chi+\delta)k}, \qquad B_{n-k}/B_n \le K_{\eta,\delta} 2^{(\chi+\delta)k}. \tag{7.3}$$

Moreover C_n, A_n, B_n all have exponent χ, and $\log R_n = o(n)$. If δ is small enough, for each compact family of buffered deterministic point geometries,

$$\mathbb{E}[(D_n(x,y;V)/B_n)^u] \le K_{u,V}, \qquad u = p, 2p, 4p. \tag{7.4}$$

The raw endpoint-chain and tame-point attachment estimates remain valid with B_n in place of $\mathcal{B}_n$ and Q in place of Q_p.

Proof. The lower ratio is immediate from the minimum defining C_n. For the upper ratio let j minimize that definition. If $j \le n-k$, the ratio equals $2^{(\chi-\eta)k}$. Otherwise, put $M = M_\delta$. When $n - k \ge M$,

$$c_j \ge c_{n-k-M} 2^{-(\chi+\delta)(j-n+k+M)}, \qquad C_{n-k} \le c_{n-k-M} 2^{-(\chi-\eta)M}.$$

Combining these with the formula for C_n gives $C_{n-k}/C_n \le 2^{(\delta+\eta)M} 2^{(\chi+\delta)k}$. When $n - k < M$, the large-gap estimate with starting index zero gives $C_n \ge c2^{-(\chi+\delta)n}$; the finitely many possible numerators absorb the remainder. Monotonicity of R_n proves the B ratio.

The common exponent gives $A_n/C_n = 2^{o(n)}$, and hence $R_n = 2^{o(n)}$ by taking maxima; all the exponent statements follow. Markov's inequality gives $b_j \le C_p A_j \le C_p B_j$. Apply the raw chain with these thresholds and the upper B ratio. For a deterministic endpoint the resulting series has u moment when $u(\xi Q - \delta) - u^2\xi^2/2 > 0$, by (4.13). Our choice of p permits $u = 4p$ after decreasing δ. The random starting core and terminal fallback are handled by their superexponential depth tails exactly as in Lemma 4.5. This proves (7.4) uniformly in the deterministic geometry. The simultaneous raw chain uses the same events, and the proof of Lemma 5.5 uses only that chain and the upper envelope ratio. □

Lemma 7.3 (Coarse upper moment comparison)**.** *For every $\varepsilon > 0$,*

$$A_{j+m} \le K_\varepsilon B_j 2^{-(\chi-\varepsilon)m}, \qquad j, m \ge 0. \tag{7.5}$$

For each fixed M, also $A_{j+m} \le K_M A_j$ for $0 \le m \le M$.

Proof. At mesh 2^{-m}, call a fine point edge good when its cost is at most $K2^{-m}B_j$. The uniform moment bound makes its bad probability arbitrarily small when K is large. On an event with complement of probability at most Ce^{-cm^2}, every relevant bad component has diameter at most $(m+1)^2$ mesh steps. Choose a sharp coarse near-geodesic from 0 to e_1 in a slightly smaller buffered cube. Replace its visits to bad components by good lattice detours. The lattice routing lemma gives polynomially many edges per detour, all within a polynomial cutoff neighborhood of the old path. To check the cost, partition the old path into successive first-exit mesh pieces. Each completed coarse piece pays at least $c2^{-m}e^{\xi H_m(z) - \xi O_m}$; its replacement has cost at most $KB_j 2^{-m}(1+m)^C e^{\xi H_m(z) + C\xi O_m}$. Here O_m is the maximal coarse oscillation on those polynomial neighborhoods. Summation over disjoint old time pieces gives total replacement mass at most

$$KB_j(1+m)^C e^{C\xi O_m} D_m^{\mathrm{sharp}}(0, e_1; B_{3/2}).$$

The $2p$ moment of the sharp coarse distance has exponent χ, and the fixed exponential moments of O_m are $e^{o(m)}$. Hölder's inequality therefore bounds the p norm of this term by $B_j 2^{-\chi m + o(m)}$.

At an endpoint, connect through the filled bad component containing that endpoint by using all fine edges in a containing lattice box. The component radius has an exponential tail. Indeed, its inner exterior boundary lies in one bad star-component by the lattice routing lemma, and contains every coordinate extreme of the filled component. If its radius lies between $2^k r$ and $2^{k+1} r$, that bad component has diameter at least $c2^k r$ and meets a box with at most $C(2^{k+1} r)^d$ sites about the endpoint. Bad-path counting therefore bounds this event by $C(2^{k+1} r)^d e^{-c2^k r}$. Summing over $k \geq 0$ proves a bound $Ce^{-c'r}$ for $r \geq 1$. The $2p$ moment bound for each fine edge, Cauchy–Schwarz, and summation over boxes of each possible radius show that this total fine cost has p norm at most $CB_j 2^{-m}$. This variable depends only on the fine field. Restoring the independent coarse endpoint value and its local oscillation bounds its p norm by

$$CB_j e^{o(m)} 2^{-m+p\xi^2 m/2}.$$

This decays faster than $B_j 2^{-\chi m}$ because $p\xi < 2Q$. The two endpoints are deterministic mesh vertices, so this use of the coarse Gaussian point moment is legitimate.

On the percolation or coarse-oscillation failure, use a fixed lattice chain. Its $2p$ norm is bounded by $B_j e^{Cm}$ using (7.4) and the full coarse supremum. The exceptional probabilities are $Ce^{-cm^2/\log(m+2)}$, so Cauchy–Schwarz absorbs this contribution. These estimates are uniform in j and prove (7.5), with bounded m absorbed into the constant.

For the final bounded-gap assertion, use the exact split at m and, in the rescaled domain, concatenate the 2^m translated unit reference point paths. For $p \leq 1$, subadditivity bounds the p moment of their sum by $2^m A_j^p$; no independence between these paths is required. The independent coarse supremum has bounded moments for $m \leq M$. This proves $A_{j+m} \leq K_M A_j$ also for convolution cutoffs. □

Proposition 7.4 (Simultaneous records)**.** *If R_n is unbounded, there are indices tending to infinity and a fixed $c_* > 0$ such that*

$$C_n = c_n, \qquad R_n \longrightarrow \infty, \qquad c_* \leq A_n/B_n \leq 1. \tag{7.6}$$

Proof. Take a record N of A_i/C_i and let j minimize the definition of C_N. Then $C_j = c_j$. Choose $\varepsilon < \eta$ in Lemma 7.3. Since $C_N = C_j 2^{-(\chi-\eta)(N-j)}$,

$$R_N \leq K_\varepsilon R_j 2^{-(\eta-\varepsilon)(N-j)} \leq K_\varepsilon R_N 2^{-(\eta-\varepsilon)(N-j)}.$$

Thus $N - j$ is bounded. The bounded-gap assertion gives $A_j/C_j \geq cR_N$. Since $R_j \leq R_N$, we have $A_j/B_j \geq c$, and $R_j \to \infty$. Relabel these j as n. □

Lemma 7.5 (Deterministic short bridges)**.** *For a fixed compact $V_0 \Subset B_2$ there are $\beta > 0$ and $M < \infty$ such that*

$$\mathbb{E}[(D_n(x, y; B_2)/B_n)^p] \leq M|x - y|^\beta, \qquad x, y \in V_0, \tag{7.7}$$

uniformly in n, for deterministic endpoints.

Proof. If $2^{-k} \leq |x - y| < 2^{-k+1}$ and $k \leq n$, restrict the path to a buffered ball of radius $C2^{-k}$ and split at depth k. The fine deterministic-point p moment is at most CB_{n-k}^p. Gaussian tilting at the coarse center, as in (4.13), gives

$$\mathbb{E} e^{p\xi \sup_{B_{C2^{-k}}} G_{k,n}} \leq C2^{p^2\xi^2 k/2}.$$

Independence of coarse and fine noise and (7.3) therefore bound the normalized p moment by $C2^{-\kappa_p k}$, where $\kappa_p = p(\xi Q - \delta) - p^2\xi^2/2 > 0$. If $k > n$, the straight segment and the cutoff Gaussian bound give

$$\mathbb{E}[(D_n(x,y)/B_n)^p] \le C|x-y|^p 2^{(p-\kappa_p)n} \le C|x-y|^{\min\{p,\kappa_p\}}.$$

Separations bounded below follow from (7.4). Take $\beta = \min\{p, \kappa_p\}$. □

8 A common core for arbitrary crossing arms

We next establish the local estimate used in resampling. It is an all-outcomes estimate; rare failures of the usual nice-mesh event are included in its random core mass.

Lemma 8.1 (Coordinate averaging)**.** *Suppose a rectifiable path P crosses a buffered coordinate interval I. At mesh $r_\ell = 2^{-\ell}$ suppose a selected point in each visited cell can attach to a common core with cost at most $\sum_{\ell\ge M} U_\ell(z_\ell)$; a completed old shell charges $L_\ell(z)$; supporting shells have overlap at most J_ℓ; and $U_\ell(z) \le A_\ell L_\ell(z)$. Then*

$$D(P, \mathit{core}) \le \frac{1}{|I|} \sum_{\ell\ge M} r_\ell A_\ell J_\ell \operatorname{len}(P). \tag{8.1}$$

The same assertion holds with the finite sum through a cutoff.

Proof. Let $x(t)$ be first hits of the coordinate levels. A half-open mesh cell can be selected for levels of total measure at most r_ℓ. The buffered interval ensures a completed shell about each selected cell. Choose one such shell subpath per cell. Bounded overlap gives $\sum_z L_\ell(z) \le J_\ell \operatorname{len}(P)$, so

$$\int_I U_\ell(z_\ell(x(t)))\,dt \le r_\ell A_\ell J_\ell \operatorname{len}(P).$$

Sum and choose a level whose total cost does not exceed its average. Repeated visits cause no extra charge because only one subpath is chosen for each cell. □

Proposition 8.2 (All-outcomes common core)**.** *Fix an annular band and a larger buffered annular region V. For $N \ge 0$ and $L \ge L_0$, there is a finite connected random core in V such that every finite collection of crossing arms can be attached to that same core at total cost at most*

$$CR_N 2^{-cL} \sum_j \operatorname{len}_{D_N}(P_j) + B_N Y_{N,L}, \qquad \sup_N \|Y_{N,L}\|_q \le e^{C_q L} \quad (0 < q < \infty). \tag{8.2}$$

The core term is one total mass, independent of the number of arms. The statement holds in the finitely many clipped orthant versions of the geometry. If V lies outside the spatial influence of a resampled fine-noise block, the construction and $Y_{N,L}$ use only unchanged noise. A core can additionally be required to contain a crossing of any fixed buffered coordinate interval in that annular geometry.

Proof. If $N < L$, use the smooth fallback described below. Suppose $N \ge L$. At depth ℓ, declare upper/nice labels using local fine point-edge thresholds $K2^{-\ell}B_{N-\ell}$, and lower labels using crossing thresholds $2^{-\ell}c_{N-\ell}$. The moment bound and the low quantile make the bad probabilities uniformly small. Apply the successive-level nice construction, not merely separate good components at each

level. Its bad-component bounds, the lower-tail estimate, and the coarse oscillation bounds imply that the first depth $M \geq L$ from which all these events hold through cutoff N satisfies

$$\mathbb{P}[M > m] \leq C \exp[-cm^2/\log(m+2)], \qquad m \geq L, \tag{8.3}$$

where $M = N + 1$ denotes terminal failure.

On $M \leq N$, raw chaining gives an upper attachment charge at depth ℓ of the form

$$C(1+\ell)^C 2^{-\ell} B_{N-\ell} e^{\xi H_\ell(z) + \xi O_\ell}.$$

The lower charge from the same mesh is $2^{-\ell} C_{N-\ell} e^{\xi H_\ell(z) - \xi O_\ell}$, up to the small exponential loss allowed by the simultaneous shell estimate. Their ratio is at most $C R_N (1+\ell)^C e^{\delta \ell}$. The factor $2^{-\ell}$ supplied by Lemma 8.1 makes the resulting relative attachment coefficient at most $C R_N 2^{-cM} \leq C R_N 2^{-cL}$. This is a simultaneous inequality for every old crossing arm.

The eligible attachment vertices at level M belong to the same good lattice component in a fixed buffer. Include a connected graph spanning all eligible vertices; there are at most $C2^{dM}$ of them, and loop-erased local routing uses at most $C2^{dM}$ edges after a fixed enlargement. Its total mass divided by B_N is at most

$$e^{CM} e^{\xi \sup_V H_M},$$

using the upper B ratio. Its fixed $2q$ moment grows at most $e^{C_q M}$. Cauchy–Schwarz on $\{M = m\}$ and (8.3) give the moment bound in (8.2). No independence of the starting depth and the field is needed. Adding fixed opposite-side mesh vertices and a connecting good route supplies the required coordinate crossing without changing this estimate.

For terminal failure one must control total core mass, not just diameter. Work in a smaller middle annular band with positive clearance from the determining region, and a larger buffer Ω still outside that region. Write $S_N = \sup_\Omega X_N$, $I_N = \inf_\Omega X_N$ and choose

$$J = N + \lceil cL + \xi(S_N - I_N)/\log 2 \rceil + C_0, \qquad h = 2^{-J}.$$

Use the entire axis-aligned lattice annulus of mesh h, or a spanning tree containing all its vertices. Every point of the middle band is within Ch of a vertex. Each old crossing arm has mass at least $a e^{\xi I_N}$, while its short straight attachment costs at most $Che^{\xi S_N}$. Their ratio is at most $C2^{-cL}$. The entire lattice graph has at most Ch^{-d} edges and total mass at most

$$Ch^{1-d} e^{\xi S_N}. \tag{8.4}$$

Since $B_N^{-1} \leq Ce^{CN}$ and $\mathbb{E}e^{q \sup_\Omega |X_N|} \leq e^{C_q(N+1)}$, the normalized $2q$ moment of (8.4) is at most $e^{C_q(N+L+1)}$. If $N < L$ this is the desired bound. If $N \geq L$, Cauchy–Schwarz and (8.3) absorb the additional factor depending on N.

Here are the boundary details used in both constructions. If a buffer touches an ambient face, enlarge the obstacle by a fixed factor and move that coordinate of its center to the face. Repeat in the other coordinates; at most d enlargements occur. At the sufficiently fine block scales used later, opposite faces cannot both be encountered. This gives only finitely many orthant geometries. Shift each lattice inward by an integer $a \geq 8$ of its mesh steps, chosen large enough that every fine-edge domain is inside the orthant. At mesh h, an inward coordinate is at least ah; every vertex of the finer mesh within $3h$ has that coordinate at least $(a-3)h \geq ah/2$, the finer inward boundary. The coarse vertices themselves belong to the finer lattice. Thus the successive-level construction uses compatible grids. Reflect the site labels across each inward grid boundary to find lattice routes, then

fold the routes and use the original interior edge paths. A site has at most 2^d reflected copies, so finite-range bad-path counting is unchanged up to constants. A reflected connected bad component folds into one old bad component, and its diameter changes by at most a dimensional factor. Thus the same local routes and depth tails apply. The smooth terminal lattice is connected in every such annular orthant: keep one coordinate outside the inner cube while changing the other coordinates, and then change the holding coordinate. Positive band widths absorb rounding and inward shifts. This also proves the stated coordinate-crossing and locality assertions. □

9 Concentration of the marked-network functional

The concentration estimate below is the quantitative step that distinguishes bounded comparison of normalizers from equality of their leading exponents. Continue to use $\max\{0, Q - 1/\xi\} < \alpha < Q$ in the barrier, and put $s_n = 1 + \log R_n$.

Lemma 9.1 (A subset coarea bound)**.** *If a nonnegative measurable Z on an interval J satisfies*

$$|\{t : Z(t) > v\}| \le |J| \min\{1, Ae^{-cv}\}, \qquad A \ge 1,$$

then for every measurable $E \subset J$ of positive measure,

$$\int_E Z(t)\,dt \le c^{-1}|E| \left(1 + \log A + \log \frac{|J|}{|E|}\right). \tag{9.1}$$

In a fixed buffered unit geometry there is a random Z_N^* with

$$\sup_N \mathbb{E}e^{qZ_N^*} < \infty \qquad (0 < q < \infty) \tag{9.2}$$

such that, simultaneously for all crossing paths P and their first-hit coordinate sections,

$$\begin{aligned} &\frac{1}{|J|}\left|\left\{t : \max_{0\le j\le N}(H_j(x_P(t)) - \alpha j \log 2)_+ > v\right\}\right| \le \min\{1, A_P e^{-cv}\}, \\ &\log A_P \le C(1 + \log^+(\mathrm{len}_{D_N}(P)/C_N) + Z_N^*). \end{aligned} \tag{9.3}$$

Proof. Layer cake bounds the left side of (9.1) by

$$\int_0^\infty \min\{|E|, |J|Ae^{-cv}\}\,dv.$$

Split this integral at $c^{-1}\log(|J|A/|E|)$.

For the second assertion, use the simultaneous lower-shell and oscillation events from a first good depth M_N onward. They satisfy $\mathbb{P}[M_N > M] \le C\exp[-cM^2/\log(M+2)]$. For $j \ge M_N$, the shell-charging proof of (5.11) in unit geometry gives

$$\frac{|\{t : H_j(x_P(t)) > \alpha j \log 2 + v\}|}{|J|} \le C\frac{\mathrm{len}_{D_N}(P)}{C_N}(1+j)^C 2^{-(1+\xi\alpha-\xi Q-\delta)j}e^{-\xi v}.$$

Choose δ so that the coefficient of j in the exponent is positive, and sum. For the earlier scales use

$$Z_N^* = C(1 + M_N) + \max_{0\le j\le \min\{M_N, N\}} \sup_V |H_j|.$$

At a deterministic M, the fixed exponential moments of the maximum through scale M are at most $e^{C_q(M+1)}$. Cauchy–Schwarz on $\{M_N = M\}$ and the depth tail prove (9.2), including $M_N = N + 1$. Increasing A_P by $e^{cZ_N^*}$ handles every earlier scale and gives (9.3) on every outcome. □

Lemma 9.2 (Surgery preserving all exterior marked points)**.** *Fix a fine-noise block of spatial side 2^{-K}, with $K \geq m + C_0$, and let $U_S \subset V_S \subset W_S$ be fixed-factor buffers containing its actual spatial influence. A marked network can be replaced by one avoiding U_S, retaining every old marked point outside V_S, and having mass increment at most*

$$CR_n 2^{-cL} e^{CG_K} \ell_S + C2^{-K} B_{n-K} e^{\xi H_K(v_S) + CG_K} Y_S. \tag{9.4}$$

Here ℓ_S is old annular-arm mass, $\sum_S \ell_S \leq C \operatorname{mass}_{D_n}(G)$, G_K is the maximum coarse oscillation on the buffers, and $\|Y_S\|_q \leq e^{C_q L}$ for every fixed q. Only constantly many core masses are absorbed into Y_S. If $\Delta_{i,S}$ bounds the increase of U_i under this surgery, then

$$\sum_S \mathbb{E}(\Delta_{0,S} + \Delta_{1,S})^2 \leq C2^{m-K}(K + s_n + L + 1)^2. \tag{9.5}$$

Proof. If G misses U_S, use the identity surgery. Otherwise retain the components of $G \setminus U_S$ that meet the complement of V_S. Each reaches the inner buffer and has a crossing arm through the intervening band. Choose one arm per retained component. Distinct components give disjoint old arms, whose total mass is at most the old mass in that band. Attach all of them to the single core of Proposition 8.2. Count the core's connected graph only once. This preserves every old point outside V_S. Apply the same operation, when necessary, to each G_i inside its target cube. Because G_i spans an interval of length comparable to $2^{-m} \gg 2^{-K}$, it has an unchanged point outside V_S; there it meets the unchanged global graph. Thus the union remains connected. Global and two local surgeries use only a fixed number of cores and a fixed multiplicity of old arms. The block buffers have bounded spatial overlap, proving the sum bound for ℓ_S. Restore the coarse field in the rescaled core estimate to obtain (9.4). Enlarging the buffers once includes all discarded points. The clipped orthant construction in Proposition 8.2 keeps every path in the ambient or target cube.

For a marked target, require the local replacement core to span a full coordinate interval J_S of length comparable to 2^{-K} containing the projection of V_S, with a fixed buffer on both sides. The middle-quarter condition on I_i and $K \geq m + C_0$ ensure that no face perpendicular to this coordinate interferes. The remaining faces are the clipped orthant cases already treated. Spanning only $J_S \cap I_i$ would not be sufficient here.

Choose a measurable minimizing section $x_i(t)$ of the old graph. Let

$$E_{i,S} = \{t \in I_i : x_i(t) \in V_S\}, \qquad \ell_{i,S} = |E_{i,S}|/|I_i|, \qquad \delta_K = C2^{m-K}.$$

Pointwise in the field,

$$0 \leq \ell_{i,S} \leq \delta_K, \qquad \sum_S \ell_{i,S} \leq C. \tag{9.6}$$

Outside $E_{i,S}$ the old section point is retained and its barrier is unchanged by refresh. Inside it use a crossing section of the replacement core. In its rescaled fine geometry, that crossing has mass at most $B_{n-K} Y_S$. Thus (9.3) has

$$\log A_P \leq C(s_n + \log^+ Y_S + Z^*_{n-K}).$$

Returning to the barrier centered at $H_m(z_i)$ adds only the coarse increments between levels m and K. Over all $O(2^{dK})$ blocks, their fixed moments are $O_q(K+1)$. The exponential moments of Z^* and fixed positive moments of Y_S, followed by a union bound, therefore give random coefficients $P_S \geq 1$ with

$$\|\max_S P_S\|_2 \leq C(K + s_n + L + 1). \tag{9.7}$$

Apply the subset estimate on the full interval J_S and drop the old nonnegative penalty. It yields

$$\Delta_{i,S} \le C\ell_{i,S}\left(P_S + \log\frac{\delta_K}{\ell_{i,S}}\right),$$

with zero defined by continuity. Since $x\log^2(\delta_K/x) \le C\delta_K$ for $0 < x \le \delta_K$, (9.6) gives

$$\sum_S \Delta_{i,S}^2 \le C\delta_K(1 + \max_S P_S^2).$$

Taking expectations and summing over i proves (9.5). The argument never conditions on a selected point being independent of the field. □

Proposition 9.3 (Network Bernstein concentration)**.** *There are constants $C, c, \kappa_0 > 0$, independent of n, m, λ and of the deterministic target centers in the fixed geometric family, such that for $\lambda \ge 1$, $m \le \min\{n, \kappa_0 s_n\}$, and $t \ge 0$,*

$$\mathbb{P}[|\log F_{n,m,\lambda} - \mathbb{E}\log F_{n,m,\lambda}| > t] \le C\exp\left[-\frac{ct^2}{s_n + \lambda^2 + t}\right]. \tag{9.8}$$

Proof. Split at depth $K \ge m + C_0$, put $N = n - K$, and denote the independent fine noise by $\mathcal{F}$. Let $f_K(\mathcal{F}) = \mathbb{E}_H \log F(H, \mathcal{F})$. There are $O(2^{dK})$ fine spatial blocks influencing the bounded ambient domain. Choose an old measurable almost-minimizing graph; all bounds below are uniform in its error, which can be sent to zero after the finite collection of blocks is fixed.

For an active block, write u_S for its core mass divided by F. Let θ_S be the old annular mass divided by F. Then $0 \le \theta_S \le C$ and $\sum_S \theta_S \le C$. Every active block has a complete old annular arm, because the graph connects the two separated targets and the local marked graphs span intervals larger than the block buffers. Let $Z_S \ge 1$ be the inverse of a local fine crossing in C_N units. It has every fixed moment by Lemma 5.2. Both Z_S and the core variable Y_S can be chosen from fine noise outside the refreshed block. The upper core mass and lower old crossing charge imply

$$u_S \le CR_nY_SZ_Se^{CG_K}\theta_S, \qquad \sum_S u_S \le CR_nM_{\mathcal{F}}e^{CG_K}, \quad M_{\mathcal{F}} = \max_S Y_SZ_S. \tag{9.9}$$

On the other hand, $F \ge 2C_n$ and the upper envelope ratio give

$$u_S \le CR_nY_Se^{CG_K}2^{-(\xi Q-\delta)K}e^{\xi H_K(v_S)}.$$

Choose $0 < \delta < \xi Q$, put $a = \xi Q - \delta > 0$, and combine these bounds to get

$$u_S \le CR_nY_SZ_Se^{CG_K}\min\{1, 2^{-aK}e^{\xi H_K(v_S)}\}. \tag{9.10}$$

The coarse oscillation satisfies

$$\mathbb{E}e^{qG_K} \le C_qe^{C_q\sqrt{K}} \tag{9.11}$$

for every fixed q: on each buffer its increment process has bounded variance and entropy, and there are $O(2^{dK})$ buffers. For $X \sim N(0, K\log 2)$, split at $aK\log 2/(2\xi)$ to obtain

$$\mathbb{E}\min\{1, 2^{-aK}e^{\xi X}\}^2 \le 2^{-aK} + e^{-a^2K\log 2/(8\xi^2)} \le 2e^{-bK}$$

for $b = b(a, \xi) > 0$. Average the coarse field *before* taking a spatial maximum. Conditional Cauchy–Schwarz in (9.10), using (9.11), and then (9.9), give, for $z_S = \mathbb{E}_H u_S$,

$$\max_S z_S \le CR_nM_{\mathcal{F}}e^{-bK/2+C\sqrt{K}}, \qquad \sum_S z_S \le CR_nM_{\mathcal{F}}e^{C\sqrt{K}}.$$

All the v_S here are deterministic mesh centers. The fixed moment bounds and the number of blocks imply

$$\mathbb{E}M_{\mathcal{F}}^2 \le C\exp[(2d\log 2/q)K + C_qL]$$

for any sufficiently large fixed q. Choose it once so that $2d\log 2/q < b/4$. It follows that

$$\sum_S \mathbb{E}z_S^2 \le CR_n^2 e^{C_1L - a_1K + C_2\sqrt{K}}. \tag{9.12}$$

Taking the maximum of the coarse field first would lose this positive-Q estimate.

The relative-arm contribution in (9.4), divided by F, is bounded by $v_S = CR_n 2^{-cL}e^{CG_K}\theta_S$. Jensen's inequality, $\sum_S \theta_S^2 \le C$, and (9.11) show

$$\sum_S \mathbb{E}(\mathbb{E}_H v_S)^2 \le CR_n^2 e^{-a_2L + C_3\sqrt{K}}. \tag{9.13}$$

Let $\mathcal{F}^S$ refresh one block. The old graph supplies an admissible new graph outside the determining region. Its old sections outside the buffers and all its new cores are unchanged by refresh. Therefore, with ΔL_S the nonnegative mass increment bound,

$$(\log F(H,\mathcal{F}^S) - \log F(H,\mathcal{F}))_+ \le \log(1 + \Delta L_S/F(H,\mathcal{F})) + \lambda(\Delta_{0,S} + \Delta_{1,S}). \tag{9.14}$$

This is the new-minus-old orientation. Exchangeability of the old and refreshed block permits precisely this one-sided form of Efron–Stein. Use conditional Jensen, $\log(1+x) \le x$, and $(x+y+z)^2 \le 3(x^2+y^2+z^2)$; no independence of mixed terms is needed. Equations (9.12), (9.13), and (9.5) give

$$\operatorname{Var} f_K \le C\left[R_n^2 e^{C_1L - a_1K + C_2\sqrt{K}} + R_n^2 e^{-a_2L + C_3\sqrt{K}} + \lambda^2 2^{m-K}(K + s_n + L + 1)^2\right]. \tag{9.15}$$

For $t \ge 1$, set $u = s_n + \lambda^2 + t$, $L = \lceil bu \rceil$, and $K = \lceil au \rceil$. First choose b large for the arm term, then a large for the core and penalty terms and for all geometric buffers. A fixed small κ_0 ensures $K \ge m + C_0$. Every term in (9.15) is at most Ce^{-cu}. All moment orders were fixed before these choices; in particular none grows with n, t, or λ. Lemma 6.1 gives conditional coarse Gaussian concentration with squared Lipschitz constant at most $C(K+\lambda^2) \le Cu$. Thus

$$\mathbb{P}[|\log F - f_K| > t/2] \le 2e^{-t^2/(Cu)}, \qquad \mathbb{P}[|f_K - \mathbb{E}\log F| > t/2] \le Ct^{-2}e^{-cu}.$$

Their sum is at most the right side of (9.8), since $t^2/u \le u$. If the prescribed K exceeds n, the full-field Lipschitz bound $C(n+1+\lambda^2) \le Cu$ gives the same conclusion directly. Increase C for $0 \le t < 1$. □

10 The bounded point-to-crossing comparison

Lemma 10.1 (Cheap paths near the reference endpoints)**.** *For every sufficiently small fixed $\delta > 0$ there are $L_\delta < \infty$ and $a_\delta > 0$, independent of n, such that*

$$\mathbb{P}[D_n(B_{\delta/8}(0), B_{\delta/8}(e_1); B_2) \le L_\delta c_n] \ge a_\delta. \tag{10.1}$$

Proof. Put $\ell = 2^{k_0}$, so the reference annulus has inner radius 4ℓ. From a cheap source crossing at cutoff $n - k_0$, stop the initial segment at its first Euclidean displacement ℓ. Cover its starting

boundary point and displacement direction by $O(\delta^{-(d-1)})$ patches each. One pair of patches has cheap probability at least $c\delta^{2(d-1)}$. The patch centers can be chosen at separation exactly ℓ, and the stopped path stays in a fixed containing ball. Translation, rotation, and scaling by ℓ^{-1} map this geometry to small balls about $0, e_1$ inside B_2.

The scaled source field has the law of the fine portion below 2^{-k_0} of X_n. Adding the independent coarse field $G_{k_0,n}$ gives exactly X_n. Conditional on the selected source path, its expected new length is at most $\ell^{-1}e^{\xi^2 k_0 \log 2/2}$ times its old length. Conditional Markov preserves a fixed fraction of the source success probability with a fixed multiplicative cost. This argument requires neither independence of the path from its fine field nor a cutoff-dependent sup-norm comparison.

It remains to replace c_{n-k_0} by c_n. At physical scale 2^{-k_0}, fine reference shells have fixed radii $4, 8$ and a fixed dependence range. A crossing of the original wide annulus that pays no good local crossing produces a bad chain of length comparable to 2^{k_0}. Choose k_0 large and then p_0 small so that its probability is less than $p_0/4$. The coarse infimum is bounded below by a deterministic constant except with probability $p_0/4$. Consequently the original crossing is at least a fixed positive multiple of $2^{-k_0}c_{n-k_0}$ with probability greater than $1-p_0/2$. Its lower p_0 quantile gives $c_{n-k_0} \le Cc_n$. The finitely many smaller cutoffs are absorbed into constants. □

Lemma 10.2 (Localized network low event)**.** *Fix $\delta > 0$ as above and $\varepsilon > 0$. Choose all shell, oscillation, and envelope slacks sufficiently small. For every sufficiently large $m = o(n)$, at contact indices $C_n = c_n$, there is a deterministic pair of targets $W_i \subset B_\delta(ie_1)$, of radius comparable to 2^{-m}, with buffered coordinate intervals, for which*

$$\mathbb{P}[F_{n,m,\lambda} \le C_n(L + 2e^{\lambda\varepsilon m})] \ge e^{-c_0 m}. \tag{10.2}$$

Here L, c_0 are fixed before λ is chosen, and the targets may depend on m, n but are deterministic.

Proof. On (10.1), a near-minimizing path has an initial portion crossing the annulus of radii $\delta/4, \delta/2$ about zero, and a disjoint final portion crossing the corresponding annulus about e_1. Each remains inside the appropriate δ neighborhood. Apply the weighted-arm argument to these fixed geometries at mesh 2^{-m}. The sharp coarse annular exponent, lower crossing tail, and coarse oscillation give $\sum_i 2^{-m\xi Q}e^{\xi H_m(z_i)} \ge e^{-\zeta m}$ outside an event of probability tending to zero. Thus a short weighted arm exists in each portion. The coarea barrier proof then gives section average at most εm, after reducing intermediate slacks. The probability of the original event intersected with these simultaneous good events is at least $a_\delta/2$ for large m.

There are $O_\delta(2^{2dm})$ target pairs and a fixed number of coordinate and interval choices for each. Pigeonhole one pair. The path and its two arms give a marked network of mass at most LC_n. Countable polygonal approximation preserves the stated bounds with arbitrarily small losses, absorbed by the prior strict slack in ε and a fixed increase of L. This proves (10.2). □

Theorem 10.3 (Bounded comparison at every positive exponent)**.** *If $Q_p > 0$, then $Q_a = Q_p = Q$ and there is $\kappa > 0$ such that*

$$c_n \ge \kappa b_n \qquad (n \ge 0). \tag{10.3}$$

In fact the envelopes in (7.1) *satisfy* $\sup_n R_n < \infty$.

Proof. Equality of the exponents was proved in Proposition 6.2. Suppose R_n is unbounded and work along the simultaneous records (7.6). Put $s = s_n \to \infty$; the common exponent gives $s = o(n)$.

We first specify the parameter order. Fix

$$\max\{0, Q - 1/\xi\} < \alpha < Q, \qquad 0 < \tau < \min\{1/16, Q\log 2/32\}.$$

Choose a tame-attachment loss $\varepsilon_1 < \xi Q \log 2/4$ and a low-event coefficient $\varepsilon < \tau/4$. Fix the moment and all envelope, lower-shell, and oscillation slacks small enough for the preceding estimates with these choices. This fixes c_* in (7.6) and M, β in Lemma 7.5. Next fix $\delta \in (0, 1/8)$ so small that

$$2M\delta^{\beta} < c_*^p/2. \tag{10.4}$$

This fixes the source event and the geometric constants in the network concentration estimate. Choose the fixed λ only after these choices, sufficiently large as specified below, and put $m = \lfloor s/\lambda \rfloor$.

The low event has probability at least $e^{-c_0 m}$ and upper threshold $\log(F/C_n) \le \lambda\varepsilon m + \log(L+2)$. For a variable with tail (9.8), such a low event implies

$$\mathbb{E}\log(F/C_n) \le \lambda\varepsilon m + C_1\sqrt{(s+\lambda^2)(m+1)} + C_1 m + C_1. \tag{10.5}$$

Indeed, if the mean exceeds the threshold by $a > 0$, its lower tail gives $ca^2/(s+\lambda^2+a) \le c_0 m + \log C$; solving this quadratic yields the stated bound. Choose $\lambda \ge 2/\kappa_0$. Then $m \le \kappa_0 s$, and (9.8) with $t = s^{3/4}$ and (10.5) show that, with probability tending to one,

$$\log(F/C_n) \le (\varepsilon + C_2/\sqrt{\lambda} + C_2/\lambda + o(1))s.$$

Take the fixed λ large enough that the coefficient before $o(1)$ is smaller than $\tau/2$. Thus $F/B_n \to 0$ exponentially in s, and a measurable near-minimizer of objective at most $F + C_n$ satisfies $U_i \le \tau m$ with probability tending to one. In each marked subgraph choose a point x_i with barrier at most $2\tau m$. The tame-point attachment estimate in B_n units gives

$$D_n(z_i, x_i; B_2)/B_n \le e^{-am}, \qquad i = 0, 1,$$

with probability tending to one, for a fixed $a > 0$. The centers z_i are deterministic, although they depend on n, so its uniform deterministic-center probability bound applies. The connected graph supplies a path between x_0, x_1 of length at most $F + C_n$. Consequently

$$Y_n := D_n(z_0, z_1; B_2)/B_n \longrightarrow 0 \quad \text{in probability.} \tag{10.6}$$

The uniform $2p$ moment bound makes this convergence also hold in L^p.

Finally use the actual reference endpoints. Since $0 < p \le 1$, the triangle inequality and (7.7) give

$$(A_n/B_n)^p \le \mathbb{E}(D_n(0, z_0)/B_n)^p + \mathbb{E}Y_n^p + \mathbb{E}(D_n(z_1, e_1)/B_n)^p \le 2M\delta^{\beta} + o(1).$$

This contradicts (7.6) and (10.4). Hence R_n is bounded. Markov's inequality gives $b_n \le (1-q)^{-1/p} A_n \le CB_n \le CC_n \le Cc_n$, which is (10.3). □

11 Quantitative output for the metric construction

Theorem 11.1 (Uniform crossing and endpoint estimates)**.** *Assume $d \ge 3$ and $Q_p > 0$. For both sharp and compact convolution cutoffs the point and annular exponents agree and are positive. With $\chi = 1 - \xi Q$ the following hold.*

1. *For every $\eta > 0$ there are positive constants c_η, C_η such that*

$$c_\eta 2^{(\chi-\eta)(t-s)} \le b_s/b_t \le C_\eta 2^{(\chi+\eta)(t-s)}, \qquad 0 \le s \le t. \tag{11.1}$$

2. *For every fixed annulus A and every fixed $q \in (0,\infty)$,*

$$\sup_{t\geq 0} \mathbb{E}\left[\left(\frac{b_t}{D_t(across\ A)}\right)^q\right] < \infty.$$

More precisely, the lower tail is bounded by $C\exp[-c(\log T)^2/\log\log T]$ at threshold b_t/T for sufficiently large T.

3. *There is $p > 0$ for which all fixed buffered deterministic point geometries have uniformly bounded $2p$ moments in b_t units. On each compact subset of a connected open domain,*

$$\mathbb{E}[(D_t(x,y;U)/b_t)^p] \leq C|x-y|^\beta \tag{11.2}$$

for deterministic nearby endpoints and some $\beta > 0$.

4. *The common core estimate* (8.2) *holds with a bounded constant in place of R_N and with b_N in place of B_N. Its relative arm coefficient can therefore be made arbitrarily small by choosing the fixed mesh depth L large.*

Proof. For integer cutoffs, (5.1), Theorem 10.3, and $A_n \leq B_n = R_nC_n$ show $b_n \asymp c_n \asymp C_n \asymp B_n$. The envelope ratios give item 1, the inverse-tail lemma gives item 2, and the point moment and short-bridge lemmas give item 3. The common-core proposition gives item 4. Since the envelope slack can initially be chosen arbitrarily small, item 1 holds with every prescribed $\eta > 0$.

We spell out uniformity over real cutoffs. Write $t = n + \theta$, $0 \leq \theta < 1$, and repeat the proof with $A_n^\theta, c_n^\theta, C_n^\theta, R_n^\theta$ for the cutoff sequence $n + \theta$. Integer scale deletion preserves θ. The residual unit-cutoff fields form a compact interval with uniform variance, increment, supremum, and finite-dependence bounds. The point and annular exponent errors are uniform in θ by the common-field leading-exponent comparison. Every constant above is consequently uniform in θ. If the ratios were unbounded over real cutoffs, choose (n_j, θ_j) exhibiting divergence and apply the simultaneous-record reduction in each fixed-θ_j sequence. The resulting contact indices tend to infinity, since all bounded-cutoff quantile ratios are uniformly bounded. The same network contradiction applies to these varying fractional parts. Thus the point/crossing comparison is uniform for all real t.

For a remaining fractional scale gap $u \in [0,1]$, split exactly at depth u. The coarse supremum has uniform fixed moments, and the residual point geometry is dilated by $2^u \in [1,2]$. The uniform point moments yield $b_{t+u} \leq Cb_t$. For the reverse bound, apply the large-gap crossing argument from $t - M$ to $t + u$, with fixed large integer M; its proof is uniform for real depths at least M and gives $c_{t+u} \geq c\, c_{t-M}$. The already proved integer-gap comparison within the fractional-part sequence gives $c_{t-M} \geq c' b_t$. Bounded initial cutoffs are absorbed by Gaussian compact-field bounds. This proves the fractional-gap bridge, hence all of (11.1). □

Proposition 11.2 (Local connectors with finite candidate vertices)**.** *Fix an annular band A and a larger annular buffer V. For every $\varepsilon > 0$ and $p_* > 0$, there are a deterministic finite set of core candidate vertices, a constant $C < \infty$, and local events of probability at least $1 - p_*$, uniformly at all sufficiently fine cutoffs, on which the following hold simultaneously for all finite-length crossing arms. Each arm P attaches within V to a candidate vertex at cost at most $\varepsilon \operatorname{len}_{D_t}(P)$, and all its possible attachment vertices lie in one connected core. For every two arms P_1, P_2,*

$$D_t(P_1, P_2; V) \leq C(\operatorname{len}_{D_t}(P_1) + \operatorname{len}_{D_t}(P_2)). \tag{11.3}$$

These events can be defined from the metric data in the buffered region. After deleting an additional fixed coarse depth, their fine-noise determining regions lie in a slightly enlarged annular buffer.

Proof. Choose L so the bounded coefficient in (8.2) is at most ε. Restrict the first good depth to $M \le M_0$, with M_0 fixed large enough that its failure probability is at most $p_*/4$. For cutoffs beyond M_0, every eligible vertex belongs to the deterministic union of the grids at depths $L, \dots, M_0$ in the fixed working region. This union is finite. Restrict the total core variable to $Y \le T$ and the annular crossing to be at least ab_t, choosing T large and a small using the core moments and inverse crossing tail. The two further failure probabilities are at most $p_*/4$ each. On this event both old arms have length at least ab_t and the connected core has total mass at most Tb_t. Attach both arms and use the core to obtain (11.3) with $C = \varepsilon + T/(2a)$.

Every actual edge path and attachment stays in V. The event asserting the displayed intrinsic inequalities and the finite candidate-core property is measurable from the density on that region: at fixed cutoff, all paths can be approximated by a countable polygonal family. It contains the constructed event and has at least its probability. For the white-noise version, delete a fixed additional number of coarse scales so every remaining kernel support used by a working tube lies in the enlarged annulus. Coarse oscillation is handled separately by the Gaussian estimates. This distinguishes locality of the metric event from locality of its determining white noise. □

12 From annular estimates to a metric

We use Theorem 11.1 and Lemma 11.2. Distances are normalized by the deterministic point normalizer; D_n denotes the resulting smooth length metric along a cutoff sequence. The following consequences will be used with fixed positive buffers.

1. Internal distances in every fixed deterministic point geometry are tight. Annular crossings and their reciprocals are tight, with negative moments of every positive order. For some $p, \beta > 0$, deterministic nearby points in a compact region satisfy

$$\mathbb{E}[D_n(x, y; U)^p] \le C|x - y|^\beta \tag{12.1}$$

   when the cutoff is smaller than their separation and U is an admissible fixed buffer.

2. Given any success probability $p_0 < 1$, a deterministic C_0 and an annular event of probability at least p_0 can be chosen uniformly in sufficiently fine cutoffs, centers, and scales. On this event every two crossing arms can be joined in a fixed compact working buffer at cost at most

$$C_0\{\mathrm{len}_{D_n}(P) + \mathrm{len}_{D_n}(R)\}. \tag{12.2}$$

   Each arm also attaches, at cost at most $C_0 \,\mathrm{len}_{D_n}(P)$, to a vertex of a deterministic finite menu. The menu has uniformly bounded cardinality and rational vertices for rational annuli. Both statements hold simultaneously for all arms, including arms chosen after observing the field.

3. These events are local modulo constants. The annular density estimate and relative normalizer bounds apply. Test bands can be widened by fixed amounts while retaining compact working buffers.

For the finite menu, impose a deterministic upper bound on the first good core depth. Its failure probability is made arbitrarily small before fixing the menu. The random-core estimate used for moment bounds is retained separately without this cap.

12.1 Simultaneous annuli

Lemma 12.1 (Density at all centers)**.** *The success probability can be chosen so that in every enhanced subsequential coupling the following holds almost surely. At every sufficiently large stage k, every center of a grid of mesh $2^{-2k}/C$ in a fixed bounded region has at least ck successful annuli among the scales 2^{-m}, $k \le m \le 2k$, in a fixed separated residue class. Consequently, simultaneously around every point of $\mathbb{R}^d$ there are infinitely many shrinking successful rational annuli with pairwise disjoint enlarged working buffers and a positive inner buffer about the point.*

Proof. Lemma 2.9 gives a failure bound e^{-BN} for a prescribed proportion of failures among N separated annuli, with B arbitrarily large when the marginal success probability is sufficiently close to one. At stage k there are $O(2^{2dk})$ grid centers and at least $k/(2L)$ tested levels per center, where L is the fixed separation. Choose B so that the union bound is summable in k.

For a fixed stage only finitely many binary indicators occur. Retain them in the coupled extraction of Lemma A.6; almost sure convergence makes each indicator eventually constant. The same failure bound passes to the limiting indicators. Borel–Cantelli proves the grid assertion. An arbitrary point is within a sufficiently small fraction of the least tested radius from a nearest grid center. Every candidate annulus therefore surrounds that point with the required inner buffer. Stages $k = 3^j$ have scale ranges sufficiently separated to give nested annuli and disjoint working buffers. A countable exhaustion proves the assertion throughout space. The event is always used at its actual grid center; no continuity of a success event in its center is assumed. □

Lemma 12.2 (Large annular barriers)**.** *The extracted normalized crossings of $B_{2^{j+1}}(0) \setminus B_{2^j}(0)$ tend to infinity almost surely.*

Proof. For $r = 2^j$, prelimit scaling gives the crossing in law as

$$\frac{r a_{\varepsilon/r}}{a_\varepsilon} e^{\xi h_r(0)} Y_{\varepsilon/r},$$

where $Y_{\varepsilon/r}$ is a normalized unit crossing. The two random factors need not be independent. The deterministic factor is at least $c_\eta r^{\xi Q_p - \eta}$ for every $\eta > 0$. Lemma 2.1 gives $\operatorname{Var} h_r(0) = \log r + O(1)$. Choose $\eta < \xi Q_p/8$ and use a sufficiently high negative moment of Y. A union bound shows that the displayed expression is at least $r^{\xi Q_p/2}$ except with probability Ce^{-cj}, uniformly in fine cutoffs. Pass each fixed closed lower-bound event to the extracted crossing and use Borel–Cantelli. □

12.2 A countable construction

Let $\mathcal{Q}$ be the cubes with rational centers and positive rational radii. Extract jointly the field, all distances between boundaries in $\mathcal{Q}$, all rational point distances, all rational-point internal distances in rational polycubical buffers, the success indicators, the normalizer ratios at positive rational scales, and the random core masses when needed. These are tight countably many coordinates. Lemma A.6 makes their convergence simultaneous in one coupling. The field coordinates retain their fixed joint law, and the smooth prelimit metrics remain measurable functions of the prelimit field copies in this representation.

Write $\dot{D}(S,T)$ for an extracted boundary-set distance. It is finite, since a fixed pair of boundary points gives an upper bound. If S, T are disjoint it is positive: cover S by finitely many rational inner cubes whose doubled cubes avoid T. Every connecting path crosses one of their annuli, so their positive limiting crossings have a positive minimum.

Set

$$Q_m(x) = B_{3\cdot 2^{-m}}\big(2^{-m}\lfloor 2^m x\rfloor\big).$$

These rational cubes are nested, contain x with positive buffer, and shrink to x. Define

$$D(x,y) = \lim_{m\to\infty} \dot{D}(\partial Q_m(x), \partial Q_m(y)) \quad (x \neq y), \qquad D(x,x) = 0. \tag{12.3}$$

Once the cubes are disjoint the terms are nondecreasing. Equivalently, $D(x,y)$ is the supremum of $\dot{D}(\partial V, \partial W)$ over disjoint rational cube neighborhoods $V \ni x$, $W \ni y$.

Proposition 12.3 (Metric limit)**.** *The function D is a measurable Euclidean lower-semicontinuous separating extended metric. The metrics D_n converge to D in the epigraph topology, jointly with*

$$D_n(p,q) \longrightarrow D(p,q) < \infty, \qquad p,q \in \mathbb{Q}^d. \tag{12.4}$$

Proof. The supremum description is a countable supremum of nonnegative constants on open product neighborhoods, extended by zero outside those neighborhoods. It proves measurability and lower semicontinuity. Symmetry is immediate; a rational annulus separating two distinct points proves strict positivity.

If $(x_n, y_n) \to (x, y)$ with $x \neq y$, the endpoints eventually belong to any fixed pair of disjoint rational neighborhoods. Their distance dominates the boundary-set distance. Taking the lower limit and then the supremum yields

$$D(x,y) \le \liminf_n D_n(x_n, y_n). \tag{12.5}$$

Nonnegativity gives this on the diagonal. Conversely, choose nearly minimizing endpoints on the boundaries in (12.3) and let the boundary index increase slowly with n, so that each fixed boundary limit has already been approximated. This gives recovery endpoints. Infinite values are also covered; on the diagonal take identical endpoints. Thus there is epigraph convergence.

For the triangle inequality, suppose $D(x,z) + D(z,y) = L < \infty$. Fix N successful annuli about z with disjoint working buffers avoiding x, y. Recovery paths for the two pairs cross all these annuli for large n. Their crossing-arm lengths sum to at most $L + o(1)$. In at least one annulus the two arm lengths therefore sum to at most $(L + o(1))/N$. Join them by (12.2), truncate, and concatenate. Equation (12.5) gives

$$D(x,y) \le D(x,z) + D(z,y) + C_0 L/N.$$

Let $N \to \infty$. This proves the triangle inequality without assuming endpoint continuity.

For rational consistency, fix rational p, and let F_j be the finite menu for the rational annulus centered at p with radius 2^{-j}. Its cardinality is uniformly bounded and its vertices are within $C2^{-j}$ of p. The small-point moment estimate gives, for the extracted variables,

$$A_j = \lim_n \max_{v\in F_j} D_n(p,v), \qquad \mathbb{E}[A_j^p] \le C2^{-\beta j}.$$

Fatou justifies the moment bound. Hence $A_j \to 0$ almost surely, simultaneously for rational p; infinitely many of these fixed-center annuli are successful by annular density.

For rational p, q, let $L = D(p,q)$. It is finite already by (12.5) and the finite extracted point value. Choose N successful disjoint annuli near each endpoint with all their $A_j < \delta$. A recovery path crosses them all. Near each endpoint one arm has length at most $(L + o(1))/N$. Attach it to its finite menu, then to the original fixed endpoint. The added cost is at most $\delta + C_0(L + o(1))/N + o(1)$. Only finitely many already convergent coordinates are involved. It follows that

$$\lim_n D_n(p,q) \le L + 2\delta + 2C_0 L/N.$$

Let $N \to \infty$ and then $\delta \to 0$. The opposite inequality is (12.5). □

12.3 Compact paths and the finite component

Lemma 12.4 (Compact separation)**.** *Disjoint Euclidean compact sets have positive distance for any separating Euclidean lower-semicontinuous extended metric. In particular, for compact K and $a > 0$,*

$$\inf\{D(x, y) : x, y \in K,\ |x - y| \geq a\} > 0.$$

There is also an eventual positive prelimit lower bound for the length of a path starting in K and making Euclidean displacement at least a.

Proof. A sequence with distances tending to zero has convergent endpoint subsequences; lower semicontinuity would give distinct limiting endpoints at distance zero. For the prelimit assertion cover K by finitely many rational inner cubes whose doubled cubes have diameter below a. A path starting in an inner cube and making the stated displacement crosses its annulus. The minimum of the positive limiting crossings gives the eventual lower bound. □

Lemma 12.5 (Compactness of paths)**.** *Paths in a fixed compact set, parameterized on a common compact interval with a uniform D_n-speed bound, have uniformly Euclidean-convergent subsequences. For such a limit,*

$$D(P(s), P(t)) \leq \liminf_n D_n(P_n(s), P_n(t)), \qquad \mathrm{len}_D(P) \leq \liminf_n \mathrm{len}_{D_n}(P_n).$$

Paths of uniformly bounded length whose starting points lie in a fixed compact set are confined to one compact set for all sufficiently large n.

Proof. Compact separation turns the speed bound into Euclidean equicontinuity. The dense-time diagonal argument in Lemma A.9 gives the subsequence. Equation (12.5) and finite partition sums give the two inequalities. The speed bound makes the limit metric-Lipschitz. For confinement choose a large annulus whose limiting crossing exceeds the common length bound, using Lemma 12.2. □

Lemma 12.6 (Simple arcs inside a compact path image)**.** *Let D be a separating Euclidean lower-semicontinuous metric. If a finite D-rectifiable path has compact Euclidean image K and distinct endpoints x, y, then K contains a simple D-rectifiable arc from x to y whose length is no greater than that of the original path.*

Proof. Let L be the infimum of lengths of paths in K joining x to y. It is finite by the given path and positive because $L \geq D(x, y) > 0$. Choose paths with lengths decreasing to L and parameterize them with constant speed on $[0, 1]$. Compact separation turns their common speed bound into a common Euclidean modulus. The dense-time extraction of Lemma A.9 gives a uniformly convergent subsequence whose limit remains in K. Lower semicontinuity at every finite partition gives limiting length at most L; the same lower inequality gives $D(P(s), P(t)) \leq L|s - t|$, so the limit is an admissible metric-continuous path. Its length is therefore exactly L. Parameterize this minimizer by arclength. If it visited the same point twice, deleting the intervening positive-length portion would give a shorter path in K, contrary to minimality. Thus it is a simple arc and has the required length bound. □

Proposition 12.7 (Finite-component geometry)**.** *Every finite-distance pair has a geodesic. All nonsingular points belong to one complete finite-distance component; its bounded sets are Euclidean bounded. For every nonsingular x, open $U \ni x$, and $\eta > 0$ there is $q \in \mathbb{Q}^d \cap U$ with $D(x, q; U) < \eta$.*

Proof. Take recovery endpoints and paths with lengths tending to $D(x,y) < \infty$. Confinement, constant-speed parametrization, and compact extraction give $D(P(s),P(t)) \le D(x,y)|t-s|$. The triangle inequality between the endpoints forces equality, so this is a geodesic.

The large barriers show that a bounded D-ball is Euclidean bounded: recovery paths to points beyond a fixed sufficiently large annulus would exceed its radius. A D-Cauchy sequence is therefore Euclidean bounded. Compact separation makes it Euclidean Cauchy. If x is its Euclidean limit, lower semicontinuity gives $D(x_n,x) \le \liminf_m D(x_n,x_m) \to 0$. The limit stays in its finite component, proving completeness.

For the density assertion take a geodesic P from x to a distinct point, of length L, and its constant-speed recovery paths from the first paragraph. Choose $t > 0$ such that $(1+C_0)tL < \eta$ and $P([0,t])$ lies compactly in U. A successful annulus about x can be chosen inside U with outer region avoiding $P(t)$. Initial recovery paths cross it and attach to a fixed finite rational menu at cost at most $C_0t(L+o(1))$. On a subsequence the menu vertex is a fixed rational q. Compact extraction inside U gives $D(x,q;U) \le (1+C_0)tL < \eta$. Thus every nonsingular point belongs to the rational component; all rational pairs have finite distance by (12.4). □

12.4 Recovery and limiting connectors

Lemma 12.8 (Path recovery)**.** *Every finite D-length path P admits prelimit paths P_n converging uniformly Euclideanly, with*

$$\mathrm{len}_{D_n}(P_n) \to \mathrm{len}_D(P).$$

Rational endpoints can be kept fixed. The recoveries can be confined to any open neighborhood of the compact image of P.

Proof. If P is constant, take the constant prelimit path at its value; this includes a constant path at a singular point. Otherwise every point of P is nonsingular, since it is connected by a finite subpath to another point of the image. Fix Euclidean error $a > 0$ and length error $\eta > 0$. Subdivide P into finitely many pieces of Euclidean diameter below $a/4$ and metric length below one quarter of the compact prelimit exit barrier at radius $a/4$. Choose rational vertices within intrinsic distance δ of the partition vertices; keep prescribed rational endpoints. Take δ so small that the total added length is below $\eta/2$ and all vertices are within Euclidean distance $a/4$ of their intended positions.

The finitely many rational-pair convergences give prelimit links whose costs are at most the piece lengths plus 2δ and arbitrarily small excess. Their costs are below the exit barrier, so each link stays within Euclidean distance $a/4$ of its initial vertex. Concatenation, with each link parameterized on its original time interval, has error at most a and total cost at most $\mathrm{len}_D(P) + \eta$. Choose a below the clearance from the prescribed open neighborhood. Let a, η decrease slowly with n. Finite partition sums give the reverse length inequality. □

Lemma 12.9 (Universal limiting connectors)**.** *At every retained successful annulus, simultaneously for all finite limiting arms crossing its wider test band,*

$$D(P,R;V) \le C_0\{\mathrm{len}_D(P) + \mathrm{len}_D(R)\}.$$

The all-outcomes random-core bounds, including a single common core mass for finitely many arms and its moment bounds, pass to the limit as well.

Proof. Recover the two arms. Their uniform convergence and the positive boundary buffers ensure crossings of the original test band. The retained success bit is eventually one. Apply its simultaneous

prelimit connector and extract a bounded-length path in the compact working buffer. Its endpoints belong to the limiting arms and its cost has the asserted bound. This is a deterministic argument on the already fixed realization, so the assertion is simultaneous over arms.

Retain each random core mass as a separate coordinate. Tightness and Fatou pass its moments. For finitely many arms, prune the prelimit connected network to a tree spanning its attachment points and suppress degree-two vertices. A tree spanning k contacts has at most $k-2$ branch vertices. Extract its finite abstract type and its bounded-length edges; the sum of limiting edge lengths is at most the lower limit of their total mass. The core is therefore charged only once. □

Lemma 12.10 (Bounded-path germ access)**.** *Let x be nonsingular, $U \ni x$ open, and $T < \infty$. For every $\eta > 0$ there are $0 < \rho < r$ with $\overline{B}_r(x) \subset U$ such that every finite path of length at most T meeting $B_\rho(x)$ and leaving $B_r(x)$ has a point joined to x inside U at cost below η. The access can be confined to any prescribed neighborhood of x.*

Proof. Choose a nonconstant initial geodesic of length δ inside the prescribed neighborhood, with $(1 + C_0)\delta < \eta/2$. Take r smaller than the Euclidean distance of its other endpoint from x. Choose N disjoint successful annuli inside $B_r(x)$, with $C_0 T/N < \eta/2$, and take ρ inside all their inner holes. The reference and tested paths cross them all. One tested arm has length at most T/N. Its connector to the reference arm costs at most $C_0(T/N + \delta)$. Prepending the reference path gives the desired cost. The choices are independent of the tested path. □

12.5 Actual intrinsic metrics

Proposition 12.11 (Intrinsic reconstruction)**.** *On the same event, for every open U the actual intrinsic metric $D^U = D(\cdot,\cdot;U)$ is lower semicontinuous in the relative Euclidean topology of $U \times U$ and, for $x, y \in U$, obeys*

$$D^U(x,y) = \liminf_{a \downarrow 0} \{ D^U(p,q) : p, q \in \mathbb{Q}^d \cap U,\ |p - x| < a,\ |q - y| < a \}. \tag{12.6}$$

Coincident rational points are allowed on the diagonal. For rational $p, q \in U$,

$$\begin{aligned} D^U(p,q) &= \inf_{V \Subset U} \liminf_n D_n(p,q;V) \\ &= \inf_{V \Subset U} \limsup_n D_n(p,q;V), \end{aligned} \tag{12.7}$$

where V ranges over a fixed countable family of rational bounded polycubical domains containing p, q. In every connected open set, any two nonsingular points have finite intrinsic distance.

Proof. For rational points in a rational cube with closure contained in U, the extracted internal coordinate is finite. Almost minimizing prelimit paths have compact limits in the closed cube. A polygonal route in a connected open set can be covered by finitely many such cubes, with rational points chosen in consecutive overlaps. Concatenation gives finite paths between all rational points, simultaneously for every connected open set, including a set chosen after the metric. Local rational density then gives the last assertion.

Suppose $x_k \to x$, $y_k \to y$ in U, with bounded nearly optimal intrinsic path lengths. If $x = y$, nonnegativity suffices. Otherwise global lower semicontinuity gives $D(x,y) < \infty$, so both endpoints are nonsingular. Apply Lemma 12.10 in disjoint small neighborhoods of the two endpoints, to initial and terminal pieces of the paths. The two selected contacts are ordered and have arbitrarily small attachment costs. Retain the original middle portion. This proves $D^U(x,y) \le \liminf_k D^U(x_k, y_k)$;

the middle portions need not stay uniformly away from ∂U. Local rational density and the triangle inequality give one bound in (12.6), and lower semicontinuity gives the other, including infinite values.

For $V \Subset U$, compact limits of paths in $\overline{V}$ show that its displayed lower limit is at least $D^U(p,q)$. Conversely, a finite limiting path has compact image in some such V. Recover it there with fixed rational endpoints by Lemma 12.8. Taking almost minimizing paths proves both equalities. □

Corollary 12.12 (Deterministic point finiteness)**.** *For every fixed deterministic connected open U and fixed $x, y \in U$, one has $D(x, y; U) < \infty$ almost surely. The continuum-finiteness axiom consequently holds.*

Proof. Retain the finitely many buffered point coordinates joining x, y along a deterministic chain of cubes in U. Their upper tightness and compact path extraction give a finite limiting path. Adding these coordinates does not change the already constructed metric, by uniqueness of its epigraph limit along the enhanced subsequence. For any two fixed nonempty compact sets choose one fixed point in each and apply this assertion. Separate probability-one statements for deterministic points do not assert simultaneous finiteness at every point. □

12.6 Boundary ports with one-sided access

Lemma 12.13 (Strict one-sided boundary access)**.** *Fix a countable family of rational cubical collectors. Almost surely every rational point y on one of their boundaries has finite metric access from both sides: in each sufficiently small neighborhood of y there are rational points q_- in the open exterior and q_+ in the interior, and finite D-rectifiable arcs from y to q_- and q_+ whose points other than y belong to their respective open sides. The arcs can be confined to fixed inward and outward coordinate cones.*

These accesses have uniform small positive moments in the local scale unit for every fixed finite family of normalized face and corner geometries. More precisely, if their dyadic construction is truncated after depth J, its remaining length has pth moment at most $C2^{-\beta J}$ in that unit, for some $p, \beta > 0$. The statement, including this tail estimate, is uniform over the enhanced limiting laws constructed above.

Proof. We give the argument on one side. Choose a rational vector v strictly inside a coordinate cone on that side, and put $y_j = y + 2^{-j} r v$, $j \geq 0$. There is a fixed rational open polycube O containing v and $v/2$ whose closure is strictly inside that cone and bounded away from its vertex. Choose $O' \Subset O$ containing these two points, and set $O_j = y + 2^{-j} r O$, $O'_j = y + 2^{-j} r O'$. Consecutive points y_j, y_{j+1} have positive clearance in O'_j, and every O_j is strictly on the chosen side. For a sufficiently small r all these buffers lie in the prescribed neighborhood.

Take $0 < p \leq 1$ small enough for the deterministic-point estimate. Its scaled version gives

$$\mathbb{E}[(L_j/s_r)^p] \leq C2^{-\beta j}, \qquad L_j = 2D(y_j, y_{j+1}; O_j), \tag{12.8}$$

with a uniform constant in the stated finite family. Here s_r is the coarse-field local unit. Before exact covariance is proved it is $c(r)e^{\xi h_r(z)}$, where $c(r) = \lim_n r a_{\varepsilon_n/r}/a_{\varepsilon_n}$ along the retained subsequence and the normalized location of $y - z$ is fixed. This use of a local unit assumes no scale covariance of the limiting metric. Indeed, apply the exact prelimit coordinate change and use the point estimate in O'_j. The coarse Gaussian factor and the normalizer ratio give the decaying exponent $p(\xi Q - \eta) - p^2\xi^2 > 0$, after decreasing p, η. Here Cauchy–Schwarz and the available $2p$ point

moment allow dependence on the local coarse normalization. This is precisely the calculation for the short-bridge estimate, now with the fixed buffered shape $(v, v/2; O')$. Compact path extraction in $\overline{O'_j} \subset O_j$ and Fatou's lemma give (12.8) for the actual limiting intrinsic distances. This also proves uniformity over enhanced limiting laws.

For every j choose a finite path in O_j with length at most L_j. Such paths may be taken continuous in metric arclength: equivalently, use the compact limits of prelimit paths in O'_j and arbitrarily small slack. Since $p \le 1$,

$$\mathbb{E}\left[\left(s_r^{-1}\sum_{j\ge J} L_j\right)^p\right] \le C\sum_{j\ge J} 2^{-\beta j} \le C'2^{-\beta J}.$$

The sum is finite almost surely. Concatenate these paths in reverse order, starting at the limiting point y, and then parameterize by total length. Their Euclidean supports shrink to y. Lower semicontinuity gives

$$D(y, y_J) \le \liminf_{k\to\infty} D(y_k, y_J) \le \sum_{j\ge J} L_j.$$

The same bound on a tail subpath proves metric continuity at y and bounds the length of the entire concatenation by $\sum_j L_j$. Every other point of the concatenation is in an open-side buffer. Lemma 12.6 gives a simple rectifiable arc in this same compact image, with the same one-sided property and no greater length. The construction on the other side is identical. There are only finitely many coordinate-cone types at cubical faces and corners. Countability of the specified ports and buffers gives the simultaneous probability-one assertion. □

Lemma 12.14 (Anchored clipped cores)**.** *The high-probability clipped-exterior connector events can be augmented as follows. Choose a deterministic boundary port y in a buffered boundary patch of the annulus, an exterior rational port q_-, and an interior rational port q_+. On the augmented event there are strict one-sided paths y to q_- and y to q_+, and open-exterior paths from q_- to every candidate vertex of the clipped core, whose total length is at most Hs_r. The clipped crossing is at least as_r, where $H < \infty$ and $a > 0$ are deterministic. The probability of this augmented event can be made arbitrarily close to one, uniformly in scale and in the finitely many boundary geometries. All paths remain in one fixed enlarged annular data region.*

Proof. Use the inward-shifted grids in the clipped-core construction. On the event that the first good depth is at most its fixed cap, every candidate vertex is strictly in the open exterior band, with positive normalized clearance. This is a deterministic finite menu. The open exterior band is connected in each face or corner geometry; choose deterministic polygonal routes from q_- to all menu vertices in a slightly larger exterior band. Finitely many positive-clearance rational polycubes cover these routes. The deterministic-point estimate in those cubes gives a uniformly bounded small positive moment for their total connecting cost. Lemma 12.13 gives the same bound for the two boundary-port chains, all confined to a small patch of the enlarged annulus.

First choose the good-depth cap. Then choose H large by Markov's inequality, and a small by the reciprocal crossing estimate. A finite union bound proves the asserted probability. The constructions use only metric data in their displayed buffers; coarse subtraction changes all costs by the bounded oscillation factors already included in the annular event. They therefore satisfy the same fine-noise annular density argument. For limiting events retain the buffered distances of the countably many port-chain links and their summed cost. Fatou's lemma bounds the sum after extraction. The tail estimate in Lemma 12.13 permits finite initial-link recovery with uniformly

small remaining cost. Hence enhanced limits preserve the augmented conclusion with any fixed strict relaxation. □

12.7 Continuous weights

For positive continuous w, define $w \cdot D$ by minimizing the integral of w against D-arclength on finite D-rectifiable paths.

Lemma 12.15 (Local weighted length calculus)**.** *The metric $w \cdot D$ separates points and is a length metric. Its finite rectifiable paths have compact Euclidean image, are D-rectifiable, and satisfy*

$$\mathrm{len}_{w\cdot D}(P) = \int_P w \, ds_D. \tag{12.9}$$

For positive continuous v, w, $v \cdot (w \cdot D) = (vw) \cdot D$. The singular points of $w \cdot D$ are those of D.

Proof. A defining path from x to the complement of a small Euclidean ball pays at least the positive minimum of w on that ball times $D(x, \partial B) > 0$. This proves separation and shows that the weighted topology controls the Euclidean topology. A weighted-continuous path on a compact interval has compact Euclidean image K.

Choose a compact buffer K' around K. Compact separation gives an unweighted exit barrier from K to $\partial K'$, and multiplication by $\min_{K'} w$ gives a positive weighted exit barrier. Weighted connections of cost below half that barrier stay in K' and cost at least $\min_{K'} w$ times their D-lengths. A finite weighted path is therefore D-continuous and D-rectifiable. In neighborhoods on which w oscillates by at most a prescribed amount, the same barriers compare short distances by the local minimum and maximum of w. Refine a path partition and decrease the oscillation. The lower and upper arclength sums give (12.9), the length property, and the composition rule.

Nonsingularity is local: a finite path to a distinct point has a nonconstant initial segment in a compact neighborhood. On that neighborhood w is bounded above and away from zero. Such a segment exists for one metric exactly when it exists for the other. □

Proposition 12.16 (Weighted recovery)**.** *Let w_n, w be positive continuous functions converging uniformly on compact sets. Finite limiting paths admit the recoveries of Lemma 12.8 with convergence of weighted lengths. For uniformly convergent paths confined to a fixed compact set, weighted length is lower semicontinuous.*

Every actual intrinsic metric $(w \cdot D)^U$ is lower semicontinuous on $U \times U$ in its relative Euclidean topology, obeys the rational envelope there, and satisfies (12.7) *with the weighted prelimit metrics. With common global upper and positive lower bounds for the weights, the global weighted metrics converge epigraphically, jointly with rational evaluations. These assertions are deterministic and simultaneous over the weights on the enhanced event.*

Proof. For upper recovery subdivide into small pieces on which w has small oscillation. The recovered links lie in the corresponding compact neighborhoods. Multiply each link cost by the maximum of w_n on that neighborhood. Uniform convergence of the weights, followed by decreasing oscillation and length errors, proves the upper bound.

For compactly confined uniformly convergent P_n, a positive lower weight bound controls the unweighted lengths. On every finite partition,

$$\mathrm{len}_{w_n\cdot D_n}(P_n) \geq \sum_j \left(\inf_{t\in[t_j,t_{j+1}]} w_n(P_n(t)) \right) D_n(P_n(t_j), P_n(t_{j+1})).$$

Pass to the lower limit and refine partitions. The limiting finite-variation path is metric-continuous: one-sided limits exist by completeness, and the compact inverse modulus identifies them with the given Euclidean-continuous values. Lower weighted arclength sums then give (12.9).

Germ access works also for weighted bounded paths in a compact neighborhood. Divide their arm lengths by the local minimum of w and multiply connector costs by its local maximum. If bounded weighted paths have convergent distinct endpoints, their initial segments to fixed small exit spheres have bounded unweighted length. Compact limits show that the endpoint limits are nonsingular. Weighted germ access attaches these endpoints at arbitrarily small cost. This proves intrinsic lower semicontinuity and the rational envelope. Compact lower extraction and weighted recovery give the buffered formulas.

Under global two-sided weight bounds, large annular barriers confine global bounded-cost paths as well. The same lower extraction and upper recovery give global epigraph convergence. All arguments are pathwise with finitely many compact local weight bounds. In particular they apply simultaneously to $w = e^{\xi f}$ and $w_n = e^{\xi(f*\rho_{\varepsilon_n})}$ for every continuous f, including a field-dependent one. Without global weight bounds, the buffered formulas, rather than global convergence, are the assertion being used. □

12.8 Conditional products before field measurability

We retain the entire white noise W in the enhanced coupling. It is a random distribution on space and scale, represented by a countable determining family of test integrals. Its restrictions to disjoint measurable white-noise regions are independent.

Proposition 12.17 (Conditional products)**.** *For pairwise disjoint deterministic open sets U_i, the limiting intrinsic coordinates D^{U_i} are conditionally independent given the entire W. The same assertion holds jointly for finite tuples of subsequential metric coordinates and for conditionally independent samples of the whole limiting metric given W. On any strictly separated finite collection of inner buffers, after a smooth coarse field has been removed, their conditional kernels involve disjoint fine-noise regions.*

Proof. First fix finitely many compact inner buffers $\overline{V}_i \subset U_i$ with positive pairwise separation. Choose a deterministic scale $s > 0$ so small that the spatial white-noise footprints

$$A_i = \{(y,t) : 0 < t < s,\ \operatorname{dist}(y, \overline{V}_i) < R_K s + \delta\}$$

are disjoint, with a positive extra margin δ. For sufficiently small convolution cutoffs the fine field $H_{(0,s)} * \rho_\varepsilon$ on V_i depends only on $W|_{A_i}$. The complement of the fine field in the pinned field is a random smooth function g_s on every compact set; it includes the pinning constant. All its local smooth restrictions are functions of W.

Retain $W|_{A_i}$, the complementary white noise, g_s, and the buffered prelimit metric coordinates for $h - g_s = H_{(0,s)}$. Before taking limits the complete blocks

$$(W|_{A_i}, \text{fine-field metric coordinates in } V_i)$$

are independent of one another and of the complementary white noise. Their joint limit retains this independence. This statement is about joint blocks, not a bare claim that conditional independence is closed under weak convergence.

Weighted buffered reconstruction identifies the actual limiting fine-field coordinates as those of $e^{-\xi g_s} \cdot D$ on smaller buffers. Conditional on the full W, the smooth correction g_s is known. Restoring

it therefore preserves conditional independence. The kernels of the fine coordinates depend only on their indicated restrictions of W.

Use countably many rational inner buffers and exhaust each U_i. For any finite number of coordinates in such an exhaustion, one deterministic scale s separates their footprints. The buffered reconstruction and rational envelope recover the actual intrinsic metrics. Their finite-dimensional conditional laws hence factor. This proves the assertion for a finite collection of open sets, including touching boundaries, and then for a countable collection. Retaining tuples throughout proves the joint assertion. Independent conditional copies of the full metric preserve the same factorization. □

Proposition 12.18 (Locality after removal of residual randomness)**.** *Suppose that the conditional comparison and resampling argument proves that an enhanced limiting metric is a measurable function of W. Then it is a measurable function of h. It admits one Borel assignment on the base field plus continuous perturbations, with actual intrinsic locality and simultaneous Weyl scaling.*

Proof. Since the original sequence uses the same W, Lemma A.7 turns joint convergence to a measurable function of W into convergence in probability on that probability space. Apply this to the global metric and its consistent global rational point distances: by (12.4), these coordinates are functions of the limiting metric. Choose a deterministic further subsequence along which that metric and that countable rational array converge almost surely.

These coordinates suffice for the pathwise reconstruction. Epigraph convergence to a separating metric gives eventual positive lower bounds on compact sets of Euclidean-separated endpoint pairs. For confinement, place two compact boundary surfaces strictly on opposite sides of a limiting large annular barrier from Lemma 12.2. The recovery-endpoint argument shows that every limiting distance between these two surfaces is at least that barrier. Compactness and the epigraph lower inequality transfer this bound to the original prelimit metrics. Thus bounded-cost paths from a fixed compact set stay in a common compact set. The rational-link argument of Lemma 12.8 and the weighted buffered reconstruction now apply on this original-space event, using only the global metric and global rational array just retained.

Every prelimit metric is a measurable function of h. An almost sure limit of these variables is measurable with respect to the completed field sigma field. Thus the limiting metric has a Borel version determined by h; the same conclusion applies to its global rational point coordinates.

For rational p, q in an open set U, define on distributional inputs

$$F_U(g; p, q) = \inf_{V \Subset U} \limsup_{n\to\infty} a_{\varepsilon_n}^{-1} D_g^{\varepsilon_n}(p, q; V).$$

For each $V \Subset U$, all sufficiently late terms depend only on $g|_U$. This is therefore a local measurable formula. On the common original-space event just obtained, weighted buffered reconstruction identifies, simultaneously for every continuous f,

$$F_U(h + f; p, q) = (e^{\xi f} \cdot D_h)(p, q; U).$$

Use the rational envelope at all endpoints. The formula depends on the input $h + f$, not on a chosen decomposition of that input, so it is consistent whenever two decompositions are possible. The weighted length calculus gives the length property and composition law.

The harmless convention on other inputs can also be made Borel without projecting an arbitrary set of distributions. Symmetrize the nonnegative rational array and put its diagonal equal to zero before taking its lower-semicontinuous envelope F on $\mathbb{R}^{2d}$. If that array is $a(q)$, its continuous

Lipschitz envelopes are

$$F_k(z) = \min\left\{k,\ \inf_{q\in\mathbb{Q}^{2d}}[a(q) + k|z-q|]\right\}, \qquad k \geq 1.$$

They are Borel functions of the input array, continuous in z, and increase to F. On every fixed compact set, the minima of an increasing sequence of continuous functions converge to the minimum of their lower-semicontinuous limit. To see this, take minimizers, extract a convergent subsequence, and first test any one fixed member of the increasing sequence. Continuous compact minima can be computed on a countable dense set. Consequently all the compact minima used below are Borel operations on the original array.

Separation is the countable condition that F have positive minimum on $\{|x|, |y| \leq m,\ |x-y| \geq 1/m\}$ for every positive integer m. Symmetry and the zero diagonal already hold. A triangle violation occurs if and only if there are rational compact cubes Q_x, Q_y, Q_z and nonnegative rational a, b, c with $a + b < c$ such that

$$\min_{Q_x\times Q_z} F > c, \qquad \min_{Q_x\times Q_y\times Q_z} \max\{F(x,y) - a, F(y,z) - b\} \leq 0.$$

The reverse implication is immediate; for the forward implication choose a, b just above the two finite short sides and use lower semicontinuity of the longer side to choose Q_x, Q_z. The second compact minimum is obtained from the increasing continuous functions with F_k in place of F. This is a countable Borel test for the triangle inequality, including infinite distances. On inputs failing one of these tests use the infinite discrete metric. This changes no relevant field-plus-continuous-function input. The local formulas identify the actual intrinsic metrics, proving locality. Their pathwise weighted identities give the simultaneous Weyl rule, including dependent continuous perturbations. □

13 Comparison, determination by the field, and rigidity

We write D^U for the actual length metric in an open set U, and $D^{[K]}$ for the length metric constrained to a closed set K. All curves in this section are continuous for the metric in which their length is measured. The compact lower barriers proved above imply their Euclidean continuity. A *crossing arm* of a cubical shell is a subcurve joining its two boundary components and otherwise contained in the shell. Its endpoints may be chosen after observing the field.

We use the following consequences of the preceding construction. They are recorded together to fix their quantifiers. The raw scale constants of a canonical subsequential assignment are denoted by c_r, and its local unit is

$$s_r(z) = c_r \exp\{\xi h_r(z)\}.$$

For a strong assignment, $c_r = r^{\xi Q}$. A fixed change between spherical and cubical averaging conventions is included in a smooth coarse field. Every normalized annular test uses a determining annulus enlarged once to include its averaging sphere, while retaining a positive inner radius. Perturbation supports stay away from that sphere. Thus the averaging sphere may lie in the inner hole of a perturbation support while remaining inside the determining field annulus; local Gaussian block estimates and exact invariance of the counting units both apply.

1. Canonical limits have one nontrivial complete finite-distance component, containing every deterministic point almost surely. Rational points are intrinsically dense in that component in

every open neighborhood. For each open U, the actual internal metric is lower semicontinuous on $U \times U$ in its relative Euclidean topology. Compact lower barriers, finite-curve recovery in open buffers, rational point convergence, and simultaneous continuous Weyl scaling hold on one enhanced realization.

2. In a fixed annular template, a finite connected rational core gives, for every $\epsilon > 0$, simultaneous connections of all crossing arms, with

$$E(A_1, A_2; U) \le \epsilon\{\mathrm{len}_E(A_1) + \mathrm{len}_E(A_2)\} + s_r Y_{r,\epsilon}, \qquad \sup_{r,z} \mathbb{E}[Y_{r,\epsilon}^p] < \infty \tag{13.1}$$

for every fixed $p < \infty$. One core mass suffices for any fixed finite collection of arms. There is also a high-probability version with a fixed finite menu of rational vertices and no additive term, after charging the core mass to a positive lower crossing. The same statements hold in the closed exterior of a cube, with rational upper paths in its open exterior. Their constants are uniform near faces and corners.

3. Normalized crossings have every negative moment; deterministic-point distances in fixed buffered shapes have some positive moment. The raw constants satisfy, for every $\eta > 0$,

$$C_\eta^{-1}(r/R)^{\xi Q+\eta} \le c_r/c_R \le C_\eta(r/R)^{\xi Q-\eta}, \qquad 0 < r \le R. \tag{13.2}$$

Coarse average increments have the Gaussian moment bounds stated above.

4. Local annular tests of sufficiently high probability occur at any prescribed density below one, simultaneously at all points of a compact set in every sufficiently large logarithmic block. The failure probability can be made $O(e^{-An})$ for any fixed A, on blocks of n scales. Before field determination, this assertion applies to enhanced conditional samples through their local white-noise product kernels. After field determination it applies to arbitrary normalized local tests by the Gaussian annular block inequality.
5. The preceding assertions are stable under enhanced subsequential extraction of rescaled canonical assignments, including pairs on the same field, rational arrays, finite annular certificates, and the buffered recovery of finite paths.

The first four assertions are conclusions of the analytic and metric construction. The enhanced-closure assertion in the fifth item is proved in Proposition 13.11, before its first use. None is an additional axiom for a competing strong metric.

13.1 Lower estimates for an arbitrary strong metric

Lemma 13.1. *Let D satisfy the five strong axioms with $Q > 0$. Finite-distance components are complete and geodesic, finite rectifiable curves are Euclidean continuous, and metric-bounded sets are Euclidean bounded. A crossing of a fixed cubical shell is positive, finite, and determined by the field on any open annular buffer of that shell.*

Proof. The distance between the two compact boundary components is positive by separation and lower semicontinuity, and finite by continuum finiteness. Exact covariance gives the fixed law of its quotient by $r^{\xi Q}e^{\xi h_r(z)}$. At the origin this crossing tends to infinity in probability as $r \to \infty$, since $h_r(0)/\log r \to 0$ in probability and $Q > 0$. Choose deterministic $R_j \uparrow \infty$, with $R_{j+1} > 4R_j$, so that the radius-R_j crossing exceeds j eventually almost surely.

A metric-continuous curve of finite length misses at least one boundary of every sufficiently large such shell. Its inverse image of a Euclidean compact set is closed: if $t_n \to t$ and $P(t_n) \to x$

Euclideanly along a subsequence, then $D(x, P(t)) \le \liminf_n D(P(t_n), P(t)) = 0$. The omitted boundaries give a countable disjoint closed partition of the parameter interval by bounded radial bands. Lemma B.1 shows that the curve lies in one bounded band. The same closed-graph argument now gives Euclidean continuity. Every path to a sufficiently remote point crosses a large barrier, so bounded metric balls are Euclidean bounded.

On a Euclidean compact set K, for every $a > 0$,

$$\inf\{D(x, y) : x, y \in K,\ |x - y| \ge a\} > 0.$$

Consequently bounded-speed near-minimizers are Euclidean equicontinuous on a common compact set. Their uniform limits, together with lower semicontinuity, give geodesics. A metric Cauchy sequence is Euclidean precompact, and lower semicontinuity at a Euclidean subsequential limit gives metric convergence. Finally trim a crossing path at its last inner-boundary visit before its first outer-boundary visit. This identifies the unrestricted infimum with the closed-shell and open-buffer infima and proves locality. □

Lemma 13.2. *For every fixed annular template and every $p > 0$, its strong-metric crossing $X_r(z)$ satisfies*

$$\sup_{z,r} \mathbb{E}\left[\left(\frac{X_r(z)}{r^{\xi Q} e^{\xi h_r(z)}}\right)^{-p}\right] < \infty.$$

The corresponding fine-field crossing at its own scale also has all negative moments. On compact sets the identity from D to the Euclidean metric is Hölder with a positive deterministic exponent.

Proof. Write $h = H_m + X_m$, where X_m uses white noise below scale 2^{-m} and H_m is smooth, including the pinning correction. In a radius-2^{-m} shell the fine-field crossing, divided by $2^{-m\xi Q}$, has one positive finite law. Locality for the shifted field $h - H_m$ makes these variables finitely dependent on the rescaled lattice. Choose a small lower threshold $a > 0$ so the probability of a bad site is below a finite-range Peierls threshold. A self-avoiding lattice path of N bad sites contains N/C independent sites, so the probability that a bad component has macroscopic diameter is at most

$$C 2^{dm} e^{-c 2^m}.$$

Outside this event every macroscopic crossing traverses a good microscopic shell. On a fixed larger cube W,

$$X_1(0) \ge a 2^{-m\xi Q} \exp\{\xi \inf_W H_m\}.$$

The coarse field has variance $O(m + 1)$ and expected supremum $O(m + 1)$. Gaussian concentration therefore gives, on choosing $m = \lfloor 2 \log_2 \log T \rfloor$,

$$\mathbb{P}[X_1(0) < T^{-1}] \le C \exp\left\{-c \frac{(\log T)^2}{\log \log T}\right\}, \qquad T \ge e^e.$$

Integrating proves every negative moment. Removing the scale-one smooth coarse field and applying Hölder's inequality gives the fine-field statement; no independence of the pinning term is used.

For the Hölder assertion, at scale $r = 2^{-n}$ union-bound over $O(r^{-d})$ grid centers. Arbitrarily high inverse moments imply that all normalized crossings exceed r^η eventually, for any fixed $\eta > 0$. Gaussian lower tails for the field averages imply that they all exceed $-A \log r^{-1}$ eventually for a sufficiently large deterministic A. Thus every grid crossing costs at least $r^{\xi Q + \xi A + \eta}$. A shell of radius comparable to $|x - y|$ separates each pair of sufficiently close points, giving the asserted inverse Euclidean Hölder bound. □

13.2 A prefix comparison

Here is the deterministic comparison mechanism. It uses no a priori finiteness at a prescribed point of the metric being bounded.

Lemma 13.3. *Suppose two finite nonconstant metric germs start at x, for metrics D, E with compact inverse Euclidean Hölder exponents $a_D, a_E > 0$. Suppose that at least κn nested successful annuli surround x at levels in $[n, 2n]$, with centers within $O(2^{-2n})$ of x. There are deterministic $\theta_D, \theta_E > 0$ and arbitrarily small successful annuli for which the first inner and outer exit times satisfy*

$$b_D \geq \theta_D a'_D, \qquad b_E \geq \theta_E a'_E.$$

Here a'_D, a'_E are exit times, not the Hölder exponents.

Proof. Order the successful annuli from outside to inside. For either germ, nesting gives $a'_{i+1} \leq b_i$, and hence

$$\sum_i \log(a'_i/b_i) \leq \log(a'_1/b_N).$$

The innermost inner exit is at Euclidean displacement at least $c2^{-2n}$. The inverse Hölder inequality bounds the last display by $(2\log 2/a_D)n + O(1)$ or $(2\log 2/a_E)n + O(1)$. Choose θ_D, θ_E so small that fewer than $\kappa n/3$ annuli fail each balance. At least one annulus satisfies both balances, for every sufficiently large block. □

Proposition 13.4. *Suppose that around every point there are positive densities of successful nested annuli as in Lemma 13.3. Suppose on each such annulus every D-arm A_D and E-arm A_E have an E-connector satisfying*

$$\mathrm{len}_E(\sigma) \leq \epsilon\, \mathrm{len}_E(A_E) + K_0 R_* \,\mathrm{len}_D(A_D). \tag{13.3}$$

Assume also that every D-arm connects in E, at cost at most $K_0 R_*$ times its length, to a nonsingular seed tending Euclideanly to the annulus center as its radius tends to zero. If $\epsilon < \theta_E$, then

$$E(x, y) \leq (K_0 R_*/\theta_D) D(x, y)$$

simultaneously for all pairs.

Proof. Write $K = K_0 R_*/\theta_D$ and let $\Gamma : [0, L] \to \mathbb{R}^d$ be a unit-speed finite D-geodesic from x to y. A successful annulus crossed before time δ gives a seed $q_\delta \to x$ and a contact $t_\delta < \delta$ with $E(q_\delta, \Gamma(t_\delta)) \leq K_0 R_* \delta$. Set $\eta = K_0 R_* \delta$ and

$$T = \{t \in [t_\delta, L] : E(q_\delta, \Gamma(t)) \leq Kt + \eta\}.$$

This is nonempty and closed. If its largest element t is below L, take an E-geodesic prefix from q_δ to $\Gamma(t)$. If these points coincide, use a nonconstant out-and-back E-germ of length at most η. Reverse this prefix and pair it with the forward D-germ of Γ. Choose a balanced annulus. Trim each arm between its last inner contact before first outer exit and that exit. If s is the prefix-tail length removed at the selected connector contact and $t' > t$ is the selected future time, then

$$s \geq b_E, \quad \mathrm{len}_E(A_E) \leq s/\theta_E, \quad \mathrm{len}_D(A_D) \leq (t' - t)/\theta_D.$$

The shortened prefix plus connector has length at most

$$Kt + \eta - (1 - \epsilon/\theta_E)s + K(t' - t) \leq Kt' + \eta.$$

An arbitrarily small infimum error is absorbed in the strict negative term. Thus $t' \in T$, a contradiction. Hence $E(q_\delta, y) \leq KL + \eta$. Let $\delta \downarrow 0$ and use lower semicontinuity. □

A connector for two different metrics

The preceding deterministic router needs an input that is stronger than a same-metric connector. We prove that input here. In this subsection E is a canonical enhanced limit. The metric D is either another canonical enhanced limit on the same white noise, or a strong metric with the same ξ, Q. In the latter case E has already been given its exact scale covariance. No comparison between D and E is assumed. Write

$$R_r = \frac{c_r^E}{c_r^D}, \qquad s_r^E(z) = c_r^E e^{\xi h_r(z)}, \qquad s_r^D(z) = c_r^D e^{\xi h_r(z)}.$$

The raw profiles have the bounds (13.2); for a strong metric $c_r^D = r^{\xi Q}$. All results below are uniform over the canonical families for which those quantitative bounds are uniform.

We fix the closed test band $A(z,r) = \overline{B}_{2r}(z) \setminus B_r(z)$ and the open working annulus $V(z,r) = B_{3r}(z) \setminus \overline{B}_{r/2}(z)$. Every arm joins the two boundary components of $A(z,r)$ and stays in that closed band. Labels are tested in a slightly larger annulus, still separated from z. A further fixed enlargement is used for the determining field and contains the averaging sphere. These enlargements depend only on dimension and the kernel. Increasing the mesh depth never changes them. All later applications of the comparison router may use this fixed test template.

Lemma 13.5 (Coordinate charging for a metric-rectifiable arm)**.** *Let D be a separating Euclidean lower-semicontinuous metric whose finite rectifiable curves are Euclidean continuous. Let $P : [0,T] \to A(z,r)$ be parameterized by D-arclength and cross the test band. For some signed coordinate j, it crosses a fixed interval I of coordinate values of length $r/2$, with a margin $r/4$ from the initial and final coordinate values. At every sufficiently fine mesh $\rho_\ell = r2^{-\ell}$, let $v_\ell(a)$ be the center of the half-open mesh cell containing the first hit $x(a)$ of level $a \in I$.*

For each mesh center v, let $L_\ell(v)$ be a nonnegative number not exceeding the D-length of any crossing of $B_{2\rho_\ell}(v) \setminus \overline{B}_{\rho_\ell}(v)$. If $U_\ell(v) \le A_\ell L_\ell(v)$, then

$$\int_I U_\ell(v_\ell(a))\, da \le J_d \rho_\ell A_\ell T, \tag{13.4}$$

where J_d is deterministic. In particular, if $\sum_{\ell \ge L} \rho_\ell A_\ell < \infty$, there is an actual contact $x = x(a)$ on P such that

$$\sum_{\ell \ge L} U_\ell(v_\ell(a)) \le \frac{2J_d T}{|I|} \sum_{\ell \ge L} \rho_\ell A_\ell. \tag{13.5}$$

The result requires no Euclidean rectifiability and no E-length bound on P.

Proof. An outer endpoint has one signed coordinate of magnitude $2r$, whereas the corresponding coordinate of the inner endpoint is at most r. Reverse and change sign if necessary, and take the middle interval of length $r/2$. For a Euclidean-continuous coordinate function starting below I and ending above it, the first-hit time $\tau(a) = \inf\{t : P_j(t) = a\}$ is nondecreasing in a. It is therefore Borel, and $x(a) = P(\tau(a))$ is a Borel section. The set of levels selecting any one half-open cell is measurable and has Lebesgue measure at most ρ_ℓ.

For every selected cell choose one of its selected contacts. The contact lies within maximum-norm distance $\rho_\ell/2$ of its center. Both endpoint coordinate values of P lie a distance at least $r/4$ from I. For ℓ large enough, the portion after that contact must therefore leave $B_{2\rho_\ell}(v)$. Between its last hit of $\partial B_{\rho_\ell}(v)$ before its first hit of $\partial B_{2\rho_\ell}(v)$ is a shell-crossing subpath P_v. Its interior lies in $B_{2\rho_\ell}(v)$, and $L_\ell(v) \le \mathrm{len}_D(P_v)$. At any time of P, at most J_d of the selected subpaths can be present, since

their centers are distinct mesh centers within $2\rho_\ell$ of the same spatial point. Integration against D-arclength gives

$$\sum_{v \text{ selected}} L_\ell(v) \le \sum_v \operatorname{len}_D(P_v) \le J_d T.$$

Multiplying by the cell-section measure bound proves (13.4). Tonelli applies to the nonnegative series. Its integral is finite under the stated assumption, and at least one level has value at most twice its average. This proves (13.5). Only continuity of the coordinate function and the finite arclength measure were used; no change-of-variables formula for Euclidean arclength is involved. □

Lemma 13.6 (Appending an endpoint to a rational chain)**.** *Let E be Euclidean lower semicontinuous. Suppose $q_\ell \to x$ Euclideanly and there are E-rectifiable paths J_ℓ from q_ℓ to $q_{\ell+1}$ with $\operatorname{len}_E(J_\ell) \le u_\ell$ and $\sum_{\ell \ge L} u_\ell < \infty$. Suppose also that the supports of J_ℓ lie in balls about x with radii tending to zero. Then their concatenation extends to an E-continuous rectifiable path from q_L to the actual point x, with length at most $\sum_{\ell \ge L} u_\ell$.*

Proof. For $k > j$ the triangle inequality gives $E(q_j, q_k) \le \sum_{\ell=j}^{k-1} u_\ell$. Passing to the Euclidean lower limit at $q_k \to x$ yields

$$E(q_j, x) \le \sum_{\ell \ge j} u_\ell.$$

The same inequality, preceded by the remaining portion of J_j, controls every point of the concatenation near its final time. The tail sums tend to zero, so arclength parametrization extends continuously in E at the endpoint x. A finite partition of the extended curve has sum of E-distances at most the total sum of the u_ℓ: approximate a partition time at the endpoint by a late chain endpoint and use the displayed tail bound. Taking the supremum over partitions proves the length estimate. This also proves directly that x belongs to the finite E component. □

Lemma 13.7 (Two-metric upper and lower mesh events)**.** *Fix $0 < \tau < 1/10$. For every $p_0 > 0$, all sufficiently large deterministic L have the following property. With probability at least $1 - p_0$, there are finite E-paths joining successive mesh representatives, a finite connected E-core in $V(z, r)$, and nonnegative charges $U_\ell(v)$, $L^D_\ell(v)$ and $L^E_\ell(v)$ for $\ell \ge L$, such that:*

1. *Every $x \in A(z, r)$ has representatives $q_\ell(x)$ with*

$$|q_\ell(x) - x|_\infty \le C(1 + \ell)^2 \rho_\ell.$$

   The representatives at level L lie in the core. The link from $q_\ell(x)$ to $q_{\ell+1}(x)$ lies in the corresponding $C(1 + \ell)^2\rho_\ell$ neighborhood of x and has E-length at most $U_\ell(v_\ell(x))$.
2. *The quantity $L^D_\ell(v)$ is a lower bound for the D-length of every crossing of the shell about v in Lemma 13.5; $L^E_\ell(v)$ has the analogous meaning for E.*
3. *Uniformly in the cells,*

$$U_\ell(v) \le C_\tau R_r 2^{5\tau\ell} L^D_\ell(v), \qquad U_\ell(v) \le C_\tau 2^{5\tau\ell} L^E_\ell(v). \tag{13.6}$$

4. *The possible level-L representatives form one fixed finite deterministic menu. The core's total E-length, on this event, is at most a random variable $Z_L s^E_r(z)$ with*

$$\sup_{z,r} \mathbb{E}[Z_L^p] \le e^{C_p(L+1)} \quad \text{for every fixed } p < \infty. \tag{13.7}$$

All the conclusions are pathwise and simultaneous over x and over the tested arms. Rational z, r give rational menus.

Proof. Write $h = G_\sigma + F_\sigma$, where F_σ is the white-noise field below scale σ and G_σ is smooth. The pinning scalar is included in G_σ. Define the fine metrics by $E_\sigma = e^{-\xi G_\sigma} \cdot E$ and $D_\sigma = e^{-\xi G_\sigma} \cdot D$. These are actual weighted metrics, including before field determination. The weighted reconstruction and Proposition 12.17 give their fine-noise local kernels.

We first record the needed uniform one-box bounds. In a box of size comparable to σ, every fixed buffered point pair satisfies, for some $p > 0$,

$$\mathbb{E}[(E_\sigma(u, v; O)/c_\sigma^E)^p] \le C.$$

The corresponding D_σ and E_σ shell crossings X_σ^D, X_σ^E satisfy

$$\mathbb{E}[(c_\sigma^D/X_\sigma^D)^a] + \mathbb{E}[(c_\sigma^E/X_\sigma^E)^a] \le C_a \quad (a < \infty).$$

For canonical metrics these follow from the deterministic-point and inverse-crossing estimates already proved, followed by the Weyl rule. Indeed $h_\sigma(v) - G_\sigma(v)$ has uniformly bounded variance, and the oscillation of G_σ on a fixed σ-box has uniformly bounded exponential moments. Cauchy–Schwarz transfers the normalized full-field estimates to the fine metrics; decrease the positive point-moment order if needed. For an arbitrary strong D, the inverse estimate is Lemma 13.2, with the same transfer. No upper point bound for D is used.

At level ℓ use mesh $\rho_\ell = r2^{-\ell}$. Declare a vertex good when every signed nearest-neighbor fine E-edge has internal distance at most $Tc_{\rho_\ell}^E$ in its fixed buffered edge box. Choose T so that the bad probability is below the threshold of Lemma 3.1. Make nice labels by requiring the same finite set of neighboring level-$(\ell+1)$ good labels as in Proposition 3.3. The one-box moment bound allows the nice-label failure probability to be made as small as required.

The labels in one level, or in the two successive levels used for niceness, have a fixed dependence range in mesh units. For a canonical conditional sample this statement follows from independence of whole fine-noise/metric blocks in Proposition 12.17. For the two levels, use the larger fine-noise cutoff: the additional half-band needed for the finer metric remains in the same bounded spatial footprint. Independence does not require the metric itself to have been proved field-measurable. The lower labels introduced below need no independence at all.

Only labels in a fixed larger annulus are inspected. Declare all other sites artificially good. This is a finite modification of the all-good lattice and keeps the dependence estimates. The bad-cluster lemma, with a union bound over $O(2^{d\ell})$ sites, shows that except on an event of probability $C2^{d\ell}e^{-c(1+\ell)^2}$, every good or nice bad component has diameter at most $C(1+\ell)^2$ mesh steps. The same successive-level routing proof as in Proposition 3.3 then chooses $q_\ell(x)$ and joins $q_\ell(x)$ to $q_{\ell+1}(x)$ by at most $C(1+\ell)^{2d}$ good level-$(\ell+1)$ edges. All these edges lie in a polynomial mesh neighborhood of x. Choose L so large that those neighborhoods, and their edge buffers, lie in $V(z, r)$. Consequently no artificially good site is used in an actual path.

Choose each good edge path with length at most twice its threshold. There are only countably many such choices. Put $\sigma_\ell = \rho_{\ell+1}$. If v is a level-ℓ cell center, restoration of the smooth field bounds the total link cost by

$$U_\ell(v) = C(1+\ell)^{2d} c_{\sigma_\ell}^E \exp\{\xi G_{\sigma_\ell}(v) + \xi O_\ell\}, \tag{13.8}$$

where O_ℓ is the maximum oscillation of this smooth field on the polynomially enlarged mesh neighborhoods. The actual link images shrink to x uniformly in the choice of x.

For both metrics require every relevant fine shell of radius $\rho_\ell = 2\sigma_\ell$ to have crossing length at least $2^{-\tau\ell}c_{\sigma_\ell}^D$ or $2^{-\tau\ell}c_{\sigma_\ell}^E$, respectively. The one-box inverse moments and a union bound give, for arbitrary fixed $a < \infty$, a failure probability at most $C_a 2^{(d-a\tau)\ell}$. Thus the lower failures after level L

have probability tending to zero as $L \to \infty$, choosing $a > d/\tau$ first. On their complement, every full-field crossing pays the charge

$$L_\ell^D(v) = 2^{-\tau\ell} c_{\sigma_\ell}^D \exp\{\xi G_{\sigma_\ell}(v) - \xi O_\ell\}, \tag{13.9}$$

and the same formula with D replaced by E. This is applied only to actual subpaths inside the displayed shell, where Weyl scaling gives the indicated minimum-weight bound.

Finally impose $O_\ell \le \tau\ell \log 2/(2\xi)$ for every $\ell \ge L$. The local oscillation estimate (2.14) has a summable tail for the failures. It applies also to the all-scale coarse field here: after rescaling by r, its contribution from scales above one has bounded increment moments and the same kernel proof applies. Summing all three failure estimates and then choosing L proves the asserted probability.

The raw profile inequalities give

$$\frac{c_{\sigma_\ell}^E}{c_{\sigma_\ell}^D} \le C_\tau R_r 2^{2\tau(\ell+1)}.$$

The shared value $G_{\sigma_\ell}(v)$ cancels when (13.8) is divided by (13.9). The lower-label loss contributes $2^{\tau\ell}$, the two oscillation factors contribute at most $2^{\tau\ell}$, and $(1+\ell)^{2d} \le C_\tau 2^{\tau\ell}$. This proves (13.6); for the same-E ratio the profile ratio is one, so the same bound is valid.

It remains to construct the core, since the attachment chains alone do not connect their first vertices. The possible $q_L(x)$ belong to a deterministic finite collection of mesh vertices in a small neighborhood of $A(z,r)$. They lie in the same infinite good component. The lattice routing lemma joins them within $V(z,r)$: first join them by polygonal routes in a fixed intermediate annulus, and choose L so that the allowed bad-cluster detours are shorter than its clearance from ∂V. The union of these lattice routes is a finite connected graph in a deterministic bounded mesh region. Take a spanning tree, with at most $C2^{dL}$ edges, and replace its edges by the already chosen actual good E-paths. The sum of their lengths bounds the core's total mass, regardless of intersections among their images. It is at most

$$C2^{dL} c_{r2^{-L}}^E \exp\{\xi \sup_{V'} G_{r2^{-L}}\}.$$

Here V' is one fixed compact buffer. Dividing by $s_r^E(z)$ and using (13.2), the Gaussian increment bound for $G_{r2^{-L}} - h_r(z)$ gives (13.7). Define Z_L by this displayed bound on every outcome, even when the good-mesh event fails. The moment calculation is then unconditional. This finishes the construction of all the promised data. □

Proposition 13.8 (Mixed connector and seed event)**.** *For every $\epsilon > 0$ and $p_0 > 0$, there are $K < \infty$ and a finite deterministic menu $\mathcal{M}(z,r) \subset V(z,r)$ such that, with probability at least $1 - p_0$, the following hold simultaneously for all finite crossing arms:*

1. *If A_D is D-rectifiable and A_E is E-rectifiable, some actual contacts on these arms are joined in E within $V(z,r)$ by a path of length at most*

$$\epsilon \operatorname{len}_E(A_E) + KR_r \operatorname{len}_D(A_D). \tag{13.10}$$

2. *Every such A_D has an actual contact joined in E to some $q \in \mathcal{M}(z,r)$, inside $V(z,r)$, at cost at most $KR_r \operatorname{len}_D(A_D)$. Every menu point is E-nonsingular almost surely.*

The constants are independent of the deterministic center and scale. The event can be taken in the completion of the local metric-block sigma field. After any fixed further coarse deletion it has a version depending only on the corresponding fine local blocks, with arbitrarily high probability and changed deterministic constants. Its geometry is unchanged as ϵ and p_0 vary.

Proof. Apply Lemma 13.7. On its success event, Lemma 13.5 and (13.6) give a selected level on each D-arm whose upper-chain series has sum at most

$$a_L R_r \operatorname{len}_D(A_D), \qquad a_L = C_\tau \sum_{\ell \geq L} 2^{-(1-5\tau)\ell} \longrightarrow 0.$$

For an E-arm the corresponding bound is $a_L \operatorname{len}_E(A_E)$. The representatives approach the selected actual arm contact, and their link images shrink to it. Lemma 13.6 therefore makes each upper chain an actual E-path to that contact. This is the step that establishes finite E-access on a D-arm without assuming finite E-length of the arm itself.

Choose L so that the mesh failure probability is at most $p_0/3$ and $a_L \leq \min\{\epsilon/4, 1/4\}$. By (13.7), choose $H < \infty$ with $\mathbb{P}[Z_L > H] < p_0/3$. Choose $b > 0$ using the inverse crossing moment so that, except with probability $p_0/3$, every D crossing arm has length at least $b s_r^D(z)$. On their intersection, the core costs at most

$$H s_r^E(z) \leq (H/b) R_r \operatorname{len}_D(A_D).$$

Following the first attachment, a core route, and the reverse second attachment gives a length at most

$$(\epsilon/4) \operatorname{len}_E(A_E) + (1/4 + H/b) R_r \operatorname{len}_D(A_D).$$

Increase K to leave a fixed strict reserve. This proves (13.10), including arbitrarily small choices of internal infimum errors. The D attachment alone ends at a level-L menu point and proves the seed assertion. The menu is finite and deterministic, and deterministic-point finiteness makes each of its points nonsingular. Rational center and radius give the rational menu needed in the router.

We give the measurability detail, since the construction used auxiliary chains whereas annular concentration requires an intrinsic local event. Put $A = A(z, r)$ and $V = V(z, r)$, and choose increasing compact constraints $K_j \subset V$ whose interiors exhaust V. Retain the closed constrained metrics $D^{[A]}$, $E^{[A]}$, and $E^{[K_j]}$. These are lower-semicontinuous random distance functions, and are Borel functions of the full metric, by Lemma A.2. The closed constraint in this statement is essential: no lower semicontinuity of the open intrinsic metric of an arbitrary strong D is assumed.

These coordinates are local. Indeed, if a finite rectifiable path P lies in an open determining buffer U, then

$$\operatorname{len}_{D^U}(P) = \operatorname{len}_D(P).$$

The inequality $D \leq D^U$ gives one direction; in the other direction, each internal distance along a partition is bounded by the D-length of the corresponding subpath. Thus the constrained metrics above are determined by the actual internal metrics on U. To make the measurable-factor assertion explicit, condition on the local block and take two independent conditional copies of the full metrics, using Lemma A.5. Locality makes their actual internal metrics on U agree, hence makes all their constrained coordinates agree. Their conditional law in the countable epigraph coordinate space is therefore a point mass: for each coordinate, two independent bounded real copies agree only if their conditional variance is zero. Their conditional means give the required measurable coordinates. This proves local measurability also for an arbitrary strong competitor.

A curve in A is encoded in the Polish space of Euclidean-continuous maps on $[0, 1]$. Its $D^{[A]}$-Lipschitz bound N is expressed by the countable inequalities $D^{[A]}(P(s), P(t)) \leq N|s - t|$ for rational s, t; lower semicontinuity extends them to all times. Taking the union over $N \in \mathbb{N}$ covers arclength reparametrizations of all finite D-arms, since the preceding length identity also holds with a closed constraint containing the path. Conversely these Lipschitz tests imply finite D-length, and the

two lengths agree. The same statements hold for E-arms. Their lengths are suprema over rational partitions by Lemma A.9. Boundary endpoints and containment in the closed test band are closed Euclidean tests, so all arm conditions are Borel.

For two arms P, R the open-buffer connection cost is

$$\inf_j \inf_{s,t\in[0,1]} E^{[K_j]}(P(s), R(t)),$$

where a constrained distance is infinity when an endpoint is outside its constraint. Each inner infimum is Borel by Lemma A.2, and the countable outer infimum is Borel. Every finite connecting path has compact image in V and is contained in some K_j, so this expression is precisely the actual open-buffer cost. For a seed use a minimum over the finite menu and one time parameter. Hence a violating arm or pair of arms is the projection of a Borel set in the curve/metric product space. Lemma A.8 makes the failure event analytic and the universal success event measurable in the completion of the local block law. The strict reserve above ensures that the intrinsic success tests, with slightly smaller constants, imply actual paths with the displayed bounds. The event's probability is at least that of the auxiliary construction. No measurable choice of its infinite chains is required.

Finally delete the coarse field at a fixed scale γr, with γ small enough that the fine spatial footprint of the determining annulus still avoids its center. If the deleted field has oscillation at most M on the determining buffer, then all internal path lengths, after removing its one common constant value, change by factors in $[e^{-\xi M}, e^{\xi M}]$. An intrinsic mixed event with constants $(\epsilon e^{-2\xi M}, K)$ for the full metrics therefore implies the fine-block event with constants $(\epsilon, Ke^{2\xi M})$. The same statement holds for the seed. The coarse oscillation has a uniform Gaussian tail. First choose its quantile M, then apply the already proved high-probability event with the smaller coefficient. This proves an arbitrarily high probability for the fine intrinsic event. Conversely a fine event restores to a full event by the same calculation. Proposition 12.17 shows that these fine events depend on disjoint whole blocks for separated annuli, including conditional canonical samples. For a strong competitor its actual intrinsic locality gives the same dependence. Thus Lemma 2.9 applies before field determination as well as afterward. □

In particular Proposition 13.8 supplies precisely the hypotheses (13.3) and the seed hypothesis of Proposition 13.4. The proof uses the D-arclength only to charge disjoint shell visits. It never applies a same-metric E connector to a path whose E-rectifiability has not been established.

For two canonical weak assignments the same argument applies on any deterministic positive-density set of scales on which R_r is bounded. If for every integer j some arbitrarily large logarithmic blocks contained a fixed positive proportion of scales with $R_r \le 1/j$, the prefix proposition would give $E \le CD/j$ for every j, contradicting separation at a fixed rational pair. Interchanging the metrics yields the opposite bound. Therefore the raw profiles are comparable on all but an arbitrarily small fraction of each sufficiently large block, and the prefix proposition gives a deterministic two-sided metric comparison. Crossing quantiles then make the profiles comparable at every radius.

Corollary 13.9. *Two canonical weak assignments coupled to the same field are deterministically bi-Lipschitz. A canonical metric with exact covariance satisfies $E \le CD$ for every strong competitor with the same (ξ, Q), for a deterministic $C < \infty$.*

13.3 Removal of residual randomness

The comparison and conditional-resampling strategy for determination by the field comes from the local-metric framework of Gwynne and Miller [10]. We give the argument for the extended metrics

and white-noise conditional kernels constructed here.

Proposition 13.10. *Every enhanced canonical subsequential metric is a measurable function of h.*

Proof. Retain the full driving white noise W. In finitely many disjoint determining space–scale regions, the coarse-subtracted prelimit internal arrays are functions of the corresponding independent fine noises. Retain these noises and the arrays as whole blocks. Their product laws pass to the enhanced limit; buffered reconstruction identifies the actual coarse-subtracted internal metrics. After disintegration, local internal metrics in disjoint open sets have product conditional kernels given W. After coarse subtraction, the local kernel in each set depends only on its determining noise. Restoring the known smooth correction g_s applies a deterministic Weyl transform given W, preserving conditional independence. To pass from buffered rational sets to arbitrary disjoint opens, exhaust each by finitely many rational subdomains, take a sufficiently fine coarse cutoff to separate their noise footprints, and exhaust.

Take two independent enhanced conditional samples given W. Their common profile has raw ratio one, so Corollary 13.9 gives a deterministic comparison $M^{-1}D \le E \le MD$. For a rational pair, the conditional essential supremum of its positive finite distance is therefore at most M times the conditional essential infimum. Thus the same comparison holds under every coupling with those conditional marginals, first on the rational array and then everywhere by reconstruction.

Let $X = D(p, q)$ for fixed distinct rational p, q. A constant-speed geodesic is coded by $P \in C([0, 1], \mathbb{R}^d)$, the endpoint constraints, and

$$D(P(s), P(t)) \le |s - t|X \qquad (s, t \in \mathbb{Q} \cap [0, 1]).$$

These are Borel tests by Lemma A.9. Lower semicontinuity extends the inequalities to all times, and the endpoint distance X forces equality by the triangle inequality. The sections are nonempty by Proposition 12.7. Lemma A.8 therefore selects such a geodesic measurably in the completed law of (W, D), before any grid is introduced. Its occupation measure is $\mu = XP_* \operatorname{Leb}_{[0,1]}$. The geodesic is injective, so this measure is nonatomic and has compact Euclidean support.

Independently choose a uniform translation $t \in [0, \delta)^d$ and use the open cells $U_i(t) = t + \delta(i + (0, 1)^d)$, $i \in \mathbb{Z}^d$. Their internal-coordinate maps are jointly measurable in t and D. Indeed, for fixed global rational p', q', encode the cell coordinate by

$$\inf_{\substack{V \text{ rational polycubical} \\ p', q' \in V,\ \overline{V} \subset U_i(t)}} D(p', q'; V),$$

with empty infimum $+\infty$. Compact containment in the translated cell is a finite set of strict coordinate inequalities. Every tested intrinsic distance is Borel by Lemma A.2, and the infimum is countable. The formula equals the actual cell distance by compactness of a finite path image. The rational envelope reconstructs the internal metric from these coordinates. Proposition 12.17 and Fubini therefore give product conditional kernels for the cells given (W, t), for almost every (W, t).

For each fixed selected path, Fubini's theorem makes its mass on grid boundaries zero for almost every translation: each spatial point belongs to a randomly translated grid hyperplane with probability zero. Conditional supportwise comparison makes this path rectifiable in every conditional sample, and makes its boundary occupation zero there as well. Equal cell internal metrics give equal length measures off the boundaries and hence determine endpoint distances. Applying the argument to each of the two samples proves that the global metric is reconstructed from these local coordinates and W.

Resample the coordinate of the open cell U_i, using its independent conditional kernel. Following the old, previously selected geodesic and using supportwise comparison gives

$$(X^{(i)} - X)_+ \le (M-1)\mu(U_i).$$

Apply the conditional version of Lemma 2.4, in the one-sided resampling form (2.12), to the independent cell coordinates. It gives

$$\operatorname{Var}(X \mid W, \text{grid}) \le \sum_i \mathbb{E}[(X^{(i)} - X)_+^2 \mid W, \text{grid}] \le (M-1)^2 \mathbb{E}\left[X \max_i \mu(U_i) \mid W, \text{grid}\right].$$

The countable-coordinate assertion is included in that lemma. For almost every W, the conditional essential range of X is bounded, by the conditional comparison already proved, so this conditional calculation is integrable.

Now average over the independent grid translation before taking a mesh limit. The conditional law of X is unaffected by that translation, whereas every cell lies in a ball of radius $\sqrt{d}\,\delta$. Hence

$$\operatorname{Var}(X \mid W) \le (M-1)^2 \mathbb{E}\left[X \sup_{z \in \mathbb{R}^d} \mu(B_{\sqrt{d}\,\delta}(z)) \mid W\right].$$

For each finite nonatomic measure on a compact set, the displayed supremum tends to zero: otherwise shrinking balls with uniformly positive mass have centers converging to an atom. It is bounded by $\mu(\mathbb{R}^d) = X$. Conditional dominated convergence therefore gives zero conditional variance. Countability and rational reconstruction make D a function of W.

Finally, Lemma A.7 shows that convergence in law of (W, D_n) to the graph $(W, F(W))$ implies convergence of D_n to $F(W)$ in probability on the original noise space. Each D_n is h-measurable; an almost surely convergent further subsequence shows that its probability limit is h-measurable. No additional white-noise dependence remains. □

Proposition 13.11. *The family of normalized deterministic recenterings and rescalings of canonical weak assignments is closed under enhanced subsequential limits whenever its quantitative constants are uniform. A limiting pair has rational evaluation consistency, buffered finite-path recovery, simultaneous Weyl scaling, local conditional white-noise kernels, and field measurability.*

Proof. For each member of a sequence, use its defining LFPP subsequence. Proposition 13.10 and the graph argument in its proof give convergence in probability on the original field space, jointly with each finite list of consistent global rational point evaluations. Choose one sufficiently fine deterministic LFPP approximation to the jth assignment so that, with probability at least $1 - 1/j$, both its global metric coordinate and its first j global rational point coordinates are within $1/j$. For a pair use two cutoffs on the same field. These diagonal LFPP pairs have exactly the same subsequential limits as the original pairs and their global rational arrays.

All construction estimates are uniform for these normalized rescalings. Retain the raw profile on rational radii, lower barriers, coarse fields, fine noises, rational arrays, upper core bounds and finite good-annulus certificates, and the shifted buffered internal arrays. The raw ratio bounds keep the limiting profile positive and finite. Tightness and the original rational-consistency argument apply to this diagonal LFPP sequence. They give a separating metric with the asserted rational evaluations.

Also retain the actual certificates and arrays of the original weak assignments, not only those of the diagonal approximants. Integer core-depth coordinates are tight. Conditional on their finite

values the numbers of path and vertex coordinates are fixed, and lower barriers make bounded certificate paths compact. The original sequence now has convergent rational distances and uniform lower barriers. Partitioning a finite limiting path into short pieces, approximating its vertices intrinsically by rational points, and joining these vertices in the approximating metric recovers that path. Exit barriers confine each link to its prescribed open neighborhood. This proves recovery for the original weak-metric sequence. Consequently actual retained all-arm successes pass to the limit by recovering the prescribed arms and extracting their connectors.

One may also retain the completed-measurable compact incoming-arm classes of Lemma 13.21 as compact hyperspace coordinates; their Hausdorff limits may be proper subsets of the limiting admissible classes. This retention uses compactness only and preserves every approximating arm limit.

Repeat the weighted buffered reconstruction with this recovery. It identifies actual shifted internal metrics and passes whole-block independence of the retained fine noises and arrays to the limit. Conditional replicas of the limiting metric have the same raw profile; the prefix comparison and Proposition 13.10 remove their residual noise. Each limiting metric and its actual intrinsic coordinates are functions of the same field. This argument uses comparison and construction estimates, but no scalar uniqueness assertion. □

13.4 Contact geometry and span estimates

Lemma 13.12. *Suppose that every point has arbitrarily many disjoint shrinking annuli on which all pairs of crossing arms have an E-connector costing at most C times their total length. If x is nonsingular, O is an open neighborhood of x, and $T < \infty$, then, for every $\eta > 0$, some $B_\rho(x) \subset B_R(x) \Subset O$ has this property: every finite E-path of length at most T meeting $B_\rho(x)$ and leaving $B_R(x)$ has an actual point joined to x inside O at cost below η.*

Proof. Take a nonconstant reference germ from x of length δ with $(1 + C)\delta < \eta/2$, confined in the desired neighborhood. Choose N disjoint successful annuli surrounding x, crossed by that germ, with $CT/N < \eta/2$. Take ρ inside all their inner holes. The arbitrary path crosses all annuli; its crossing-arm masses sum to at most T. At one annulus its arm costs at most T/N. Connect it to the reference arm and prepend the reference initial segment. The total cost is below $(1 + C)\delta + CT/N < \eta$. □

The same proof gives the useful relative version

$$E(x, P; O) \le \eta + \epsilon \operatorname{len}_E(P) \tag{13.11}$$

for all paths meeting one sufficiently small neighborhood of x and leaving O. A finite cover gives the corresponding assertion for any compact set of nonsingular points. All certificates use finitely many annuli and finite reference germs. They therefore persist under enhanced convergence, with every fixed strict relaxation. In a closed cubical exterior use the clipped annular cores; a short rational reference arm in the open exterior is obtained from the finite connected core in a wide successful clipped annulus. Its cost is bounded by a fixed multiple of the crossing cost of an initial closed-exterior germ. It can therefore be made arbitrarily small and has ordinary open-buffer recovery.

Lemma 13.13. *For $0 < \lambda \le \Lambda < \infty$ there is $\beta > 0$ such that geodesics for all continuous weights $w \in [\lambda, \Lambda]$ satisfy a common positive-power bound for first-to-last visit spans of small balls away*

from their endpoints. For fixed rational endpoints the bound also holds near those endpoints, with an almost surely finite random constant.

Proof. Let $T_{\mathrm{in}}, T_{\mathrm{out}}$ be the first-to-last weighted visit spans of nested inner and outer regions of a successful annulus. The incoming and outgoing arms have total original length at most $(T_{\mathrm{out}} - T_{\mathrm{in}})/\lambda$. An E-connector of weighted cost at most $(\Lambda C/\lambda)(T_{\mathrm{out}} - T_{\mathrm{in}})$ shortcuts the intervening geodesic. Thus

$$T_{\mathrm{in}} \le \frac{\Lambda C}{\lambda + \Lambda C} T_{\mathrm{out}}.$$

There are linearly many successful nested annuli between radii comparable to r and $\sqrt{r}$, simultaneously around every point. Iterating gives $T(B_r(x)) \le ALr^{\beta}$ when both endpoints are outside a fixed multiple of $B_{\sqrt{r}}(x)$.

At a fixed rational endpoint p, deterministic-point positive moments and a finite core menu imply, by a grid union bound and Borel–Cantelli, that every candidate vertex v on the relevant scale-2^{-m} core satisfies $E(p, v) \le 2^{-\alpha m}$ eventually, for some $\alpha > 0$. Attaching the outgoing arm to such a vertex gives the same contraction with an additive $\Lambda 2^{-\alpha m}$. Iteration supplies a positive power near p. Combine the two endpoint estimates with the interior estimate. Coercivity confines all the weighted geodesics to one random Euclidean compact set. □

Lemma 13.14. *Given a target density below one and a prescribed exponential failure rate, one can choose a small relative core coefficient and a deterministic $A < \infty$ so that the corresponding span-envelope bounds hold at that density of scales at every point. These bounds hold uniformly for geodesics whose annular arms and connector buffers are unchanged by other perturbations.*

Proof. If the core coefficient is ϵ, geodesicity gives

$$T_j \le qT_{j+1} + (1-q)H_j, \qquad q = \frac{\epsilon}{1+\epsilon}, \quad H_j = s_j Y_j.$$

Hence an inner span is at most $q^m L + \sum_{j<m} q^j H_j$. Choose q after a desired finite moment order so that Gaussian unit-ratio moments and (13.2) make

$$\left\| \sum_{j \ge 0} q^j H_{k-j}/s_k \right\|_p < \infty$$

uniformly in k. To obtain density rather than just a marginal bound, write the envelope recursively as $A_k = Y_k + e^{-a} r_k A_{k-1}$, where $a = -\log q$ and $r_k = s_{k-1}/s_k$. Set

$$B_k = \log(1 + A_k), \quad W_k = \log(2(1 + Y_k)), \quad V_k = \log r_k + \log 2.$$

For $C_k = (B_k - b)_+$ and $Z_k = (W_k - b)_+ + (V_k)_+$,

$$C_k \le Z_k + (C_{k-1} - a)_+, \qquad a\#\{k \le n : C_{k-1} \ge a\} \le C_0 + \sum_{k \le n} Z_k.$$

The annular Gaussian block inequality, with a fixed finite coloring, bounds exponential moments of sums of $(V_k)_+$ by K_t^n, independently of a. Choose a large enough for the desired drift and failure rate. Only then choose b: every fixed moment of Y_k implies that the relevant exponential moments of $(W_k - b)_+$ tend uniformly to one. Exponential Markov and the initial envelope moment make the fraction with $B_k > b + a$ arbitrarily small, with failure Ce^{-Bn}. The annular physical geometry

was fixed before these choices. A union bound over $O(L_*^{dn})$ finest-grid chains gives the simultaneous all-point assertion after increasing B.

For separated perturbation sites, iterate outward from a working ball to the mesoscopic separation scale. All these annular arms and cores are untouched. The terminal term $q^m L$ is smaller than any fixed power of the microscopic-to-macroscopic ratio by choosing q sufficiently small first. This choice is independent of the eventual discount depth. □

Eligible contacts of bounded original geodesics with a fixed cubical collector have a compact intrinsic contact class. Indeed short incoming arms attach to finite rational cores. A Euclidean limit of contacts is therefore nonsingular. Applying Lemma 13.12 to that limit and using the occupation contraction on untouched collars proves intrinsic convergence of a subsequence. Quantitatively, two nearby contacts satisfy

$$E(x, y; \text{closed exterior collar}) \le K|x - y|^{\beta' s}. \tag{13.12}$$

To prove this, take one common successful annulus at a scale between $|x-y|^{1/2}$ and $|x-y|^{1/4}$. Both incoming arms have small length by Lemma 13.13; the clipped connector and the retained terminal portions give (13.12). This applies only to actual path contacts, not to all points on the boundary.

13.5 A protected star

Lemma 13.15. *Let $O \subset \mathbb{R}^d$ be connected and open, and let A be a finite union of compact sets each homeomorphic to a closed subset of a compact interval. Then $O \setminus A$ is connected. The same holds when $A \cap O$ is replaced by any relatively closed subset of itself; the resulting complement is connected and open. In particular, finite unions of compact simple arcs may be removed. The arcs need not be contained in O and may meet its boundary.*

Proof. This is Theorem B.8, including its relatively closed subset clause. Its proof by finite cochains and compact-cover limits applies to arbitrary embeddings and requires no Euclidean length or dimension estimate for their images. □

Lemma 13.16. *Let V be a rational cubical collector, whose two-sided boundary collar lies compactly in the local data region U. Let $x_1, \dots, x_N$ be distinct eligible contacts on ∂V, and fix an access tolerance $\eta > 0$. There are exterior rational ports q_i, interior rational ports p_i, rational boundary ports y_i in open faces of ∂V, and finite simple access arcs F_i from q_i to p_i, such that*

$$F_i \cap \partial V = \{y_i\}, \qquad F_i \setminus \{y_i\} = F_i^{\text{out}} \dot\cup F_i^{\text{in}},$$

where the two pieces lie respectively in the open exterior and open interior of V. The F_i have pairwise disjoint images and E-length below η. Each x_i has closed-exterior access to q_i of cost below η. All these paths lie in arbitrarily small chosen disjoint neighborhoods of the x_i, and hence avoid any fixed interior trial boxes. Appending them to the contact net preserves the prescribed tolerance after an arbitrarily small increase.

Proof. We first recall why a deterministic boundary point has strictly one-sided access, a stronger property than unconstrained rational density. Take a rational point y on an open face and rational points $v_j = y + 2^{-j} r v$ in its inward cone. Each adjacent pair can be joined in a buffered open box of diameter $O(2^{-j} r)$ lying strictly on that side of the boundary. Write its length as L_j.

Deterministic-point moment bounds, (13.2), the Gaussian increment bound and Cauchy–Schwarz give, for sufficiently small $p \in (0,1)$ and some $\eta' < \xi Q$,

$$\mathbb{E}[(L_j/s_r(z))^p] \le C\, 2^{-j\{p(\xi Q-\eta')-p^2\xi^2\}}, \qquad |y-z| \le Cr.$$

Choose p so that the exponent is positive. Since $(\sum_j L_j)^p \le \sum_j L_j^p$, the infinite concatenation has finite length and its tail length tends to zero. Lower semicontinuity at the Euclidean limit y gives metric continuity there. Every point of the concatenation other than its endpoint lies strictly on the chosen side. The same construction in an outward cone gives an exterior path. Finite additional buffered rational links connect the first vertices to any prescribed finite local core menu. This is the boundary-port construction of Lemma 12.13. It uses $Q > 0$ only through the positive exponent above.

Choose disjoint neighborhoods of the x_i, all in a thin collector collar disjoint from the trial boxes. In each shrinking clipped annular template choose finitely many rational boundary ports on nearby open faces, with their inward and outward cone boxes. Near an edge or corner, move the port a distance comparable to the annular scale along one face, away from its other faces. A finite face-and-corner menu supplies uniform shapes and constants. The ports and all their cone boxes remain in the annular data buffer; the infinite tails are around the ports, away from the annulus center.

Augment the clipped exterior-core event by an upper bound Hs_r on the finitely many strict-side port paths and their exterior core links, and by a lower crossing bound as_r. The moment estimate makes H finite at any prescribed probability; choose $a > 0$ small afterward. Annular density applies to these local tests, including their arbitrarily fine data near the deterministic ports. Thus every eligible closed-exterior germ crosses arbitrarily small successful clipped annuli. If its initial arm has length ℓ, the lower crossing gives $s_r \le \ell/a$. The core connection, the retained part of the germ, and the added port links have total length at most a fixed multiple of ℓ. The initial arm length tends to zero. Consequently x_i has arbitrarily short closed-exterior access to an exterior rational q_i, while q_i is joined to a rational boundary port y_i, and y_i to an interior rational p_i, by paths of arbitrarily small total cost. Every exterior or interior port path stays strictly on its side except at y_i.

Erase loops separately from the two side paths: a compact rectifiable image, with its intrinsic length metric, contains a minimizing simple arc between its endpoints. The resulting two arcs are disjoint except at y_i, because their open sides are disjoint. Their concatenation is the required simple F_i. The chosen neighborhoods give mutual disjointness and disjointness from all marks. Contact-to-x_i connections from (13.12) remain in the closed exterior and can be followed by the new access to q_i.

For later compactness arguments retain the whole finite arc F_i as a rational-endpoint path in the open buffer U, together with its positive clearance from all fixed wall regions. Its strict one-boundary-contact property is used to construct the limiting tree. Recovered approximating paths need only retain the wall clearance and length bounds, which open-buffer recovery provides; no unsupported assertion about recovery inside a closed half-space is needed. □

Lemma 13.17. *Let $O = V \cap U$ be the connected collector interior in the local data region; in the annular construction it is a cubical annulus. Suppose each of finitely many disjoint trial boxes compactly contained in O, away from its outer collar, contains a simple finite E-arc R_i with rational endpoints and a protected middle. Let p_i, F_i be as in Lemma 13.16. There is a finite embedded E-rectifiable tree $T \Subset O$ through the ports and all the marked arcs such that every route from p_i to the central part crosses its protected middle. Moreover, for any $\eta > 0$, T meets $F = \bigcup_i F_i$ only*

on initial leaf stubs of total length below η. After slightly extending and deleting those stubs, the remaining central tree is compact and disjoint from F.

Proof. Inside O, each access arc is a pendant slit with its sole boundary root y_i on the outer boundary. There are no inner-boundary contacts, since the whole access collar lies compactly in U. Near p_i, every point of F lies on the single arc F_i and connects to p_i along an arbitrarily short subarc. Choose disjoint such neighborhoods, disjoint also from all marks.

Choose a short open terminal parameter stub of F_i at p_i, entirely contained in its chosen neighborhood. For each leaf, connect p_i to the left rational endpoint of R_i, forbidding the compact initial part of F_i complementary to that stub, all other access arcs, all other marked arcs, the protected and right part of R_i, and previously constructed leaf attachments. The forbidden set is a finite union of compact simple arcs and closed subarcs. Lemma 13.15 therefore makes its complement in O connected and open, including when an access arc meets the outer boundary.

Local rational rectifiable connectivity in this open complement gives a finite path. Read it from the mark toward F, stop at its first F contact, and retain the portion after its last R_i contact before that time. Erase loops and append the short terminal F_i stub to p_i. The forbidden set forces the F contact into the allowed terminal stub. The mark contact lies in its unprotected left stub; the retained joining arc has no other contact with R_i, and that stub can be followed toward the protected middle. The attachments are disjoint and do not bypass a protected middle.

For central wiring forbid all of F, all leaf attachments, and the protected-left portions of the marks, leaving their right stubs available. Their complement in O is connected by the same lemma. This forbidden set stays fixed during central wiring; the growing central tree is allowed to be met and is never added to the forbidden set. Join the right rational marked endpoints to a rational hub in the connected open complement. Take the portion between the last contact with one existing component and the first contact with another whenever a new connection is added, and erase its loops. This constructs a finite embedded tree while avoiding all forbidden parts. Every leaf-to-central route traverses its own marked middle. The only intersections with F are the previously chosen short terminal stubs. Extending each by a tiny positive amount in the leaf attachment and then deleting it leaves a compact central subtree disjoint from F. All routes and stubs are in O, so the required strict collector and local-data clearances hold. □

Lemma 13.18. *An embedded finite tree $T \subset \mathbb{R}^d$ admits a continuous ambient retraction $\pi : \mathbb{R}^d \to T$.*

Proof. Apply Lemma B.3 to the identity on the closed set T. Its proof includes the scalar extension construction and the continuity of the tree interpolation. □

13.6 The signed local contract

Suppose $c_*D \le E \le CD$, with $0 < c_* < C$. Fix collector and working regions $V \Subset W$, a local annular data region U, and a protected star $T \Subset V \cap U$. Write L_T for its total E-length. Each protected middle $R = [u_R, v_R]$ has

$$\operatorname{len}_E(R) < c_0 a_R, \qquad D(u_R, v_R) > a_R > 0, \qquad c_0 < c_1 < C. \tag{13.13}$$

All exterior access paths join eligible original collector contacts to the leaf ports at cost at most η_0 and lie outside V or in F. Fix a weighted visit-span bound A, insertion charge $a > 0$, and inactive restoration budget $\epsilon_0 > 0$, with $2\eta_0 < c_* a/4$.

Proposition 13.19. *There is a smooth signed f supported in $V \cap U$, and constants $b > 0$, $A_{\mathrm{rest}} < \infty$, with the following properties.*

1. *Every eligible pair of original collector contacts has a connection in $e^{-\xi f} \cdot D$ of cost below $a/2$. Retained pieces outside the collector have unchanged cost.*
2. *Every finite path with endpoints outside W and weighted first-to-last W span at most A either has a path interval of weighted D-length above b, contained in U, with an $e^{-\xi f} \cdot E$ shortcut of ratio below c_1, or all its positive-discount visits can be deleted at total original E-cost below ϵ_0.*
3. *Without the inactive assumption all positive-discount visits can be deleted at original E-cost at most A_{rest}.*

Other perturbations supported in disjoint working regions do not affect these conclusions. Restoration remains valid after inserting canonical paths in those other regions, using the original target contacts.

Proof. Take disjoint initial leaf stubs K_i of total length much smaller than ϵ_0, before every branch and mark. The central tree after their deletion is disjoint from F. Around it choose a fixed thin compact body $L \Subset V \cap U$ disjoint from F, with an inner body L_{in}. Two distance levels around the central tree define an outer wall. For each protected endpoint choose a short tree cut arc and use the retraction of Lemma 13.18 to extend its monotone cut coordinate to the ambient body. The $-1/2$ and $1/2$ level sets are disjoint compact sets. Their D distances, and the distance between the outer-wall levels, have a common positive lower bound a_{wall}.

Choose a wall multiplier $B > 4A/a_{\mathrm{wall}} + 2$ before shrinking any holes. Protected cut neighborhoods O_u, O_v can be made disjoint with $D(O_u, O_v) > a'_R$ and

$$\mathrm{len}_E(R) + 2\eta_{\mathrm{cut}} < c'_1 a'_R, \qquad c_0 < c'_1 < c_1.$$

Choose $\lambda > 0$ so small that $(2\eta_0 + \lambda L_T)/c_* < a/2$. Then choose $\delta > 0$ so small that, for every mark,

$$2B\delta + \lambda\big(\mathrm{len}_E(R) + 2\eta_{\mathrm{cut}}\big) < c'_1 \lambda a'_R, \qquad \sum_i (2\delta + \mathrm{len}_E\, K_i) < \epsilon_0/2.$$

Lemma 13.12, with path-length bound $2CA/\lambda$, gives tiny holes about the leaf and cut arcs and compact carriers of actual path contacts having original E-access below δ to those arcs. Choose a sufficiently thin neighborhood N of T so that its intersections with the fixed walls lie only in these holes, and so that $N \subset L_{\mathrm{in}} \cup \bigcup_i H_i$, where H_i are the leaf holes. Choose a smooth $w_0 \in [1, B]$, equal to B on all wall slabs and to one near F and outside $V \cap U$. With a cutoff χ equal to one near T and supported in N, put

$$w = \lambda\chi + (1 - \chi)w_0, \qquad f = -\xi^{-1} \log w.$$

Then $w = \lambda$ on T, $w \le 1$ on F, and $w = B$ on every wall outside its holes. The compact carrier $\operatorname{supp}\chi$ contains every positive-discount visit.

Insertion follows from the two access paths and the tree route, using $D \le E/c_*$. The entire signed support lies in the collector, so retained original pieces incur no increase. After earlier deletions in disjoint regions, every newly exposed first or last collector contact still has an original incoming or outgoing exterior arm and is eligible.

A weighted path of span at most A cannot cross a wall outside its holes. Any excursion in L from one leaf gate to another therefore visits the two holes of a protected mark in tree order. This order follows from the retraction: for two successive cut coordinates Φ_u, Φ_v, one has $\Phi_u < 1 \Rightarrow \Phi_v = -1$ and $\Phi_v > -1 \Rightarrow \Phi_u = 1$. The finite access certificates supply two ordered actual path contacts

with original E-access below δ to the cut arcs. Their intervening weighted D-length exceeds $\lambda a'_R$, whereas the route through the protected middle costs less than $c'_1\lambda a'_R$. This is an active shortcut. Choose $b < \lambda \min_R a'_R/2$.

Consequently an inactive excursion through the body uses only one leaf gate. Carrier visits outside the inner body already lie in a leaf hole. For each gate retain the compact set of original times whose positions have access at most δ to its short leaf arc. Every carrier visit assigned to that gate has a preceding and following such contact. Process the earliest remaining carrier visit, and replace from a preceding contact to the last original contact at that gate. Its cost is at most $2\delta + \operatorname{len}_E K_i$, and the gate is never needed again. Disjoint gate buffers keep inserted paths out of all other gates. Thus total inactive cost is below ϵ_0. In the active case, all needed contacts connect through the full tree at cost at most $2\delta + L_T$; one replacement between first and last necessary contacts proves the general bound.

All contacts and access arms lie in the target working region. Inserted images in other disjoint regions miss them. Deletion endpoints are on retained original pieces, so an earlier inserted path is either retained whole or swallowed whole. Fresh bracketing contacts exist after each prior deletion because the current original point lies outside the next target working region. This proves the final assertion without applying a competitor-length estimate to any inserted canonical path. □

13.7 Uniform certificates and two-sided rigidity

We give the stability argument used to select deterministic perturbations. The distinction between a certificate and the event which it certifies is useful. A certificate is finite geometric and metric data. The working event for a fixed smooth perturbation is its universal path contract. We prove measurability of that event directly; we do not project a universal path condition through a further existential quantifier.

In this subsection all metrics are restricted to fixed open buffers when locality is at issue. A compact constraint is written in square brackets. Suppose

$$0 < c_* \le C < \infty, \qquad c_* D \le E \le C D. \tag{13.14}$$

The same inequalities hold for path lengths and actual intrinsic metrics: apply the inequalities to every finite partition of a path, then take the supremum over partitions and the infimum over paths. In particular the two metrics have the same rectifiable paths and the same nonsingular points. Local rational density for E gives local rational density for D. The germ-access proof of intrinsic lower semicontinuity applies to D as well: use the E access paths and divide their costs by c_*. Consequently the actual internal metrics used below have the same Borel rational reconstruction as in Proposition 12.11.

Definition 13.20 (Incoming-arm codes)**.** Fix a compact cubical collar $K \Subset U$, closed subsets $A_0, A_1 \subset K$ at positive Euclidean distance, a closed cubical exterior F, constants $0 < \ell_- < \ell_+ < \infty$ and $H < \infty$, and positive exponents α, β. An incoming-arm code is a pair (P, ℓ), where $P \in C([0,1], K)$ and $\ell \in [\ell_-, \ell_+]$, satisfying

$$P(0) \in A_0 \cap \partial F, \quad P(1) \in A_1, \quad P([0,1]) \subset F, \tag{13.15}$$

$$|P(t) - P(s)| \le H|t-s|^\alpha, \quad \ell|t-s| \le H|P(t) - P(s)|^\beta, \qquad s,t \in [0,1], \tag{13.16}$$

$$D^{[K]}(P(s), P(t)) \le \ell|t-s|, \qquad s,t \in [0,1]. \tag{13.17}$$

It is harmless to replace $D^{[K]}$ in the last display by the length of the indicated subpath; for compactness we use the weaker displayed condition. We use finite unions of these classes for the

finitely many collar orientations. Constants are fixed on a common high-probability environmental event. The lower crossing bound supplies ℓ_-, compact lower barriers supply the first modulus, and the occupation estimate supplies the second. Thus the original incoming geodesic arms needed in the counting argument belong to these classes, after arclength reparameterization. We include their time reversals when outgoing arms are needed.

The second inequality in (13.16) will only be used as a bound for a short terminal part. For example, before the first exit from $B_R(P(0))$, the D-time is at most $H(2R)^\beta$. By (13.14), the E-length of that part is at most $CH(2R)^\beta$. One may equivalently impose this bound as a closed occupation condition, using slightly larger rational balls. The inverse modulus formulation makes the closedness transparent.

Lemma 13.21 (Compact arm families and their enhancement)**.** *For fixed data in Definition 13.20, the incoming-arm class is compact in the uniform-path topology times the usual topology on $[\ell_-, \ell_+]$. Suppose (D_n, E_n) converge in an enhanced coupling, with uniform compact inverse moduli and* (13.14)*, and let $\mathcal{K}_n$ be measurable compact subfamilies of the corresponding incoming-arm classes. After a further extraction the compact sets $\mathcal{K}_n$ converge in Hausdorff distance to a compact set $\mathcal{K}$. Every member of $\mathcal{K}$ is an admissible limiting incoming arm. Every member of $\mathcal{K}$ has approximating members of $\mathcal{K}_n$, and every limit of members of $\mathcal{K}_n$ belongs to $\mathcal{K}$.*

The assertion holds simultaneously for a countable list of collars and rational quantitative parameters. Empty classes can be retained by adding one isolated symbol to the compact hyperspace. In particular, the full admissible classes are measurable random hyperspace coordinates after completion of the underlying probability space.

Proof. The first modulus and the compact range give a common compact ambient path space by Lemma A.9. The position and modulus conditions are closed. The last condition is closed under uniform path limits by Lemma A.2, first at rational times and then at every time. This proves compactness for one metric.

For completeness, the nonempty compact subsets of a compact metric space form a compact space in Hausdorff distance. To see total boundedness, choose a finite 2^{-j}-net of the underlying space; a compact set is approximated within 2^{1-j} by a nonempty subset of this finite net. For a Hausdorff Cauchy sequence (K_n), take the set of all limits of sequences $x_{n_k} \in K_{n_k}$. Compactness of the ambient space shows that this set is nonempty and compact, and the Cauchy property shows that K_n converge to it. Total boundedness and completeness give the claimed compactness.

The lower-limit inequality for D_n, and compact path extraction for the constraints $[K]$, imply that a uniform limit of admissible arms satisfies (13.17) for D. All other conditions pass by continuity. The two approximation assertions are the definition of Hausdorff convergence. A diagonal extraction handles countably many classes. In particular, retaining these compact classes is an additional compact coordinate, not an additional stochastic estimate.

Here is the measurability detail for the full admissible classes used below. Their graph in the metric and ambient path coordinates is Borel: position and modulus conditions are closed, and (13.17) is the countable intersection of its rational-time Borel tests, by Lemma A.2. For every open subset O of the compact ambient path space, the event that the class meets O is an analytic projection. It is universally measurable by Lemma A.8. These hit events, for a countable open basis, generate the Borel sigma field of the compact hyperspace. Indeed a Hausdorff neighborhood is specified by finitely many open hits and containment in an open neighborhood, and containment is the complement of hitting a closed set; the latter is the intersection of the hits on its shrinking open neighborhoods. The empty-class event is treated in the same way. Thus the hyperspace-valued

map is measurable after completion. Under each individual law it has a Borel version: approximate it measurably by values in successive finite nets, replace the countably many completed measurable fibers by Borel representatives, and take the almost sure limit. These versions suffice for enhanced extraction and do not alter any probability. No claim of a common Borel version for all laws is needed. □

The limiting set $\mathcal{K}$ need not equal the entire arm class of the limiting metric. Equality is unnecessary. It contains every possible limit of the approximating contacts which must be covered. Conversely, selecting representatives from $\mathcal{K}$ supplies actual approximating arms. This observation avoids a closed-half-space recovery assertion.

Lemma 13.22 (Persistence of finite metric data)**.** *The following data persist under enhanced convergence, with every fixed strict relaxation of their numerical bounds.*

1. *A lower bound $D^{[K]}(A, B) > a$, for fixed compact $A, B, K \Subset U$, provided every tested path lies in K.*
2. *A finite graph of E-rectifiable paths, each compactly contained in its prescribed open buffer, together with upper bounds on their total lengths and continuous weighted lengths. Common vertices can be recovered as common vertices, and rational vertices can be kept fixed. Required positive clearances from fixed compact obstacles persist. The recovered graph is not required to be embedded.*
3. *Finitely many retained universal annular or clipped-annular connector successes, including their finite rational candidate menus, anchored core paths, and strict local-unit bounds.*

Proof. If the first assertion fails, choose paths of length at most $a + o(1)$ in K with endpoints in A and B. The common inverse modulus gives a uniformly convergent subsequence. Its limit lies in K and has D-length at most a, contradicting the strict lower bound. A separate lower exit-strip test can be included if a global lower bound is to be deduced from a compactly constrained one.

Here are details concerning common vertices in the second assertion. For each vertex v, choose rational points $q_{v,j}$ converging intrinsically to v inside a small prescribed neighborhood of v. Use $q_{v,j} = v$ for a rational vertex which must remain fixed. The connecting access paths have length tending to zero and hence images shrinking to v, by compact lower barriers. For a fixed j, attach these paths to each incident edge. The resulting finite family has rational endpoints and admits simultaneous buffered recovery. Let j tend sufficiently slowly to infinity along the approximating sequence. The recovered edges share the endpoints $q_{v,j}$; their images converge uniformly to the original edges and their lengths have the required upper limits. Weighted recovery gives the same assertion for continuous positive weights. A positive distance between a compact path image and a fixed forbidden compact set persists under uniform convergence. This proves every assertion in item 2.

For item 3, a retained binary success coordinate is eventually equal to its limiting value. The corresponding universal conclusion applies to the actual approximating arms; it is not inferred by recovering one chosen limiting arm. The finite core paths recover by item 2, and their rational endpoints stay fixed. Local units are retained scalar coordinates and converge. A fixed strict upper or lower bound on a unit therefore persists. □

Lemma 13.23 (Persistent access to finitely many exterior ports)**.** *Let $\mathcal{K}_n \to \mathcal{K}$ be the compact incoming-arm enhancement of Lemma 13.21. Fix $\eta > 0$. On a limiting environment with the clipped-annular successes and anchored boundary ports proved above, there are finitely many rational*

exterior ports, rational interior ports, and strict one-boundary-contact access arcs with pairwise disjoint images, such that the following holds. Every contact of every arm in $\mathcal{K}$ has closed-exterior E-access of cost less than η to one exterior port. The exterior port has an access arc to its interior port of length less than η. The latter arcs lie in prescribed disjoint small contact neighborhoods. For all sufficiently late n, the same exterior-port menu covers all contacts of $\mathcal{K}_n$ with cost less than 2η; the access arcs have buffered recoveries with all fixed wall clearances.

Proof. Write $x = P(0)$ for a member (P, ℓ) of $\mathcal{K}$. Choose a small parameter $t_0 > 0$, so that the initial reference segment $P|_{[0,t_0]}$ has E-length at most $C\ell_+ t_0 < \eta/20$. The inverse modulus makes this segment nonconstant. Choose N successful clipped annuli surrounding x, crossed by this reference segment, with pairwise disjoint enlarged buffers. Make their outermost radius so small that the terminal occupation bound for every admissible incoming arm is less than $\eta/20$ in E units. Choose N large enough that $C_0 C\ell_+/N < \eta/20$, where C_0 is a common symmetric connector constant. These choices are possible by annular density and by decreasing the radii after fixing N. Choose the annular inner and outer boundaries with a fixed positive clearance from the two ends of the reference segment. These are strict crossing tests, and hence survive uniform path convergence.

The interiors of all their inner holes contain one Euclidean neighborhood H_x of x. Any incoming arm whose contact lies in H_x crosses all the chosen annuli. Its crossing-arm masses have total at most $C\ell_+$, while the reference crossing-arm masses have total at most $C\ell_+ t_0$. At one annulus their sum is at most $C\ell_+(1 + t_0)/N$. The connector there, the reference initial segment and the incoming terminal segment show that its contact has closed-exterior E-distance less than $\eta/4$ from x. Only the bounds for these retained original subpaths are being used; no recovery of the closed-exterior reference path has been asserted.

The neighborhoods H_x cover the compact contact set of $\mathcal{K}$. Choose a finite subcover with distinct representatives $x_1, \dots, x_m$ and representative arms $P_1, \dots, P_m$. Each P_i has approximating arms $P_{i,n} \in \mathcal{K}_n$. Uniform convergence makes the approximating arms cross the same finite annuli. Their initial-segment D_n-lengths are at most $\ell_+ t_0$ directly from their Lipschitz clocks, so the same E_n length bounds hold by (13.14). The retained successes therefore give the same $\eta/4$ comparison, with a fixed relaxation, between every approximating contact in H_{x_i} and $P_{i,n}(0)$. Every contact of $\mathcal{K}_n$ is in this finite union for late n: otherwise an uncovered sequence has, by Hausdorff compactness, an uncovered limiting arm in $\mathcal{K}$.

After fixing the finite representatives, choose disjoint tiny Euclidean neighborhoods of them, away from every interior trial box. Inside the ith neighborhood use an anchored clipped core at a sufficiently small successful radius r_i. Its lower crossing is at least $a s_{r_i}$ and its total port and core cost is at most $H s_{r_i}$. The reference germ crossing that annulus can be made arbitrarily short. Hence its crossing charge gives s_{r_i} arbitrarily small, and the strict upper bound on this unit can be retained. Attaching the reference arm to the core and following the anchor paths gives a fixed rational exterior port q_i with closed-exterior access from x_i of cost less than $\eta/4$. The boundary-port part of the same construction gives a simple arc F_i from q_i through one rational boundary point to an interior rational port, of length less than η. The arcs have disjoint images because their entire neighborhoods are disjoint.

In the approximating metrics, the actual reference arm $P_{i,n}$ crosses this anchored core. Its retained terminal length is small by its Lipschitz clock and occupation bound. The retained universal success, the convergent strict unit bound and the recovered finite core paths give access from $P_{i,n}(0)$ to the same q_i at cost less than $\eta/2$. Combining with the wide-annulus comparison gives the asserted 2η bound for all approximating contacts. Finally, each whole F_i is a finite rational-endpoint path in the open data buffer. Recover it there by Lemma 13.22. Its one-boundary-contact property was

used to construct the limiting pendant tree. Its recovery is used only for upper lengths and wall clearances, both of which persist in that open buffer. □

Definition 13.24 (Finite wall certificate)**.** A finite wall certificate records the following data from the construction in Proposition 13.19.

1. A finite abstract tree, its leaf stubs and marked middles; a finite family of realizing E-paths; the exterior access paths; and open buffers for all these paths. It records upper bounds on the sum of all edge lengths, on each marked-middle length, and on the sum of the leaf-stub lengths. The abstract incidence relations are part of the data.
2. The collector, the working buffer, a compact body L, an inner body L_{in}, pairwise disjoint leaf-gate buffers, and the two cut-gate buffers for each protected middle. It records the continuous outer-wall coordinate and the continuous cut coordinates. The outer coordinate is a continuous distance level around the central subtree: its lower and upper level sets separate L_{in} from the exterior of L. Each cut wall is the slab between the $-1/2$ and $1/2$ levels of its cut coordinate on L; the coordinate takes values -1 and 1 on its recorded leaf and central sides. For successive cuts these satisfy

$$\Phi_u < 1 \implies \Phi_v = -1, \qquad \Phi_v > -1 \implies \Phi_u = 1.$$

   Each leaf gate has the leaf-side values for its own marked middle and the central-side values for every other leaf's marked middle. A recorded compact carrier containing $\{f > 0\}$ is contained in the union of L_{in} and the inner leaf-gate holes. The part of every wall on which the weight can be below its wall lower bound is contained in the recorded inner gate holes. All these inclusions have prescribed positive Euclidean buffers.
3. Lower metric bounds between the two levels of every wall and between the two cut-gate neighborhoods of each marked middle. Each bound is constrained to a fixed compact buffer in U containing every path segment to which it will be applied. If an exit is possible, its crossing is listed as an additional lower test.
4. For each gate, finitely many successful annuli with disjoint buffers and short reference germs, certifying actual-contact access below δ to the corresponding leaf or cut arc for every local arm of the prescribed length bound. References are part of the finite network, with rational reference anchors and recoverable connecting paths. Collector-contact access is certified separately by the finite cover of Lemma 13.23.
5. A fixed smooth profile f, with $w = e^{-\xi f}$, bounds

$$\begin{aligned} &0 < \lambda_- \le w \le B_+, \\ &w \le \lambda_+ < 1 \quad \text{on open tubes around the upper routes,} \\ &w \ge B_- > 1 \quad \text{on the wall slabs outside their holes.} \end{aligned}$$

   The profile vanishes off the collector and the data annulus. The access paths avoid the wall slabs; their actual weighted upper costs are included in the insertion budget.

The certificate has a *strict reserve* if all numerical inequalities used below have a fixed positive margin. In particular, writing a_{wall} for the wall lower bound, L_T for the total tree-length upper

bound, ℓ_R for a marked-middle upper bound and a'_R for its cut-neighborhood lower bound, require

$$B_- a_{\mathrm{wall}} > 2A, \tag{13.18}$$

$$(2\eta_0 + \lambda_+ L_T)/c_* < a/2, \tag{13.19}$$

$$2B_+\delta + \lambda_+(\ell_R + 2\eta_{\mathrm{cut}}) < c'_1\lambda_- a'_R, \qquad c'_1 < c_1, \tag{13.20}$$

$$\sum_i (2\delta + \ell_{K_i}) < \epsilon_0, \qquad 0 < b < \tfrac{1}{2}\lambda_- \min_R a'_R. \tag{13.21}$$

Upper path and access costs in this definition are allowed to use slightly smaller budgets than the displayed ones. This makes all subsequent finite relaxations simultaneous.

The cut coordinates and the compact bodies are geometric data, not functions which must be chosen measurably from the metric. At one limiting realization they are supplied by the embedded tree and its retraction. During recovery they remain fixed. The recovered graph need only obey the recorded incidences, upper costs, tubes and clearances.

Lemma 13.25 (Finite sufficiency and persistence)**.** *A finite wall certificate with strict reserve implies the signed local contract of Proposition 13.19, with its recorded constants. Under enhanced convergence, including the compact incoming arm families, such a certificate supplies the same fixed profile and the same relaxed contract in all sufficiently late approximations. The active interval remains confined to its recorded annular buffer. The restoration conclusions continue to hold after finitely many canonical insertions in disjoint working regions.*

Proof. We spell out which data are used in each conclusion. Insertion follows from the two exterior-port accesses, their access arcs, and an upper route through the finite graph. Inequality (13.19), together with $D \le E/c_*$, gives the required cost. Retained original pieces lie outside the collector, where $f = 0$.

Consider a path of weighted first-to-last working-region span at most A. If it crosses a wall slab without meeting one of its recorded inner gate holes, its portion between the last lower-level hit and the first upper-level hit stays in the slab. On that portion $w \ge B_-$, and its original D-length is at least a_{wall} by the corresponding compact lower test. This contradicts (13.18). Thus every relevant wall crossing visits a gate hole.

On each local gate arm, the original E-length is at most CA/λ_-. The finite annular access certificate applies to that bound. At a cut hole it gives an actual path contact, in the recorded cut neighborhood, with original E-access below δ to the cut arc. If a body excursion goes from one leaf gate to a different leaf gate, continuity of the cut coordinates gives two successive cut traversals for one protected middle. The displayed compatibility implications order these traversals. Choose the access contacts on the corresponding traversing arms; the gate buffers are disjoint, so the contacts are chronologically ordered. The interval between them stays in the compact body together with its gate buffers, and hence in U.

The compact cut lower test makes its original D-length at least a'_R, and therefore its weighted length at least $\lambda_- a'_R$. The two access paths, the short cut arcs and the upper marked route have weighted E-cost at most $2B_+\delta + \lambda_+(\ell_R + 2\eta_{\mathrm{cut}})$. By (13.20) and (13.21), this is an active shortcut of ratio below c_1 and length threshold above b. If the same two cut traversals occur during an excursion returning to its original leaf gate, the identical conclusion applies.

If no active shortcut exists, each excursion through the inner body uses only one leaf gate. Positive-carrier visits outside that body are already in a leaf-gate hole. Because the complete path has endpoints outside the working region, every such visit has a preceding and a following gate arm.

Their finite access certificates give actual contacts with access below δ to that leaf stub. Process the earliest remaining visit to the compact recorded carrier and replace from its preceding gate contact to the last necessary original contact at the same gate. This costs at most $2\delta + \ell_{K_i}$. Its image lies in that gate's buffer, disjoint from every other gate buffer. The gate is never used again. After at most the number of leaf gates, every carrier visit is removed, and (13.21) bounds the total original E-cost by ϵ_0. Without inactivity, the first and last necessary gate contacts can instead be joined through the whole graph; this gives the general restoration bound $2\delta + L_T$, enlarged by the finite recorded access budgets if needed.

Now suppose that canonical paths have already been inserted in other disjoint working regions. Their images miss the target gate buffers. Every surviving target carrier visit is on a retained original piece. Reading that piece backward and forward to the boundary of its gate buffer gives fresh original gate arms: an inserted image cannot stop one of these readings inside the buffer. Apply the same finite access certificate to these arms. A new deletion either retains an old inserted path whole or swallows it whole, since the new endpoints lie on original pieces and the other images avoid the target buffers. Thus the restoration algorithm and its cost ledger remain valid. No competitor-length bound is applied to an inserted canonical path.

Finally apply Lemma 13.22 to the finitely many lower tests and upper routes, and Lemma 13.23 to insertion contacts. All wall coordinates, bodies, profiles and gate buffers stay fixed. Their strict geometric clearances persist under upper-network recovery; the finite annular successes apply to all actual approximating arms with the prescribed relaxed length bounds. Strict reserves in (13.18)–(13.21) absorb every length and access error. The proof just given therefore applies in each sufficiently late approximating metric, with the same profile. It has used no stability of a geodesic selector and no preservation of the embedding of the recovered graph. □

Lemma 13.26 (A countable profile family and measurable working events)**.** *There is a countable family of smooth signed profiles, supported in each prescribed rational collector and data buffer, sufficient for all strict wall constructions above. For a fixed profile and rational numerical contract parameters, the event that its insertion, active-or-inactive, and restoration contracts hold is measurable in the completed local field sigma field. The assertion includes the finite hybrid-restoration tests used in the counting argument.*

Proof. Write $f = f_+ - f_-$ with nonnegative parts. First truncate a small positive level of each part. Approximate the truncated parts uniformly by finite nonnegative smooth bump sums with rational coefficients and rational parameters, with supports compactly contained in the corresponding open sign regions. These sums form a countable family. In particular the approximating profile has no positive support outside the old positive-discount carrier. If $\|f' - f\|_\infty < t$, every weighted path length changes by a factor in $[e^{-\xi t}, e^{\xi t}]$, every shortcut ratio changes by at most $e^{2\xi t}$, and the bound for the original lengths of admissible weighted arms increases by at most $e^{\xi t}$. Construct the access charts with this fixed larger bound. Choosing t below the finitely many strict reserves preserves a relaxed contract. The original inactive bracketing covers the smaller positive carrier of f'.

We give the measurability argument independently of certificate selection. Code curves in the complete separable uniform space $C([0,1], \mathbb{R}^d)$, together with a total-length parameter. Take a countable union over integer bounds on the Euclidean range and the total length. Compact modulus classes are needed for extraction, but are not needed for this analytic-projection argument. Lemma A.9 makes all length and local rectifiability conditions Borel. The actual intrinsic distances are Borel functions of the local metric data and their endpoints, by the rational reconstruction. All path inequalities below therefore define Borel relations before any universal quantifier is imposed.

Activity can be tested at rational pairs of path times. Indeed the path's weighted arclength has no atoms. On a segment compactly contained in U, the endpoint intrinsic E-distance is continuous along the path: short retained subarcs have arbitrarily small D length, hence arbitrarily small E length by (13.14). Strict threshold and ratio inequalities therefore survive a sufficiently small perturbation of the two times to rational times. Containment of the segment in U survives because its image is compact. The active predicate is consequently a countable union of Borel tests.

Restoration can likewise be tested on finite ordered lists of rational time intervals. A strict-cost restoration has endpoints in fixed open gate buffers. First merge touching deletion intervals, concatenating their restoration routes; this changes no total cost. The remaining finitely many intervals have positive mutual gaps. Enlarge each one by less than one third of the smallest such gap to rational endpoints; local endpoint continuity preserves the strict cost bound. The enlarged intervals still have their endpoints in the required buffers and cover all old carrier visits. For a fixed finite list, coverage is the Borel condition

$$\max_{t\in[0,1]\setminus\bigcup_j(s_j,t_j)} f(P(t)) \le 0.$$

An infimum internal connection cost strictly below the budget is equivalent to the existence of a connection strictly below that budget, so no minimizing path need be selected. Summing the finitely many endpoint costs gives a Borel test. A countable union over rational lists gives the restoration predicate.

For hybrid restoration, add finitely many original time intervals and continuous inserted curves as code coordinates. Require their endpoint incidences and their confinement to the other disjoint working regions. These are Borel conditions. Target deletion endpoints have positive clearance from those other regions, so the same rational-time enlargement can be performed within retained original pieces. The coverage and cost tests are again Borel, with a countable union over the number of insertions and deletions.

The failure of the fixed-profile contract is now the existence of a curve code, or a finite hybrid code, satisfying its Borel admissibility conditions and violating its Borel conclusion. It is an analytic projection. The working event is coanalytic and is therefore universally measurable by Lemma A.8. Insertion failure has the same description using two incoming-arm codes and their endpoint internal cost. All the data lie in the specified annular buffer; hence these events belong to its completed field sigma field. Equivalent Cameron–Martin shifts have the same null sets, so completion causes no ambiguity in the counting estimates. This proof never projects a coanalytic certificate-validity relation. □

Lemma 13.27 (Sequential finite-prefix principle)**.** *Let $(A_j)_{j\ge1}$ be measurable working events for a countable profile family. Suppose every sequence of laws under consideration has an enhanced almost surely convergent subsequence with the following property. There is a measurable limiting construction event G of probability strictly greater than p such that, at every realization in G, some finite strict certificate for some A_j implies A_j for all sufficiently late approximations. Neither the chosen index nor the chosen certificate is required to be selected measurably. Then some $J<\infty$ satisfies*

$$\inf_{\mu} \mu\Big(\bigcup_{j\le J} A_j\Big) > p.$$

Proof. If the conclusion fails, choose a law μ_n whose first n events have union probability at most

$p + 1/n$. Extract the specified coupling, writing X_n for the approximating objects, and put

$$B_j = \bigcup_{m\ge 1} \bigcap_{n\ge m} \{X_n \in A_j\}.$$

Each B_j is measurable. Pointwise certificate persistence gives $G \subset \bigcup_{j\ge 1} B_j$. Continuity from below therefore supplies $J < \infty$ with $p' := \mathbb{P}(\bigcup_{j\le J} B_j) > p$. On this fixed measurable event, X_n belongs to $\bigcup_{j\le J} A_j$ for every sufficiently large n. Fatou's lemma gives

$$\liminf_n \mu_n\Big(\bigcup_{j\le J} A_j\Big) \ge p' > p,$$

contradicting the choice of μ_n for $n \ge J$. The argument uses only measurable working events in the approximations and the known measurable construction event G; it makes no assertion about measurability of the existential set of all finite certificates. □

Lemma 13.28. *Consider an enhanced relatively compact family of canonical pairs with fixed comparison constants and uniform construction estimates. Suppose a common finite rational menu of the local strict low-ratio witnesses in Lemma 13.30 has probability at least $q > 0$ under every law in the family. Require common upper path-length caps, fixed compact path and testing constraints, fixed open exit targets, and fixed positive reserves. In particular the local endpoint and exit tests imply a strict global rational endpoint lower bound. Then, for each prescribed failure probability, there is a deterministic finite family of signed profiles satisfying Proposition 13.19 with at least the prescribed success probability under every law of this family. The insertion gain, inactive budget, shortcut ratio and threshold, and general restoration bound are uniform.*

In the weak-metric application, the reverse capped-witness statement of Lemma 13.29, the rational-menu construction, and Lemma 13.30 supply exactly these quantitative hypotheses on an admissible positive-density set of radii. Thus strict discrepancy of the lower optimal ratio gives the stated finite family on that set. For an exactly covariant comparable strong pair, a positive-probability strict discrepancy at one scale gives the same conclusion at every scale.

Proof. First make the discrepancy a finite strict local test. Choose a finite cheap E-path and a strictly larger D endpoint threshold. Approximate its endpoints intrinsically by rational points; lower semicontinuity of D preserves a relaxed threshold. Confine the upper path to a compact rational buffer. The original global endpoint lower bound implies the same lower bound for every compactly constrained distance used by the construction. Record these compact constraints in the local data region. Only these constrained lower bounds are used for active intervals, which stay in the recorded body and gate buffers. If a global lower bound is needed for an auxiliary step, record its separate exit-strip lower test as in Lemma 13.22. Trim with a strict reserve to obtain a protected middle and two nonempty stubs. Countably many rational endpoint, buffer, threshold and margin choices suffice, so a finite choice has positive probability.

The required quantitative witness data are uniform by the explicit hypotheses. For their verification in the subsequent rigidity proof, Lemma 13.29, the polynomial-menu paragraph and Lemma 13.30 provide the finite rational menu, the common positive probability, bounded upper paths and fixed strict reserves. These arguments do not use the present lemma, so this forward reference is not circular. On an enhanced subsequence retain the success indicator, its finite menu index, upper paths and unit, as well as the *global* rational endpoint distances. For each fixed open exit target also retain all global distances from the menu endpoints to its rational points. Rational evaluation consistency passes these lower bounds to the limit. Intrinsic rational density inside

the open targets passes their infima, as proved in Lemma 13.30. Compact extraction passes the upper paths. The larger recorded reserves therefore give a relaxed local witness with the same positive-probability lower bound in the limiting law. This is a different direction from Lemma 13.22, which transfers a strict *limiting* lower test to the approximations and is used later for the constructed wall certificate.

Amplify those witnesses in finitely many disjoint interior trial boxes. Split the field at scale L times the trial radius. After subtraction of one coarse value, the coarse oscillation on a trial box tends to zero in probability as $L \to \infty$. Simultaneous Weyl scaling shows that a strict full-field witness yields a relaxed fine-field witness except on the small coarse-oscillation event. Fine fields in disjoint enlarged footprints are independent, jointly for the two metric coordinates by locality and the conditional-product construction. First choose the number of trials for the desired failure probability, then choose L, then choose a fixed microscopic ratio small enough that the footprints fit. Restoring the coarse fields uses only the retained strict reserves. Every prescribed interior box therefore contains a protected mark with the desired high probability. For weak pairs perform this construction on the admissible radii where the quantitative witness has its fixed positive probability. The fixed microscopic ratio shifts the geometric indices by a fixed amount and preserves their positive density.

Fix a high-probability common inverse-modulus, occupation and annular event. Its quantitative constants give compact incoming-arm classes and a deterministic bound on the number of contact representatives at any fixed intrinsic access tolerance. Retain those compact classes as in Lemma 13.21. Lemma 13.23 gives disjoint pendant access arcs for a finite representative menu. The protected-star construction applies for every $d \geq 3$. All its upper costs are finite; no bound for those costs is required before the star is constructed.

Run Proposition 13.19 with strict reserves, recording the finite wall certificate of Definition 13.24. The order of choices is essential: fix the desired span bound and restoration budget, then the wall levels and their positive lower bounds, then the wall multiplier, then the discount depth, and finally the finite gate-access charts and their holes. These choices give (13.18)–(13.21). The network has open discount tubes and the exterior access arcs have positive wall clearance. Approximate the resulting profile by the countable family of Lemma 13.26, using only part of the strict reserves.

To obtain one finite family uniformly, suppose that no finite prefix has the required success probability. Choose a bad law for each prefix and extract an enhanced convergent subsequence, also retaining the compact incoming-arm classes and the quantitative strict witness. The limiting arm class may be a proper subset of the limiting metric's full class, but the preceding construction applies to any such compact class and its approximating representative arms. The limiting strict-construction success probability is larger than the prescribed target. Lemma 13.25 transfers each of its finite certificates with the same fixed smooth profile. Lemma 13.27 contradicts the choice of the bad laws.

Take the minimum positive threshold and insertion gain, and the maximum finite general-restoration bound, over this finite prefix, decreasing or increasing them within their strict reserves as appropriate. The inactive budget and active ratio were fixed before selecting the prefix. All working events are completed-local measurable by Lemma 13.26. In the exactly covariant case the normalized pair has one law at every radius; continuity from below already gives a finite prefix, and exact covariance transports it to every scale. □

Here a *buffered macroscopic witness* at ratio c consists of a finite $A^{[K]}$-geodesic P in a fixed

compact rational annulus K, with its image a positive Euclidean distance from ∂K, and

$$B^{[K]}(P(0), P(L)) > cL.$$

Its length is between two fixed positive multiples of the local unit. The central parameter half of P is separated by a fixed positive Euclidean distance from its two endpoints. All these positive constants, as well as the probability of the witness, are fixed before the microscopic scale is sent to zero. Countably many rational annuli, length bounds and positive rational separation constants suffice to express this event. The geometric separation exists for any one compact geodesic because it is injective and Euclidean continuous. It permits all sufficiently small connector buffers used on the central half to remain inside K.

Lemma 13.29. *Let $B \le C_0 A$ be an optimally comparable canonical pair. There are $a_H > 0$, $A_H < \infty$ and $p_* > 0$, independent of $c \in [C_0/2, C_0)$, with the following implication. A positive-probability buffered macroscopic witness at a threshold sufficiently near C_0 forces at least three quarters of the levels of each sufficiently small logarithmic block to have probability at least p_* of a capped annular witness at threshold c.*

A capped witness at z, r uses the fixed compact constraint

$$K_{z,r} = \{r/4 \le |x - z|_\infty \le 4r\}.$$

It is an $A^{[K_{z,r}]}$-geodesic parent crossing from radius $r/2$ to $3r$, of length at most $A_H s_r$, with a middle crossing P_{uv} from r to $2r$, contained in that closed middle shell, such that

$$B^{[K_{z,r}]}(u, v) > c \operatorname{len}_A(P_{uv}).$$

All these tests are local to one open annular buffer of $K_{z,r}$. A fixed positive probability of a capped witness also gives a fixed positive probability of a buffered witness, uniformly in scale and in the ratio threshold.

Proof. Choose the lower crossing quantile a_H and the high-density span envelope first, and then A_H larger than twice its cap. At a site visited by a macroscopic geodesic, take its last radius-$r/2$ contact before the first subsequent radius-$3r$ contact following an inner visit. This parent stays in $K_{z,r}$ and is a constrained geodesic there, since $K_{z,r}$ lies inside the macroscopic constraint. If the full outer span is at most $A_H s_r$ and the capped witness fails, its middle crossing has A-length at least $a_H s_r$ and admits a B-connection in $K_{z,r}$ of cost at most c times that length, with arbitrarily small infimum error. Retaining the other pieces with comparison $B \le C_0 A$ saves at least

$$\gamma(c)|I|, \qquad \gamma(c) = \frac{(C_0 - c)a_H}{A_H} > 0$$

on its full outer visit interval I.

If a quarter of the candidate levels had witness probability below a sufficiently small fixed p_*, the event consisting of no witness and a lower crossing bound would have high probability on those levels. Annular concentration and the span envelopes cover the central half of every buffered macroscopic witness by successful outer visit intervals, outside a probability tending to zero. Select a finite inclusion-minimal interval cover. No point lies in more than two of its intervals: a middle interval among three with a common interior point is covered by the other two and is unnecessary. Thus one of two disjoint-interior colors has total length at least a quarter of the parent length. Replacing that color gives a B-connection of cost at most $(C_0 - \gamma(c)/4)$ times the parent length, contradicting a buffered witness at threshold $C_0 - \gamma(c)/8$. All probability and geometric choices preceded $C_0 - c$.

Optimality supplies a buffered witness at some macroscopic scale: take a strict finite-pair discrepancy, a geodesic witnessing its length, and a compact annular constraint containing that path with clearance. Countably many rational choices preserve positive probability. For the last assertion, the capped middle length lies in $[a_H s_r, A_H s_r]$. Occupation bounds around its endpoints, applied to the full parent whose endpoints are separated from them, keep the central half away from small endpoint balls. A fixed Gaussian quantile compares the local and buffered units. This gives fixed positive buffer and probability parameters. □

For later use, rational endpoint menus have polynomial size. At accuracy δs, choose successful annuli with radii a fixed power of δ and attach each eligible contact to a rational core vertex. Occupation bounds make the retained terminal piece and connector cost at most δs. There are polynomially many deterministic annuli. A core-depth cutoff $O(\log \delta^{-1})$ also has polynomially many candidate vertices, and its superexponential tail makes its total failure smaller than any prescribed power of δ. Hence each endpoint collar has a deterministic menu of size at most $C\delta^{-K_0}$. The constants and K_0 are fixed before the counting parameter below.

Local witnesses with persistent global rational tests

The passage from a local constrained witness to a protected mark needs one additional step. We give it explicitly, since lower bounds for constrained distances do not in general pass from approximations to a limit merely by lower semicontinuity.

Lemma 13.30 (Quantitative localization of reverse witnesses)**.** *Suppose $c_* D \le E \le CD$, with fixed positive comparison constants, and the pair has the uniform canonical estimates used above. Fix*

$$1/C < b_0 < b_1 < b_2 < 1/c_*.$$

Suppose that the local capped witness for the reverse ordered pair (E, D), with ratio b_2, has probability at least $q > 0$ at the radii under consideration. There are a fixed finite rational endpoint menu, a finite positive list of thresholds a, a number $\tau > 0$, fixed compact path and testing buffers, and fixed open exit strips, with the following properties.

With probability at least a fixed $q' > 0$, one menu pair q_-, q_+ and one threshold a have a cheap finite E-path and a strict global *rational lower test,*

$$D(q_-, q_+) > (1 + 2\tau)as, \qquad \mathrm{len}_E(P) < (1 - 2\tau)as/b_0. \tag{13.22}$$

Their endpoint distances to the open exit strips exceed a fixed multiple of s larger than the endpoint threshold. The event can be tested using only actual internal metrics in one fixed determining annulus. The upper path-length caps and all reserves are uniform. The construction admits a protected middle, retaining nonempty end stubs, with ratio below $1/b_0 < C$.

Under enhanced convergence, the same positive-probability witness survives with every fixed smaller reserve, provided that global rational distance values, its upper path, and its unit are retained. This last assertion uses the consistency of rational evaluations, not a converse to lower semicontinuity.

Proof. Work first at center zero and radius one; all distances below are measured in the common local unit s. Put

$$K = \{1/4 \le |x|_\infty \le 4\}, \qquad M = \{3/4 \le |x|_\infty \le 9/4\},$$

and choose the open exit target

$$S = \{1/3 < |x|_\infty < 1/2\} \ \cup \ \{3 < |x|_\infty < 7/2\}.$$

Every Euclidean-continuous path from M to K^c reaches S before leaving K. All these sets lie in a fixed open annulus $U = \{a_d < |x|_\infty < 7\}$, where $0 < a_d < \min\{1/16, 1/(4\sqrt{d})\}$. This determining annulus also contains the Euclidean unit sphere defining the local unit.

The normalized lower crossing bounds give positive constants a_{mid}, β such that, outside an event of probability less than $q/8$,

$$\ell \geq a_{\mathrm{mid}} s, \qquad D(M, S) \geq 4\beta s. \tag{13.23}$$

Here ℓ is the E-length of any middle crossing in the capped witness, and its upper bound is $\ell \leq A_H s$. For the second estimate, cover the compact middle region by a fixed finite grid of small cubes separated from S. A path from the middle region to S crosses one of their fixed enlarged shells. Uniform lower crossing quantiles, and a finite union bound, give the stated estimate. Equivalently one can use the compact lower barriers from the construction. The infimum $D(M, S)$ equals the infimum constrained to K: a path that leaves K before reaching its proposed endpoint in S has an earlier contact in S, and can be truncated there. For an open strip, take contacts just after the first entrance. Thus this barrier event is local as well.

Choose an integer N so large that

$$A_H/(c_* N) < \beta/8.$$

Divide the capped middle E-path into N equal-arclength pieces, each of length $t = \ell/N$. Their endpoints lie in $\{1 \leq |x|_\infty \leq 2\} \subset M^\circ$, and

$$D(u, v) \leq t/c_* < \beta s/8$$

for every such pair. A global D near-minimizer of this cost cannot reach S, by (13.23); it stays in K. Consequently its global and K-constrained endpoint distances coincide. If every piece satisfied $D(u, v) \leq b_2 t$, concatenate these constrained connections, with arbitrarily small total infimum error. This would give $D^{[K]}$ between the middle endpoints at most $b_2 \ell$, contradicting the reverse capped witness. At least one piece therefore satisfies

$$D(u, v) > b_2 t, \qquad a_{\mathrm{mid}} s/N \leq t \leq A_H s/N. \tag{13.24}$$

This is the step that turns a local constrained discrepancy into a global discrepancy.

Use the rational-menu construction preceding this lemma at one fixed sufficiently small accuracy η. The two ends of the selected piece are points of the full parent, a fixed Euclidean distance from its endpoints. The local proof of the occupation estimate applies to this constrained parent because every annulus used near these points lies inside K. It and the rational core attachments therefore give, outside a further failure of probability less than $q/8$, rational points $q_-, q_+ \in M^\circ$ in one deterministic finite menu, with E-access at most ηs to u, v. The access paths lie in M, using its fixed margin around the middle region. Their D-costs are at most $\eta s/c_*$. Choose η so that

$$2\eta(b_1 + c_*^{-1}) < (b_2 - b_1) a_{\mathrm{mid}}/(4N), \qquad 2\eta/c_* < \beta/8.$$

If P is the concatenation of the two accesses with the selected piece, and $L_P = \mathrm{len}_E(P)$, then

$$L_P \leq t + 2\eta s, \qquad D(q_-, q_+) > b_2 t - 2\eta s/c_* > b_1 L_P, \qquad D(q_-, q_+) < \beta s/4.$$

The last inequality follows by using P as a competitor and the chosen bounds for t/c_* and $2\eta/c_*$. Also $L_P \geq t \geq a_{\mathrm{mid}}s/N$, since this particular concatenated path contains the entire selected piece.

Choose $\tau > 0$ sufficiently small that

$$\frac{b_1}{1+2\tau} > \frac{b_0}{1-2\tau}.$$

The interval of allowed thresholds

$$\frac{b_0 L_P}{(1-2\tau)s} < a < \frac{D(q_-, q_+)}{(1+2\tau)s}$$

has a deterministic positive minimum length, by the lower bound for L_P/s and the strict reserve just chosen. Its endpoints lie in a fixed bounded positive interval, and its upper endpoint is below $\beta/4$. A fixed sufficiently fine finite positive grid therefore contains an allowed a. This proves (13.22), with $a < \beta/4$. Both rational endpoints lie in M, so their global distances to S are at least $4\beta s$.

Here is the precise local test. Require an upper E-path in the compact constraint M with the indicated strict bound, the endpoint test

$$D^{[K]}(q_-, q_+) > (1+\tau)as,$$

and the two exit tests

$$\inf_{y \in S} D^{[K]}(q_\pm, y) > 2\beta s.$$

These are functions of actual internal data in U. The compact upper and lower tests are Borel; for the open target S, exhaust it by compact sets and take a countable infimum, as in Lemma A.2. Any global path from $q_\pm$ reaching S has a first such contact before leaving K, so the exit infima are also the global exit infima. Since $(1+\tau)a < 2\beta$, a putative global endpoint path below the tested threshold could not exit K. The local endpoint test thus implies the global lower bound. The construction above satisfies these local tests with the larger reserves 2τ and 4β. Its probability is at least $q' = q/2$, after allocating the remaining quantile and menu failures within $q/4$. All choices were uniform in the radius and in the family of laws.

Loop erasure gives a simple E-arc from q_- to q_+ with no greater length. Delete short positive terminal subarcs of total E-length much smaller than $c_*\tau as$. The D endpoint lower bound loses at most that length divided by c_*, while the upper E length can only decrease. The remaining middle retains a fixed strict ratio below $1/b_0$, and the deleted parts supply nonempty stubs. All of this lies in $M \Subset U$.

We finally prove the limiting assertion. Use the *stronger* local success relation: the upper bound has reserve 2τ, the constrained endpoint distance is greater than $(1+2\tau)as$, and both exit infima are at least $4\beta s$. The construction above gives this relation probability at least q'. It is Borel before projection of the upper path coordinate. Choose a successful menu index and upper path measurably using Lemma A.8, and take a Borel version under each law. Retain the success indicator, the finite index, the path, the unit and all global rational distances in an enhanced coupling. The indicator and finite index are eventually constant on their limiting success event. Its probability is at least q'. Bounded upper path lengths and compact lower barriers give a uniformly convergent path subsequence; its limit has no greater E length. Rational evaluation consistency gives

$$D(q_-, q_+) \geq (1+2\tau)as.$$

For every rational $y \in S$, it also gives

$$D(q_\pm, y) \geq 4\beta s.$$

The infimum of $D(q_{\pm}, y)$ over the open set S equals its infimum over rational $y \in S$. Indeed a finite-distance point of S has intrinsically arbitrarily close rational points inside S, and the triangle inequality gives equality of the two infima; points at infinite distance cannot lower either infimum. Thus the exit bounds hold for all of S in the limit. The smaller reserves in the displayed local tests follow. Notice that every lower bound was passed by convergence of fixed rational values and rational reconstruction; none was passed by reversing a lower-semicontinuity inequality. Restoring the center and radius gives the result. □

Proposition 13.31. *Two canonical weak assignments for the same field and parameters are deterministic positive multiples of each other. More generally, an exactly covariant comparable pair consisting of a canonical metric and a strong metric has this conclusion.*

Proof. Let $0 < c_* \le C$ be the optimal deterministic constants with $c_* D \le E \le CD$. Suppose $c_* < C$. Choose protected-mark ratio thresholds strictly between c_* and C, with reserves. Apply Lemma 13.29 to the reverse ordered pair to obtain positive-density admissible levels and Lemma 13.28 to obtain a finite signed family. Fix a baseline insertion gain $g > 0$, an inactive budget $\epsilon_0 < Cg/16$, active ratio $c_1 < C$, active threshold $b > 0$, and restoration bound $A_{\mathrm{rest}} < \infty$.

Fix a macroscopic constrained endpoint problem with

$$\ell_0 S \le L = D^{[K]}(q_1, q_2) \le MS, \qquad E^{[K]}(q_1, q_2) \ge CL - \delta S. \tag{13.25}$$

The endpoints are selected from the polynomial rational menus in two untouched, Euclideanly separated collars. A common lower crossing quantile between these collars gives the deterministic $\ell_0 > 0$, outside an arbitrarily small prescribed probability. Here $S = c_R e^{\xi h_{L_* R}(0)}$ uses a macro averaging sphere outside every possible microscopic support. For a label $i = (z, r)$, $s_i = c_r e^{\xi h_r(z)}$ uses a sphere inside its annular support's inner hole. Thus all units are unchanged by all counting shifts; a global repinning multiplies every distance and unit by one common factor. Every counting geodesic is selected measurably from the metric divided by S and its endpoint data, with curves parameterized on $[0, 1]$. Use Lemma A.8 on the compact minimizing-curve class. The selector therefore has exactly the same input after common scalar repinning; the selected geometric path is unchanged. The same convention is used for the compact endpoint problems below.

Choose the raw-unit exponent B_0, occupation exponent $\theta_0 > 0$, and mesoscopic separation exponent $0 < \nu < 1$ first. Choose the span recurrence coefficient so small that its terminal term, iterated from radius $r \in [\varepsilon^2 R, \varepsilon R]$ to $\varepsilon^\nu R$, is $o(\varepsilon^{B_0+1} S)$. Then choose its high-density cap, and only afterward the finite profiles. Intersect profile levels, span-envelope levels, and lower crossing events. The resulting inner balls cover a fixed central portion of the original D-geodesic. Its outer visit intervals have lengths at most As_i. Summing and binning by radius, disjoint spatial color, profile, and dyadic unit gives N visited labels with

$$\varepsilon^{B_0}/2 \le q \le \varepsilon^{\theta_0}, \qquad s_i \in [qS, 2qS], \qquad Nq \ge v_\varepsilon := c(\log \varepsilon^{-1})^{-2}. \tag{13.26}$$

The lower unit bound follows from (13.2) and Gaussian maxima. The upper bound follows from occupation and an untouched local crossing charge. The deterministic bin universe has at most $C\varepsilon^{-2d}$ labels.

Restrict to k-subsets whose centers have pairwise distance above a fixed multiple of $\varepsilon^\nu R$. For fixed k this loses a vanishing fraction: local crossing charges and occupation show that one label excludes at most $C(\varepsilon^{\nu\beta} + \delta)/q$ others, whereas $Nq \ge v_\varepsilon$. Take $\delta = \varepsilon^{B_0+1}$. Outward annuli at each selected label are untouched by all selected bumps up to the mesoscopic scale, so Lemma 13.14 bounds every discounted output's local span independently of discount depth.

Insertion and surviving outputs. For each allowed k-subset I, chronological collector replacement and Proposition 13.19 give

$$D^{[K]}_{h-f_I}(q_1, q_2) \le L - gkqS. \tag{13.27}$$

Each designated original crossing is deleted either by its own replacement or by an earlier one. Since working regions are disjoint, its interval is swallowed whole; the sum of deleted crossing charges pays for all inserted routes. Retained pieces have no negative-support cost increase, because the collector contains the entire signed support.

If a selected minimizer avoids the positive part of one remaining profile, remove that label and recompute. Its current path has $f_i \le 0$, so removing the profile cannot increase its cost. This stops at $Z \subset I$ visiting all remaining positive supports, still satisfying (13.27). Restore active labels at cost $A_{\mathrm{rest}} s_i$ and inactive labels at cost $\epsilon_0 s_i$. On retained pieces the perturbation is nonpositive, so original E-length is at most C times discounted D-length. If $m = |Z|$ and a_Z is its number of active labels, (13.25) implies

$$Cgk - \delta/q \le 2A_{\mathrm{rest}} a_Z + 2\epsilon_0(m - a_Z). \tag{13.28}$$

For small ε, fixed $\kappa_1, \kappa_2 > 0$ therefore satisfy $m \ge \kappa_1 k$ and $a_Z \ge \kappa_2 m$.

Let $A_m(h)$ count the necessary output events of size m, including the macro near-maximality, all local certificates and actual span bounds, and

$$D^{[K]}_{h-f_Z}(q_1, q_2) \le D^{[K]}_h(q_1, q_2) - gmqS. \tag{13.29}$$

An output of size m has at most $\binom{N-m}{k-m}$ input preimages. The asymptotically full set of allowed inputs and (13.26) give

$$\sum_{m=\lceil \kappa_1 k \rceil}^{k} q^m A_m(h) \ge c_k (v_\varepsilon/(2k))^k. \tag{13.30}$$

Indeed divide the preimage bound by N^k, use $\binom{N-m}{k-m} \le N^{k-m}/(k-m)!$, and then multiply A_m/N^m by $(Nq)^m \ge v_\varepsilon^k$.

All shifted outputs use one path. Evaluate the event for Z at $h + f_Z$. Its discounted metric is the common D_h, so every output uses one consistently selected original geodesic. By (13.29), its shifted original endpoint distance is at least L. A visited label pays an untouched confined crossing charge aqS; hence the common visited universe has $O(q^{-1})$ labels.

Fix one shifted target output Z_0, $|Z_0| \le k$. Every label outside Z_0 active in any shifted output gives an E_h shortcut along the common path. Its interval and image are in its own working region, so such intervals are disjoint. If there are n of them, their total saving against C times original length is at least $(C - c_1)b(n-k)qS$. Insert these shortcuts and restore the target output using its original certified contacts. A shortcut is retained whole or swallowed whole, and target restoration adds at most $2A_{\mathrm{rest}} kqS$. Thus

$$(C - c_1)b(n-k)qS \le 2A_{\mathrm{rest}} kqS + \delta S.$$

Since $\delta/q \le 2\varepsilon$, the entire shifted active universe has $O(k)$ labels. If $B_m(h)$ counts shifted size-m output events, at least $\lceil \kappa_2 m \rceil$ of their labels are in this bounded active universe; therefore

$$B_m(h) \le C_k q^{-(m - \lceil \kappa_2 m \rceil)}. \tag{13.31}$$

Gaussian counting. Each deterministic profile has uniformly bounded Cameron–Martin norm under translations and dilations. A sum of at most k such profiles also has uniformly bounded norm.

Gaussian change of measure and Hölder give, for any fixed $\alpha < 1$,

$$q^m \mathbb{E} A_m \le C_{\alpha,k} \varepsilon^{-2dm(1-\alpha)} q^{m-\alpha(m-\lceil \kappa_2 m \rceil)}.$$

This follows by summing the individual estimate $\mathbb{P}(A_Z(h)) \le C_{\alpha,k} \mathbb{P}(A_Z(h+f_Z))^\alpha$ over the deterministic label universe and applying finite-sum Hölder, followed by (13.31). Choose α near one so

$$\gamma = \alpha\kappa_2\theta_0 - 2d(1-\alpha) > 0.$$

Then $q^m \mathbb{E} A_m \le C_{\alpha,k} \varepsilon^{\gamma m}$. Pinning is harmless: replace f_Z by $f_Z - \bar{f}_Z$ in the Gaussian formula. The common scalar Weyl factor preserves every normalized event, unit ratio and selected geometric path. Monotonicity was used only through (13.29) in the unpinned representative.

There are at most $C\delta^{-2K_0}$ endpoint-menu pairs. Choose fixed k after all profile constants so that $\gamma\lceil \kappa_1 k \rceil > 2K_0(B_0+1)+1$, and then send $\varepsilon \downarrow 0$. Equations (13.30) and the Gaussian bound, together with the endpoint-menu union bound, suppress a near-maximal upper witness. Other bins and v_ε^{-k} contribute only logarithmic factors. The common environment failure is chosen once below the fixed witness probability.

For an exact pair, Lemma 13.29 and scale covariance give a fixed positive probability of near-maximal capped witnesses at every radius, a contradiction. For weak pairs use the two density sets. The reverse optimality argument gives reverse-buffered good radii at density at least three quarters in sufficiently long outer blocks; at each such radius its inner low-ratio admissible levels give the profile family and upper-witness suppression at one fixed threshold below C. Upper optimality gives capped upper witnesses at that fixed threshold on at least three quarters of sufficiently long outer blocks. The two base radii shift geometric indices by a fixed number, so their density-three-quarter sets intersect in long blocks, contradicting suppression. Thus $c_* = C$, and equality holds simultaneously on all finite and infinite pairs. The multiplier is deterministic by the definition of the optimal comparison constants. □

13.8 Exact covariance and identification of the normalization

Lemma 13.32. *For almost every deterministic $R > 0$, almost surely simultaneously for all $x, y \in B_R$,*

$$E^{B_R}(x,y) = E^{[\overline{B}_R]}(x,y). \tag{13.32}$$

If the metric law modulo a positive scalar is dilation invariant, this holds for every fixed deterministic R.

Proof. For rational p, q, $F_{p,q}(R) = E^{B_R}(p,q)$ is finite and nonincreasing once $p, q \in B_R$. Compactness of a path image makes it left-continuous. Compact extraction of bounded paths in decreasing larger balls gives

$$F_{p,q}(R+) = E^{[\overline{B}_R]}(p,q).$$

A finite monotone function has countably many jumps. Tonelli and countability give (13.32) on all rational pairs for almost every deterministic radius. If x, y have finite closed-ball distance, they are nonsingular. Intrinsic rational approximations within B_R, the equality for rational pairs and the triangle inequality extend it to x, y. Infinite distances and diagonal pairs need no approximation. The event (13.32) is invariant under multiplying the metric by a scalar. Dilation invariance makes its probability independent of R, so its almost-everywhere value one is its value at every fixed radius. □

Proposition 13.33. *For every fixed $r > 0$ and $z \in \mathbb{R}^d$, almost surely on all pairs,*

$$E_h(rx + z, ry + z) = E_{h(r\,\cdot+z)+Q\log r}(x, y).$$

Proof. The transformed assignment $E_g^{[r]}(x, y) = E_{g(\cdot/r)}(rx, ry)$ is in the canonical weak class by the LFPP coordinate change, uniform normalizer ratios and Proposition 13.11. Proposition 13.31 gives $E^{[r]} = k(r)E$ for a deterministic positive $k(r)$. Composition gives $k(rs) = k(r)k(s)$. With $T_r h = h(r\cdot) - h_r(0)$,

$$E_h(rx, ry) = k(r)e^{\xi h_r(0)} E_{T_r h}(x, y).$$

The normalized field $T_r h$ has the original law. Therefore the median of $e^{-\xi h_r(0)}$ times a fixed-shape radius-r crossing is exactly $k(r)$ times its radius-one median. Uniform crossing quantiles make $k(r)$ comparable to c_r, with constants independent of r. For fixed $0 < r < 1$, apply this comparison and (13.2) to $k(r)^n = k(r^n)$. Taking nth roots gives, for every $\eta > 0$,

$$r^{\xi Q+\eta} \le k(r) \le r^{\xi Q-\eta}.$$

Let $\eta \downarrow 0$. Thus $k(r) = r^{\xi Q}$ for $r < 1$, while multiplicativity gives the same result for $r > 1$ by inversion and for $r = 1$. Translation covariance and the simultaneous Weyl rule finish the proof. □

Lemma 13.34. *The internal distance $E_h(0, e_1; B_2)$ has an atomless law with a unique median.*

Proof. Choose a nonnegative smooth Cameron–Martin direction f supported in the Euclidean ball of radius $1/2$ about 0, and bounded below by $c > 0$ on a closed annular band separating 0 from e_1. Its average on the unit pinning sphere is zero. For

$$X(t) = E_{h+tf}(0, e_1; B_2)$$

the Weyl rule gives finiteness, continuity and, for $t_2 > t_1$,

$$X(t_2) \ge X(t_1) + (e^{\xi c(t_2-t_1)} - 1)E_{h+t_1 f}(\text{across the band}) > X(t_1).$$

The same lower crossing shows $X(t) \to \infty$ as $t \to \infty$. Thus each good affine Gaussian line maps continuously and strictly increasingly onto (L, ∞) for some $L \ge 0$. Disintegrate the pinned field along f and its independent orthogonal complement. On almost every line choose a coefficient giving a good field, and use simultaneous Weyl scaling to define the whole line. Conditional distributions have no atoms and give positive mass to every interval meeting (L, ∞). The unconditional distribution is atomless and strictly increasing above the essential infimum of the conditional lower endpoints. Its median is unique. □

Theorem 13.35. *With the internal point-median normalizer a_ε, the entire normalized LFPP family converges in probability, jointly with the field, to one field-measurable metric E_h. Its actual internal distance $E_h(0, e_1; B_2)$ has median one. For every fixed $\lambda > 0$,*

$$a_{\lambda\varepsilon}/a_\varepsilon \longrightarrow \lambda^{1-\xi Q}.$$

The assignment satisfies the five strong axioms, including their simultaneous formulation for arbitrary continuous field perturbations.

Proof. Weak uniqueness already gives dilation invariance modulo a scalar, so Lemma 13.32 applies at $R = 2$. Buffered recovery gives the upper limit for internal distances with rational endpoints. Compact extraction of near-minimizers in $\overline{B}_2$ gives the lower limit, identified with the open-ball distance by (13.32). Consequently the prelimit internal normalizing variable converges along every enhanced subsequence to the actual internal distance of its metric limit. Its prelimit median is one; Portmanteau makes one a limiting median. Lemma 13.34 then fixes the deterministic multiple in Proposition 13.31. All subsequences have one joint graph limit, and Proposition 13.10 gives convergence in probability.

For completeness the assignment is defined on continuous perturbations by the same Borel local reconstruction. Choose a deterministic LFPP subsequence on which the global metric and the countable global rational distance array converge almost surely for h. The buffered reconstruction then identifies the lower limits and exhaustions of the internal arrays; it does not require every raw internal array to have an almost sure graph limit. For an input distribution g, form the countable regularized internal distances in rational buffers along that subsequence, then their buffered lower limits, exhaustions, and rational lower-semicontinuous envelopes. These are Borel operations. Apply the explicit countable metric-validity tests in Proposition 12.18 to the resulting lower-semicontinuous function. Use it when the tests pass and the fixed fallback metric otherwise. This gives a Borel rule whose values are separating lower-semicontinuous extended metrics on every input.

On every good base field h, the deterministic weighted-recovery theorem identifies these formulas at $g = h + f$, simultaneously for every continuous f, with the actual internal metrics of $e^{\xi f} \cdot E_h$, so the validity tests pass. Exhaust the space only after local reconstruction. Thus the formula depends on g itself and agrees for any two decompositions into a good field plus a continuous function; no measurability of an orbit saturation or choice of a decomposition is required. This supplies simultaneous Weyl scaling also for a possibly dependent random continuous perturbation. Locality is inherited from the buffered formulas, and exact coordinate covariance extends by the same Weyl identity. A finite connection between two deterministic points in any connected open domain follows by concatenating deterministic-point connections in finitely many buffered cubes; it implies continuum finiteness. Continuous weights preserve its finite length because its image is compact.

For fixed r, any limit of the bounded positive factors $ra_{\varepsilon/r}/a_\varepsilon$ is identified by the exact LFPP coordinate change and Proposition 13.33 as $r^{\xi Q}$, evaluating at one positive finite pair. Every subsequential factor limit has that value, so the factors converge. Put $r = \lambda^{-1}$. □

13.9 Thin-shell restoration for a one-sided competitor

For the remainder let E be the canonical strong metric and D an arbitrary strong competitor with the same parameters. Corollary 13.9 gives $E \le CD$. We prove rigidity from this one-sided comparison. No upper point estimate for D will be used.

Lemma 13.36. *Fix nested cubes with radii $r_j \uparrow r_\infty$ and gaps $r_{j+1} - r_j \asymp 2^{-j}$, all inside one annular working region. For every $\epsilon > 0$ there are local costs $H_{j,\epsilon}$ such that every two D-arms crossing the jth thin shell have an E-connector within that same shell, between selected actual contacts, with cost*

$$\epsilon\{\mathrm{len}_D(A_-) + \mathrm{len}_D(A_+)\} + sH_{j,\epsilon}. \tag{13.33}$$

For some $p \in (0, 1]$ and $B < \infty$, independent of ϵ,

$$\mathbb{E}[H_{j,\epsilon}^p] \le C_\epsilon e^{Bj}.$$

Proof. An arbitrary D-arm has E-length at most C times its D-length. Anchor the canonical all-arm core to a fixed rational point by joining that point to a second rational point across the local shell, applying (13.1) to the arbitrary arm and a crossing portion of this reference path, and returning along the reference path. The resulting anchor cost is at most $\epsilon \operatorname{len}_D(A) + s_\rho Z_\epsilon$. A small fixed $2p$ moment follows from the deterministic-point positive moment and the all-moment core bound. The moment order is chosen independently of ϵ.

Put a rational mesh of spacing comparable to 2^{-j} on the midpoint cube boundary of the thin shell. Its nearest-neighbor graph is connected with $O(2^{(d-1)j})$ edges. Every shell-crossing arm crosses a small annulus about a nearby mesh point, with its enlarged buffer strictly in the thin shell. Anchor that crossing arm and connect all mesh anchors by finite rational internal E-paths in similarly buffered small boxes. This one connected network gives (13.33). All images stay in the hollow shell, including at corners. Dividing each cost by its own local unit gives a uniform $2p$ moment. The Gaussian increment from the local scale to the fixed working scale and exact scaling contribute only $e^{B_0 j}$ to its p moment. Since $p \le 1$, the p moment of a sum is at most the sum of the p moments. The number of mesh terms changes the exponential constant but not its independence from ϵ.

Use actual local infima, with an added unit of slack, to make these costs local functions of the canonical metric. For example the least all-arm excess is the supremum over finite Lipschitz crossing arms of the positive part of their internal connection cost minus the charged arm lengths. Its strict superlevel events are analytic projections of Borel curve tests. Hence it is measurable in the completed local field sigma field and is dominated by the retained core variable. No auxiliary randomness remains in $H_{j,\epsilon}$. □

Lemma 13.37. *Fix $0 < c < C$ and choose ϵ so that $q = \epsilon/(c+\epsilon)$ satisfies $q^p e^B < 1$. Set $H = \sum_{j\ge 0} q^j H_{j,\epsilon}$, which is finite almost surely. Consider a continuous path carrying any finite nonatomic chronological measure μ, equal to competitor arclength on the untouched thin-shell arms. Its mass elsewhere may be arbitrary. If T_j is its first-to-last mass span of the inner collector, the shell contacts bracket an interval J_j with $\mu(J_j) \ge T_j$.*

1. *Some shell connector has E-cost at most $c\mu(J_j) + As$, where $A = 1 + (1-q)H$. Its image is outside the discount support and deletion of J_j removes all visits to that support.*
2. *If no interval of mass at least bs, $0 < b \le 1$, has an E shortcut in the working region of ratio less than c, then*

$$T_0 \le \left(\frac{1}{1-q} + \frac{H}{c+\epsilon}\right) s.$$

Proof. The shell-arm masses sum to at most $T_{j+1} - T_j$. If every connector cost exceeded $c\mu(J_j) + As$, (13.33) would give

$$T_j < qT_{j+1} + \frac{H_{j,\epsilon} - A}{c+\epsilon} s.$$

Iterate and let the number of shells tend to infinity. The terminal term vanishes because μ is finite. The resulting upper bound for T_0 is $\frac{s}{c+\epsilon}(H - A/(1-q)) < 0$, a contradiction. For the second assertion, absence of a shortcut yields $T_j \le bs + qT_{j+1} + H_{j,\epsilon} s/(c+\epsilon)$. Iterate, let the terminal term vanish, and use $b \le 1$. □

The useful bookkeeping consequence is exact. If an original competitor path of length L has already had canonical paths inserted in disjoint working regions, give retained original pieces their original D time and inserted pieces their E-arclength divided by c. The measure is finite even when

an inserted path has infinite D-length. At a fresh target region all shell arms are retained original pieces. If its deletion swallows original time t and inserted canonical cost e, Lemma 13.37 replaces them at cost at most $ct + e + As$, whereas their former comparison ledger was $Ct + e$. Restoration increases the ledger by at most As and preserves every previous saving. Contacts are selected on the current hybrid curve, rather than on possibly erased original intervals.

13.10 Protected trees with small unmarked pieces

Lemma 13.38. *Fix a compact buffer $K \Subset O$, $h > 0$, $L < \infty$, and $\omega(r) \downarrow 0$. Suppose $\alpha_n : [-h, h] \to K$ converge uniformly Euclideanly, are E-Lipschitz with constant L, and satisfy*

$$|s - t| \le \omega(|\alpha_n(s) - \alpha_n(t)|).$$

Then $E(\alpha_n(0), \alpha(0); O) \to 0$.

Proof. Lower semicontinuity makes the limit E-Lipschitz, and the inverse modulus makes it non-constant. Apply Lemma 13.12 at $x = \alpha(0)$ with path bound $2Lh$. Choose its outer radius R so small that $L\omega(3R)$ is below half the target error and the two endpoints of each late α_n lie outside the outer ball. The selected contact has small intrinsic access to x. Its parameter distance from 0 is at most $\omega(3R)$, so the retained stub to $\alpha_n(0)$ has small length as well. □

Lemma 13.39. *For $d \ge 3$, assume a strict local mark of the form* (13.13) *has positive probability. There are arbitrarily high-probability annulus-local events with this consequence: every finite canonical geodesic passing through the inner test region, with endpoints outside the outer buffer, admits a simple replacement arc of length at most Ks_r, joining two ordered contacts, whose interior misses the entire old geodesic and which traverses a protected marked middle. The constants and determining annulus are fixed before the geodesic is chosen.*

Proof. Work in normalized coordinates. Let $\Omega = \{1 < |x|_\infty < 4\}$ and enlarge its field buffer slightly, including the Euclidean unit averaging sphere. Collapse Ω^c to a cemetery point $*$. The quotient is compact; a useful metric is

$$d_Y(x, y) = \min\{|x - y|, \operatorname{dist}(x, \Omega^c) + \operatorname{dist}(y, \Omega^c)\}, \qquad d_Y(x, *) = \operatorname{dist}(x, \Omega^c).$$

Choose nested annular buffers $O \Subset K_1^\circ \Subset K_2^\circ \Subset \Omega$ and a continuous cutoff χ, equal to one on K_1 and supported in K_2°. Quantile a thin-shell envelope from Lemma 13.37, applied with $D = E$ and $c < 1$. A canonical geodesic has no such strict self-shortcut, so its first-to-last outer visit span is at most As_r on this local event.

Local lower shell bounds and occupation give deterministic $M, a, \beta > 0$, outside an arbitrarily small additional probability, so the quotient curve parameterized by normalized time satisfies

$$d_Y(\alpha(s), \alpha(t)) \le M|s - t|^a, \qquad \chi(\alpha(s))\chi(\alpha(t))|s - t| \le M d_Y(\alpha(s), \alpha(t))^\beta. \tag{13.34}$$

Its time interval is extended at $*$ to $[0, A]$. Choose two separated crossing times in a fixed port band, with fixed positive-size windows in that band. Lower shell bounds give a common positive window size. Include the two times as compact coordinates. The class of curves obeying (13.34) and these closed window conditions is compact, independently of the field. For fixed E, impose Lipschitz bounds for the actual internal metric E^Ω on the two windows; intrinsic lower semicontinuity gives a closed compact subfamily. The windows of an original geodesic qualify because their whole subarcs

lie in Ω, and their arclength bounds the corresponding internal endpoint distances. All tests defining this subfamily are therefore measurable from the fixed local annular field.

The trace in K_1 is homeomorphic to a closed subset of the parameter interval, by the second inequality in (13.34). Such a compact set contains no open ball, as proved in Theorem B.8. In a prescribed trial region inside O, each such trace misses some rational closed box by a positive distance. This avoidance persists in a neighborhood of the quotient curve. Compactness gives a finite deterministic box catalog, chosen before the metric. Amplify local marked configurations in every catalog box as in Lemma 13.28, with total failure as small as desired.

Fix the metric, the resulting finite marks, and one obstacle curve. Select a missed box. Forbid the old trace outside the left attachment window and the protected-right portion of the chosen mark. The forbidden compact set is a closed subset of a finite union of compact sets homeomorphic to closed interval subsets; its complement in O is connected by Theorem B.8, including its relatively closed subset clause. The left window anchor is nonsingular and separated from the forbidden set. Intrinsic rational density followed by rational connectivity provides a finite E-path from it to the left rational marked endpoint in that complement. Trim at the first old-trace contact when read from the mark, and at the last appropriate mark contact; erase loops and append permitted short stubs if needed. For the right connector also forbid the entire left connector, the old trace outside the right window, and the protected-left marked portion. The same nonseparation statement gives a disjoint finite right connector. Trimming produces the asserted simple replacement with unique ordered old-trace contacts and the full protected middle.

Its least possible total cost is finite and upper semicontinuous on the compact obstacle class. To verify the latter, choose a nearly minimizing connector pair for a limiting obstacle. Their compact images have positive clearance from each other's images and their forbidden trace portions. Retain disjoint open tubes with these clearances. Uniform quotient convergence preserves the forbidden-trace clearances. Lemma 13.38 connects each moving window anchor to its limiting anchor at cost tending to zero within its tube. Concatenating with the fixed connectors proves the upper limit inequality. Hence the cost supremum over the compact obstacle class is finite.

This supremum is a local measurable random variable. Express each connector cost as the infimum of the actual intrinsic distances over countably many pairs of disjoint rational tubes satisfying the strict avoidance tests. Mark curves can be retained by a universally measurable selector from their Borel Lipschitz path codes. Under the local law take a Borel version of this selector before forming the joint metric-obstacle relation. The existence of an obstacle whose cost exceeds a threshold is then an analytic projection; completed measurability suffices. Quantile the finite supremum to get deterministic K. Choose, in order, the desired success probability, the span cap, the moduli and window size, the finite catalog, the mark-amplification trials, and finally K. The outer determining buffer never changes. This proves the required uniform local event. □

Proposition 13.40. *Under the strict-discrepancy hypothesis, for finitely many nonsingular ports in a connected open region U and every $\delta > 0$ there is a finite embedded E-rectifiable tree through the ports with protected middles satisfying* (13.13)*, such that*

$$E(u, v; U) < \delta \quad \textit{whenever the tree arc from } u \textit{ to } v \textit{ contains no full protected middle.} \tag{13.35}$$

Proof. First promote Lemma 13.39 to random geodesic points. At stage n use the common grid of mesh $2^{-2n}/16$ and candidate radii 2^{-m}, $n \le m \le 2n$. Annular concentration and a grid union bound give high-density successful local detour events at every grid center, eventually almost surely. All inverse crossing moments also give, for any fixed $\eta > 0$,

$$X_r(w) \ge r^{\eta} s_r(w)$$

simultaneously at every such center and candidate radius, eventually almost surely. Fix an interior point x of a canonical geodesic J and choose the nearest common grid center $w_n(x)$. It is within $r/16$ of x at every candidate radius. An actual subarm of J crosses that center's test shell, so Lemma 13.13 gives

$$r^{\eta} s_r(w_n(x)) \le X_r(w_n(x)) \le \operatorname{span}_J(B_{3r}(x)) \le C_{J,x} r^{\beta}.$$

Choose $\eta < \beta$. Consequently $K s_r(w_n(x)) \to 0$ on successful annuli. Their detours lie in shrinking Euclidean neighborhoods of x, and their contact times lie in its shrinking original geodesic visit span. Thus every geodesic interior point has arbitrarily cheap protected replacements in arbitrarily short intrinsic subarc neighborhoods. This argument uses the actual grid-center unit; it does not transport a measurable event to a random center.

Join the ports by finitely many finite paths compactly in U. Partition into short pieces and replace them by ambient geodesics; lower exit barriers keep these pieces in U. Add ports one at a time, stopping each new concatenation at first contact with the existing tree and erasing loops. A finite concatenation of simple arcs can be loop-erased by successive first/last intersection deletion while retaining finitely many subarcs. The resulting embedded tree has edges that are finite concatenations of geodesic subarcs.

Subdivide it so every edge has length below $\delta/10$ and every change of geodesic piece is a vertex. Choose an interior point of each edge and disjoint Euclidean replacement neighborhoods meeting no other edge or vertex. Use the preceding cheap replacement in each neighborhood, with cost below $\delta/10$. This preserves tree topology and gives every final subdivided edge a protected middle and length below $\delta/5$. A tree arc containing no full middle can pass through at most one subdivision vertex; otherwise it traverses a full marked edge. It therefore lies in at most two final edges and has length below $2\delta/5 < \delta$, proving (13.35). □

13.11 Germ gates and a pure-discount contract

Lemma 13.41. *Let $T \Subset U$ be a finite embedded E-rectifiable tree. Fix finitely many degree-two cut points, an original path-length bound, and $\eta > 0$. There is an arbitrarily thin open neighborhood N of T with finite germ charts such that:*

1. *a path approaching N from outside the chart buffers has a selected contact with E-access at most η to a tree anchor;*
2. *each cut has a small trigger ball with a germ chart giving access at most η to the cut;*
3. *outside its trigger, the two branches of N have disjoint labels; an exterior excursion changing labels without hitting the trigger has two ordered actual contacts with access at most η to opposite-branch anchors.*

All charts can be confined inside any prescribed open neighborhood of T.

Proof. First choose the cut charts using Lemma 13.12, in disjoint short edge-stub neighborhoods, and choose a trigger radius ρ below one quarter of the resulting inner radius. Outside $B_{\rho/2}$ the two compact tree branches have positive Euclidean separation. Give them disjoint open neighborhoods and cover the tree outside $B_{3\rho/4}$ by finitely many germ inner balls whose outer balls and connection regions stay in the appropriate branch neighborhood. Their union with $B_{7\rho/8}$ is a neighborhood of the whole tree. Independently take a finite global germ cover. Only now choose N, with compact closure inside all these finite unions and the prescribed thin tube.

For an exterior excursion with opposite labels, the chart outer balls at its two endpoints are disjoint. Its initial portion exits the first, and its terminal portion enters the second. Applying the corresponding germ charts gives ordered selected contacts on the excursion itself. The half-, three-quarter-, and seven-eighth-trigger radii ensure no uncovered annular gap near the trigger. □

Give T its intrinsic edge metric d_T. Compactness, separation and lower semicontinuity imply

$$\omega_T(t) := \sup\{d_T(u,v) : u,v \in T,\ E(u,v;U) \le t\} \longrightarrow 0. \tag{13.36}$$

Indeed a violating pair has Euclidean convergent subsequences with zero limiting E distance, hence the same limiting tree point.

Choose $\lambda_n \downarrow 0$ and $\eta_n = o(\lambda_n)$ and take the gate construction at each n. If an exterior excursion of original D-length $O(\lambda_n)$ changes branch labels, its selected opposite anchors satisfy $E(k_-, k_+; U) \le O(\lambda_n) + 2\eta_n$. Their tree arc passes through the cut. Under a weight equal to λ_n on T and at most one everywhere, either selected contact therefore reaches that cut at weighted E-cost

$$\eta_n + \lambda_n \omega_T(O(\lambda_n) + 2\eta_n) = o(\lambda_n). \tag{13.37}$$

This is a selected-path statement. It asserts no access from arbitrary boundary points.

Proposition 13.42. *Assume a strict low-ratio discrepancy between E and D. For every prescribed high success probability, there is a deterministic finite family of nonnegative smooth annular profiles with the following local contract. Each has a fixed insertion gain on a confined baseline crossing, an active ratio $c_1 < C$, a threshold $b > 0$, and an inactive cost $\epsilon_0 > 0$.*

Every finite multiply discounted competitor path with endpoints outside the working region, visiting the local positive support, either has an interval of actual weighted length above bs, with endpoints inside the working region and a canonical shortcut there of ratio below c_1, or has two ordered actual contacts bracketing all local positive-support visits, with original canonical connection cost below $\epsilon_0 s$. An active interval contains a positive own-region crossing charge or lies wholly in its own region and supplies its own threshold charge. The contract is uniform over other discounts in disjoint working regions, and remains valid if exterior chronological mass is arbitrary and finite.

Proof. Set the local unit to one. Use collector B_3, working ball B_4, and annular path buffer $U_c = \{1 < |x|_\infty < 15/4\}$. Fix a smooth baseline weight $w_0 \in [a_0, 1]$ whose nontrivial compact carrier G is inside $\{7/4 < |x|_\infty < 12/5\}$ and gives a strict gain on a fixed confined crossing. Use attachment band $\{13/5 < |x|_\infty < 14/5\}$. Quantile positive lower crossings of the gap from this band to G, and untouched inner and outer exit bands. The thin shells expand from radius $7/2$ to $15/4$, outside the collector. Lemma 13.37 supplies a local high-probability no-shortcut span cap A, independent of the future discount depth and threshold.

For each $M < \infty$ and $\eta_0 > 0$ there is a finite rational menu in the attachment band such that every canonical crossing arm of length at most M has a selected contact with intrinsic E-access below η_0 to a menu point. To prove this, parametrize arms with speed at most one on a common compact interval and retain a terminal crossing time. They form a compact Euclidean curve class by lower barriers and intrinsic lower semicontinuity. For one limiting arm choose a nonsingular interior point, a short rational access, and a germ chart. A uniform neighborhood of the arm retains the chart crossing. Finitely many such neighborhoods cover the class. All data are local.

Take $M = CA/a_0$ and η_0 much smaller than the untouched weighted gap charge $\delta_{\mathrm{gap}} = a_0 a_{\mathrm{gap}}$. By Proposition 13.40 construct a finite protected tree through this menu, compactly inside $\{3/2 < |x|_\infty < 29/10\}$, with no-mark endpoint image below $\epsilon_0/4$. Write its total length as L_T, and choose

$T_0 > 2L_T/c_1$. Take $\lambda_n \downarrow 0$, $\eta_n = o(\lambda_n)$ and gate charts for a bound exceeding both CA/a_0 and CT_0. After choosing the charts choose N_n shrinking to T, and a smooth weight

$$\lambda_n \le w_n \le w_0 \le 1, \qquad w_n = \lambda_n \text{ on } T, \qquad w_n = w_0 \text{ outside } N_n.$$

Its compact positive-discount carrier is contained in $G \cup \overline{N}_n$.

Consider an inactive path. The thin-shell cap bounds its entire collector visit span by A. Read it from each end up to first contact with $G \cup \overline{N}_n$. If the first contact is with G, an incoming attachment-band arm avoiding N_n has canonical length at most CA/a_0. The menu gives a selected *core contact* of access at most η_0 to the tree. Between it and G a confined untouched gap costs at least δ_{gap}. If the first contact is with $\overline{N}_n$, the first germ chart supplies a selected *germ contact* of access at most η_n. All such contacts lie inside the collector and bracket every positive-support visit. Ties use the germ case.

If either end is of core type, the interval between the two selected contacts has weighted D length at least δ_{gap}, while the canonical route through their anchors and T costs at most $2\eta_0 + 2\eta_n + \lambda_n L_T < c_1\delta_{\mathrm{gap}}/2$ for large n. This is active. Thus both contacts of an inactive path are germs. Their weighted canonical tree connection costs at most $2\eta_n + \lambda_n L_T$. If the intervening weighted D length exceeds $T_0\lambda_n$, it is active. Otherwise it is at most $T_0\lambda_n$.

For large n the untouched exit barriers confine this short middle to U_c. Its original D length is at most T_0, and its original length outside N_n is $O(\lambda_n)$. If inactive restoration failed for infinitely many n, compactness of these original bounded-speed paths would give a limiting D-path in T. Every exterior excursion has an endpoint near T and original length tending to zero; canonical lower semicontinuity shows that the entire limit lies in T. Its endpoint anchors converge in the finite tree topology.

If the tree arc between the limiting endpoints contains no full protected mark, its intrinsic canonical endpoint distance is below $\epsilon_0/4$. Vanishing endpoint access and short tree tails then give original canonical restoration below ϵ_0. Otherwise it contains a full mark $[u_R, v_R]$. Choose ordered transition intervals across its two cuts, in disjoint stub neighborhoods. At a trigger hit the germ chart gives weighted access $\eta_n = o(\lambda_n)$ to the cut. A bypass changing branch labels gives (13.37). The selected contacts are chronologically ordered and converge to u_R, v_R. Lower semicontinuity makes their original D distance greater than a_R for all large n. Their intervening weighted length is at least $\lambda_n a_R$, while the route through the protected middle costs

$$o(\lambda_n) + \lambda_n \operatorname{len}_E(R) < c_1\lambda_n a_R.$$

This is active, a contradiction. Choose b below half the minimum of 1, δ_{gap}, $T_0\lambda_n$, and $\lambda_n \min_R a_R$, for one sufficiently large fixed n.

Thin-shell active intervals contain an untouched collector-to-shell crossing. Core intervals contain their gap crossing. A marked or long germ interval lying in the working ball supplies its own threshold mass. If such an interval exits the ball, its endpoints lie in the fixed interior set $B_{19/5}$, so it contains an untouched confined exit crossing. Thus all activity types have a positive charge inside their own working region. Decrease b, if necessary, below half the fixed positive lower quantiles for the untouched collector-to-thin-shell and exit crossings. Every active interval then has weighted D-mass at least b inside its own working region, in the current unit normalization.

For locality, allow any finite nonatomic measure μ on the path parameter which dominates the local weighted competitor arclength on intervals in the annular buffer. Outside that buffer its mass is arbitrary. Every shell, attachment, gap and germ-arm estimate above uses only local arclength bounded by μ. If the selected middle has mass at most $T_0\lambda_n$, exit barriers confine it to the annular

buffer, where the compactness argument applies. Active ratios only improve when extra mass is present. This proves the claimed exterior-mass formulation and allows arbitrary changes outside the determining annulus.

The construction is measurable from local data. Lipschitz arms, finite port covers, protected paths, compact lower tests and chart certificates have the Borel or analytic codes described in Lemma 13.28. Admissibility of (P, μ) is tested on rational intervals in compact rational subdomains; nonatomicity and lower semicontinuity extend the tests to all intervals and imply local length domination. Active and inactive strict tests can be restricted to rational times by local metric continuity. Failure is an analytic projection, so the universal contract is completed-local measurable. Approximate the logarithmic discount by a nonnegative smooth profile supported inside its old positive support, first truncating a tiny level and then using rational smooth bumps. Relative weight errors preserve strict ratios and thresholds; the old inactive bracketing still covers the new support. A countable deterministic family covers the high-probability construction event, and a finite subfamily has the prescribed success probability. Exact covariance gives the same family at every scale. □

13.12 Strong witnesses and the final one-sided count

Lemma 13.43. *Let $P : [0, L] \to K$ be a unit-speed constrained D-geodesic with $E^{[K]}(P(0), P(L)) \geq CL - \Delta$. For an interval $I = [a, b]$ put*

$$d(I) = C(b - a) - E^{[K]}(P(a), P(b)) \geq 0.$$

Disjoint-interior intervals have total defect at most Δ. For every $\alpha > 0$, the union of intervals with $d(I) > \alpha|I|$ has length at most $5\Delta/\alpha$. Outside this union every interval containing the given time has endpoint $E^{[K]}$ distance at least $(C - \alpha)$ times its length.

Proof. Concatenate endpoint connections of the disjoint intervals with the retained original pieces; the triangle inequality gives the total defect bound. Endpoint distances along P are continuous because $E \leq CD$. Strict bad intervals can therefore be slightly enlarged, and rational endpoints give a countable family with the same union. Here is the required covering argument. Group its intervals by dyadic length, starting with the largest group. In each group enumerate the intervals and retain one exactly when its interior is disjoint from the previously retained interiors. Every rejected interval meets a retained interval of at least half its length, and is contained in that interval's fivefold concentric dilation. Thus the fivefold dilations cover the original union. The retained intervals have disjoint interiors; every finite subfamily has total defect at most Δ, so their lengths sum to at most Δ/α. Their dilations have total length at most $5\Delta/\alpha$, giving the claim. The last assertion is the definition of the complement. □

At a time outside the bad union, canonical connectors give span control for the competitor path. If $c_g = C - \alpha$, incoming and outgoing arms of nested regions satisfy

$$c_g T_{\text{in}} \leq \epsilon(T_{\text{out}} - T_{\text{in}}) + sY, \qquad T_{\text{in}} \leq \frac{\epsilon}{c_g + \epsilon} T_{\text{out}} + \frac{sY}{c_g + \epsilon}.$$

On a symmetric connector event the additive term is absent and the coefficient is strictly below one. Thus the envelope and occupation arguments apply at every good time, with constants depending on c_g, not on Δ.

Lemma 13.44. *If C is the optimal deterministic constant in $E \le CD$, there are $a_H > 0$, $A_H < \infty$ and $p_* > 0$ such that at every deterministic radius and every $c \in [C/2, C)$ a capped annular D-witness with ratio above c has probability at least p_*.*

Proof. Use $\alpha = C/2$ in Lemma 13.43. Fix the high-density good-time span envelope and a lower crossing quantile, then take A_H twice the envelope cap. Put $\gamma = (C - c)a_H/A_H$. Optimality supplies a positive-probability buffered macroscopic witness of defect less than $\gamma L/128$. Its bad times have total length below $L/32$. If capped witnesses had probability below a sufficiently small fixed p_*, exact covariance would make this true at every microscopic radius. The local no-witness and lower-crossing event has high probability, independently of $C - c$. Annular concentration, intersected with the span-envelope levels, covers all good central times by successful intervals. A finite subcover and two-color selection give disjoint intervals of total length at least $L/8$. Each saves γ times its length, by the capped-witness failure calculation in Lemma 13.29. The resulting canonical connection costs at most $(C - \gamma/8)L$, contradicting the macroscopic threshold $C - \gamma/128$. Hence the stated uniform p_* exists. □

Lemma 13.45. *A sufficiently near-maximal capped witness gives, outside one common small-probability event, one of polynomially many shift-invariant compact endpoint problems satisfying*

$$\ell_0 S \le L = D^{[K]}(K_1(h), K_2(h)) \le MS, \qquad E^{[K]}(K_1(h), K_2(h)) \ge CL - \delta S. \tag{13.38}$$

Its minimizing competitor geodesic is attained. No deterministic endpoint finiteness for D is needed.

Proof. Take two separated fixed radial bands strictly inside the witness middle annulus. Impose a positive D crossing lower bound $a_0 S$ for each. If $10\Delta/C < a_0 S$, each crossing interval contains a good time of Lemma 13.43. These two good contacts are separated from each other and from the old parent endpoints. Fix disjoint *compact* rational polycubical endpoint collars V_1, V_2, with the two contact bands in their interiors, and open rational buffers U_i such that $V_i \subset U_i \Subset K^\circ$. Good-time occupation and the canonical rational-core attachment give, at accuracy ηS, rational points q_i from deterministic polynomial menus in the interiors of these collars, with

$$E(x_i, q_i; U_i) \le \eta S.$$

Define

$$K_i(h) = \{x \in V_i : E_h(x, q_i; U_i) \le \eta S_h\}.$$

They are nonempty Euclidean compact sets by intrinsic lower semicontinuity, with canonical intrinsic diameter at most $2\eta S$. The original witness gives a finite competitor connection, and the sets lie in the two fixed Euclideanly separated collars. A common lower crossing quantile between those collars gives $L \ge \ell_0 S$ for a deterministic $\ell_0 > 0$, outside an arbitrarily small prescribed additional failure probability. Compact path extraction gives an attaining constrained competitor geodesic and $L \le \ell$, where ℓ is the original intervening witness length. For arbitrary $u_i \in K_i$,

$$E^{[K]}(u_1, u_2) \ge E^{[K]}(x_1, x_2) - 4\eta S \ge C\ell - \Delta - 4\eta S \ge CL - \Delta - 4\eta S.$$

Choose $\Delta + 4\eta S \le \delta S$. All microscopic supports avoid $U_1 \cup U_2$ and the macro averaging sphere defining S. Weyl scaling therefore leaves both endpoint sets exactly unchanged under every counting shift. Common repinning factors also cancel in their definition. □

Theorem 13.46. *For every $d \ge 3$ and every parameter with $Q > 0$, the five strong axioms have a unique metric assignment up to a deterministic positive scalar. With the internal point-median normalization the scalar is fixed, and the assignment is the limit in Theorem 13.35.*

Proof. Take the optimal C in $E \le CD$. It is positive: continuum finiteness supplies a finite connection with distinct endpoints, whose canonical distance is positive.

If $E = CD$ fails, there is an E-finite pair with $E < c'D$ for some deterministic $c' < C$, including the possibility $D = \infty$. Indeed, $E = \infty$ forces $D = \infty$ by $E \le CD$, so any strict discrepancy lies in the finite component of E. Choose numbers $c' < c'' < c_{\mathrm{mark}} < c_1 < C$ and a finite positive threshold a such that

$$E(x, y) < c'a, \qquad a < D(x, y).$$

When $D(x, y) = \infty$, take any sufficiently large finite a. Lower semicontinuity supplies Euclidean neighborhoods of x and y on whose product $D > a$. Choose a finite E-path of length below $c'a$, and use intrinsic rational density near its endpoints to attach rational points q_1, q_2 at total additional cost below $(c'' - c')a$. Thus there is a finite rational-endpoint path P of length below $c''a$ and $D(q_1, q_2) > a$. No finiteness of either deterministic D point has been asserted.

Here is the localization needed for a field-measurable finite trial. Confine P to a compact rational buffer. Coercivity supplies a sufficiently distant outer exit strip whose D distance from small endpoint neighborhoods exceeds a. Choose a rational center outside a still larger compact set, so that a small inner exit strip about that center also has endpoint distance above a and misses P. A sufficiently large rational outer radius about that center encloses P and keeps its outer exit strip beyond the first compact set. This gives a fixed rational annular buffer: any path from an endpoint that leaves it must first pay one of the two exit costs above a. The constrained endpoint lower test and these confined exit-strip tests are consequently local and imply the global lower bound. All choices can be made with strict margins. Moving a short distance into the two ends of P preserves a lower threshold slightly below a, by lower semicontinuity; deleting those short end stubs and slightly increasing c'' therefore gives a protected middle with ratio below c_{mark}. Countably many rational buffers, endpoints, thresholds, stub lengths, and margins suffice, so one such local configuration has positive probability. Translate and dilate that fixed annular template to locate its trials inside the designated boxes. Exact covariance transports its positive probability to every scale.

Proposition 13.42 now gives the finite pure-discount family, with fixed insertion gain $g > 0$, small inactive cost $\epsilon_0 < Cg/16$, threshold $b > 0$, and the thin-shell general-restoration quantile A_*. The latter is chosen before the tree and discount depth.

Take the endpoint problem (13.38). Along its selected competitor geodesic, total defect is at most δS. Lemma 13.43 removes only $O(\delta S)$ bad time. The common good-time covers, span envelopes, local profile successes and crossing lower bounds yield a bin of N visited labels with the same inequalities (13.26). There are at most $C\varepsilon^{-2d}$ deterministic labels and polynomially many endpoint-menu pairs. Fix $\delta = \varepsilon^{B_0+1}$. Every fixed-size subset is now allowed; the thin-shell local contract is uniform in all other disjoint-region discounts, so mesoscopic sparsity is unnecessary.

For each k-subset I, following the original path through the baseline discounts gives directly

$$D^{[K]}_{h-f_I}(K_1, K_2) \le L - gkqS.$$

Remove a label whenever the consistently selected discounted minimizer avoids its positive support, recomputing after removal. Since the discounts are pure, its minimum is unchanged. The final output Z visits all retained supports and keeps the full gain.

Restore active labels using the hybrid ledger following Lemma 13.37. Restore inactive labels using their cheap canonical contacts from Proposition 13.42. Fresh contacts must be chosen after earlier deletions: use the remaining original incoming chunk preceding the first surviving local support visit, and the outgoing chunk following the last. Their outer endpoints are outside the target working region because inserted images lie in disjoint regions. The original path's no-active

span cap bounds each local arm. If either fresh contact were of core type, its original gap crossing would force an active shortcut on the original interval between the fresh contacts. Thus both are germs. Absence of activity forces that original middle to have weighted length at most $T_0\lambda$, and the gate argument gives the original canonical connection of cost below $\epsilon_0 s_i$. These contacts survive, so their current hybrid interval may be deleted. No inactivity or competitor-rectifiability claim is made about the hybrid curve itself.

The resulting upper ledger is

$$CL_{\mathrm{discounted}} + 2A_* a_Z qS + 2\epsilon_0(|Z| - a_Z)qS.$$

Comparing with (13.38) gives the survival bound (13.28), with A_* in place of A_{rest}. Hence fixed positive κ_1, κ_2 give $|Z| \geq \kappa_1 k$ and at least $\kappa_2|Z|$ active labels. The same preimage calculation gives (13.30).

Under the shift $g = h + f_Z$ all outputs use one selected D_h geodesic: the compact canonical endpoint sets are invariant, and the discounted metric is D_h. Local units are unchanged. Pure-discount monotonicity gives $L_g \geq L_h$ in the unpinned representative. Every visited label pays a confined untouched own-region crossing, so there are $O(q^{-1})$ visited labels.

Active intervals need not remain inside their working regions. Their shortcut images do, and each interval contains an own-region charge of length at least bqS by Proposition 13.42. Charges at distinct labels are disjoint. A finite family of intervals containing disjoint charges of total length $nbqS$ has a disjoint-interior subfamily of total interval length at least $nbqS/5$: keep a longest remaining interval and delete every intersecting interval; every deleted interval lies in the triple dilation of the kept one. The union of the intervals has length at least the total charge, proving the estimate.

Fix one shifted target output Z_0, of size at most k. If n other labels are active in any shifted output, insert the disjoint shortcuts just selected. Their images avoid target working regions and save at least $(C - c_1)nbqS/5$ in the comparison ledger. Restore the targets by the hybrid thin-shell lemma. This adds at most $2A_* kqS$ and preserves all prior savings, even if inserted paths have infinite competitor length. Comparing to the target endpoint lower bound gives

$$\frac{(C - c_1)b}{5} nqS \leq 2A_* kqS + \delta S.$$

Thus the shifted active universe has $O(k)$ labels and (13.31) again follows. All shifts are deterministic members of the finite profile family. Their scale-invariant Cameron–Martin norms satisfy explicitly

$$\|\Pi f_Z\|_{\mathcal{H}} \leq \sum_{i \in Z} \|\Pi f_i\|_{\mathcal{H}} \leq kM_{\mathcal{H}},$$

where Π subtracts the pinning average. Gaussian change of measure therefore has a constant depending on k, α but not on the microscopic radius. The same finite-sum Hölder calculation as in Proposition 13.31 gives

$$q^m \mathbb{E} A_m \leq C_{\alpha,k} \varepsilon^{\gamma m}, \qquad \gamma = \alpha\kappa_2\theta_0 - 2d(1 - \alpha) > 0.$$

Choose k large enough to dominate the polynomial endpoint-menu exponent, and then let $\varepsilon \downarrow 0$. The lower count suppresses capped near-maximal witnesses, contradicting the radius-uniform positive probability of Lemma 13.44. Hence strict discrepancy is impossible.

This excludes every strict discrepancy at an E-finite pair, including pairs with infinite D distance, by the countable rational/compact localization above. If $E(x, y) = \infty$, the initial comparison already forces $D(x, y) = \infty$. Thus equality holds simultaneously for all pairs. Simultaneous Weyl scaling transfers it to every continuous perturbation of a good base field. Applying the conclusion to any two competitors gives the theorem. □

Completion of the proof of Theorem 1.2. The exponent comparison and Theorem 10.3 give the uniform estimates of Theorem 11.1 throughout the range $Q > 0$. Proposition 12.3 constructs a subsequential extended metric, and the intrinsic and weighted reconstruction results identify its actual internal distances and continuous Weyl perturbations. Proposition 13.10 removes residual randomness. Proposition 13.31, Proposition 13.33, and Theorem 13.35 give one normalized strong assignment, full convergence, the actual internal median, and regular variation of the normalizer. Theorem 13.46 gives uniqueness among all assignments in Definition 1.1. Finally, Proposition 12.7, Corollary 12.12, and Lemma 12.2 give the stated finite-component geometry, deterministic-point finiteness, and Euclidean boundedness of metric-bounded sets. Every conclusion of the theorem follows. □

A Compactness, conditional laws, and measurable choices

This appendix records the general probabilistic and coding facts used in the construction. All measurable spaces carrying random coordinates are complete separable metric spaces with their Borel sigma fields; completion is taken with respect to the probability law under discussion. The countable coordinate encodings of a field or white noise use a countable determining collection of test functions. Standard operations of measure theory, including Lebesgue integration, conditional expectation, and the Radon–Nikodym theorem, have their usual meanings. The arguments below supply the compactness, representation, projection, and selection results needed in addition to those operations. General treatments of conditional distributions and probability representations are given in Kallenberg [13]; analytic measurability and uniformization are treated in Kechris [14]. The forms used here are proved below.

Lemma A.1 (A complete separable model for the field coordinates)**.** *The fields and white noises used in this paper can be realized in complete separable spaces of local negative Sobolev distributions. Their Borel sigma fields are generated by countably many smooth test pairings. The same spaces contain every continuous perturbation of the field. Thus retaining all field coordinates in a coupled extraction retains the entire field.*

Proof. Choose a nested smooth exhaustion χ_j of the domain, each compactly supported in the interior of a larger rectangular torus, with $\chi_j = 1$ on the jth compact set and $\chi_j\chi_{j+1} = \chi_j$. Fix an integer s larger than half the dimension of the domain. A torus distribution with Fourier coefficients a_k belongs to H^{-s} when

$$\sum_{k\in\mathbb{Z}^m} (1 + |k|^2)^{-s}|a_k|^2 < \infty.$$

This is a separable Hilbert space, identified with a weighted ℓ^2 space. Take the countable product of these Hilbert spaces for $\chi_j g$, and impose the consistency equations $\chi_j(\chi_{j+1}g) = \chi_j g$. Multiplication by a smooth cutoff is continuous in these norms: its Fourier coefficients decay faster than every power, so the claim follows by convolution and Cauchy–Schwarz. Between different tori this operation first multiplies by χ_j, whose support is away from the smaller torus boundary, and then uses its periodic extension on that torus; the same estimate proves continuity. The consistency equations therefore define a closed subspace of a complete separable product. A consistent tuple defines a unique distribution, since any compactly supported test is contained in a region where some $\chi_j = 1$.

The Fourier coordinates of the localized fields have uniformly bounded second moments at each fixed j. For white noise this is the isometry applied to χ_j times the Fourier basis. For the pinned log-correlated field, the localized covariance kernel is bounded in absolute value by $C_j + |\log|x - y||$.

The pinning potentials are bounded on compact sets, as in Lemma 2.1. The integral of this bound in either variable over the fixed compact box is bounded. Cauchy–Schwarz with that integral bound makes the covariance quadratic form at most $C_j\|\varphi\|_2^2$, giving the same coordinate bound. Summing the expected weighted coefficient squares proves almost sure membership in H^{-s}, since $s > m/2$. A continuous function is locally in L^2, hence in this space as well.

Fourier coordinates after localization are smooth compact-test pairings. They generate the Hilbert and product Borel sigma fields and reconstruct the distribution by consistency. Every smooth test pairing is a continuous linear functional on the appropriate local H^{-s} space, because the test has finite H^s norm. This proves the countable determining property. For white noise on $\mathbb{R}^d \times (0,\infty)$, use compact rectangles bounded away from scale zero in the exhaustion. Its isonormal integrals are obtained from the same distribution by the L^2 isometry and approximation by smooth tests. □

Lemma A.2 (Measurability of compactly constrained distances)**.** *Let $K \subset \mathbb{R}^d$ be compact. For separating Euclidean lower-semicontinuous metrics D, the constrained length distance $D^{[K]}(x,y)$ is jointly lower semicontinuous in (D,x,y) for epigraph convergence of D. Its infimum over endpoints on two continuously parameterized compact curves in K is lower semicontinuous as well. For an open U, the analogous actual internal distance and compact-curve infimum are Borel functions of these coordinates, whenever the tested curves lie in a fixed compact subset of U.*

Proof. Suppose $D_n \to D$ epigraphically, $x_n \to x$, $y_n \to y$, and the constrained distances under consideration have finite lower limit M. Pass to a subsequence along which these distances converge to M, and choose the nth path in K with length at most $D_n^{[K]}(x_n,y_n) + 1/n$. After discarding a finite prefix, these lengths are at most $M+1$. Parameterize the paths on $[0,1]$ with metric speed at most $M+1$. For every Euclidean $a > 0$, compactness and the epigraph lower inequality imply

$$\liminf_n \inf\{D_n(u,v) : u,v \in K,\ |u-v| \geq a\} > 0.$$

Otherwise minimizing endpoint pairs have subsequential limits at Euclidean separation at least a and zero limiting D distance, contrary to separation. This gives a common Euclidean modulus for the paths along a further tail at each accuracy. The dense-time extraction of Lemma A.9 yields a uniformly convergent subsequence. The limit is D-Lipschitz by the lower inequality, joins x to y, remains in K, and has length no greater than the lower limit: apply the lower inequality to every finite partition and then take the supremum over partitions. This proves the first assertion, including infinite values. If endpoints are selected on two uniformly converging curves, first extract their two time parameters in the compact parameter intervals and apply the same argument. This proves the compact-curve assertion.

Choose increasing compact sets $K_j \subset U$ whose interiors exhaust U; for example use the sets with $|x| \leq j$ and $d(x,U^c) \geq 1/j$, with harmless enlargements of their indices. Every path with compact image in U lies in one K_j. Thus $D^U(x,y) = \inf_j D^{[K_j]}(x,y)$ for endpoints eventually in these sets. A countable infimum of Borel functions is Borel. The infimum over endpoint parameters commutes with $\inf_j$, giving the last assertion once the fixed tested curve images are contained in K_j for all sufficiently large j. □

Lemma A.3 (Regularity and tightness of a single law)**.** *A finite Borel measure on a complete separable metric space is inner regular by compact sets and outer regular by open sets.*

Proof. Replace the metric by its minimum with one. Let (x_i) be a countable dense set. For each k, the balls $B(x_i, 2^{-k})$ cover the space. Choose N_k so that the union of the first N_k balls has

complement of measure less than $\epsilon 2^{-k-1}$. The closed set

$$K = \bigcap_{k\geq 1} \overline{\bigcup_{i\leq N_k} B(x_i, 2^{-k})}$$

is totally bounded and complete, hence compact: a sequence has a Cauchy subsequence by successively choosing one of the finitely many balls at each scale, and completeness supplies its limit. The complement of K has measure at most ϵ.

For an open set O, the closed sets $\{x : d(x, O^c) \geq 1/n\}$ increase to O. For a closed set F, its $1/n$ neighborhoods decrease to F. Continuity of a finite measure under monotone limits therefore gives closed inner approximation of open sets and open outer approximation of closed sets. The collection of sets admitting both approximations is closed under complements and countable unions. For a union, first approximate finitely many terms from inside, and approximate every term from outside with summable errors. It consequently contains every Borel set. Intersect a closed inner approximation with the compact set constructed above to obtain a compact inner approximation. □

Lemma A.4 (A countable model of epigraph convergence)**.** *The nonnegative lower-semicontinuous functions on $\mathbb{R}^m$ with values in $[0,\infty]$ have a compact metrizable epigraph topology. Evaluation at a point is jointly lower semicontinuous in the function and point. This topology is exactly the lower-limit and recovery-sequence topology used in* (1.1)–(1.2)*.*

Proof. Use $\vartheta(t) = t/(1+t)$, with $\vartheta(\infty) = 1$, and associate to F the closed set

$$A_F = \{(x, u) : \vartheta(F(x)) \leq u \leq 1\} \subset \mathbb{R}^m \times [0, 1].$$

It contains $\mathbb{R}^m \times \{1\}$ and is upward closed in its last coordinate. Record the values of $d(\cdot, A_F)$ on a fixed countable dense set, using the product Euclidean metric. These values belong to $[0, 1]$. The distance functions are 1-Lipschitz, so coordinate convergence is equivalent to uniform convergence on every compact set. The coordinates determine A_F as the zero set of their continuous extension, hence determine F.

Every sequence of codes has a coordinatewise convergent subsequence. Let f be the limiting 1-Lipschitz function. Its zero set A is nonempty and contains the top layer. For every fixed z, choose $z_n \in A_{F_n}$ with $|z - z_n| \leq d(z, A_{F_n}) + 1/n \leq 1 + 1/n$. A subsequence of these points converges to z_*. Local uniform convergence gives $f(z_*) = 0$ and $|z - z_*| = f(z)$, so $d(z, A) \leq f(z)$. The Lipschitz inequality and $f|_A = 0$ give the reverse inequality. Thus $f = d(\cdot, A)$. If $(x, u) \in A$, approximating zero-distance points of A_{F_n} show that $(x, v) \in A$ for every $v \geq u$. Hence $A = A_F$ for one lower-semicontinuous F. This proves compactness of the code space, and the product metric supplies metrizability.

If $x_n \to x$ and a subsequence of $\vartheta(F_n(x_n))$ has limit u, then $(x, u) \in A_F$, so $F(x) \leq \liminf F_n(x_n)$. Conversely, approximating the point $(x, \vartheta(F(x)))$ by points of A_{F_n} gives recovery points when $F(x) < \infty$, using the lower inequality to obtain equality. If $F(x) = \infty$, the lower inequality already gives $F_n(x) \to \infty$. These are precisely the stated two conditions. They also give joint lower semicontinuity of evaluation. Conversely, if the lower and recovery conditions hold, every code-convergent subsequence has that same epigraph limit by these two conditions. Compactness of the code space then implies convergence of the original codes, proving the equivalence. □

Lemma A.5 (Sampling and conditional distributions)**.** *Let X and Y take values in complete separable metric spaces S, T. There is a Borel probability kernel $K(x, \cdot)$ on T such that*

$$\mathbb{E}[f(X)g(Y)] = \int_S f(x) \left(\int_T g(y) K(x, dy) \right) \mathcal{L}_X(dx)$$

for bounded Borel f, g. A probability kernel can be sampled by a measurable function of its base coordinate and one independent uniform random variable. If finitely many pairs (X_i, Y_i) are independent, their conditional law given all the X_i is the product of the individual conditional kernels. An independent complementary random coordinate can be included in the conditioning without changing this statement.

Proof. Choose refining countable Borel partitions $\mathcal{P}_m$ of T whose nonempty cells have diameter at most 2^{-m}. Such partitions are obtained by ordering a countable cover by small balls, subtracting the preceding balls, and refining the previous partition. Fix one representative in each cell. For a cell A, the Radon–Nikodym theorem gives a Borel function $p_A(x)$ satisfying

$$\mathbb{E}[\mathbf{1}_{\{X\in B\}}\mathbf{1}_{\{Y\in A\}}] = \int_B p_A(x)\mathcal{L}_X(dx).$$

There are only countably many cells. Outside one null set their probabilities are nonnegative, sum to one at each level, and each parent probability is the sum of its child probabilities. Assign any fixed consistent system on that null set.

Starting with an interval of length one, divide it into consecutive half-open intervals of lengths $p_A(x)$ for the first partition. Subdivide every such interval according to its child probabilities, continuing recursively. A uniform variable outside the countable set of subdivision endpoints selects a nested chain of nonempty cells. Its representatives form a Cauchy sequence, since two representatives at levels at least m belong to the same level-m cell. Their limit exists in T. The limit is a measurable function $\Phi(x, u)$: every finite representative is Borel in (x, u), and pointwise limits of measurable maps into a metric space are measurable. On the jointly Borel set of subdivision endpoints define Φ to be a fixed point of T; this set has zero conditional Lebesgue measure for each x and changes none of the asserted laws.

Let $K(x, \cdot)$ be the law of $\Phi(x, U)$. For bounded continuous g, the step functions which assign to a cell the value of g at its representative converge pointwise to g. Dominated convergence therefore gives the required identity from the cell identities. Continuous functions determine finite Borel measures: approximate indicators of closed sets by $\max\{0, 1 - nd(\cdot, F)\}$, then use Lemma A.3. The identity hence extends to bounded Borel g by the monotone-class argument. The same interval construction, using the kernel's cell probabilities, proves the sampling assertion for an arbitrary given kernel.

For independent pairs, test the proposed product conditional law against products $\prod_i f_i(X_i)g_i(Y_i)$. Independence and the individual identities give equality. Finite linear combinations and monotone limits extend it to all bounded Borel tests. This also proves the assertion about an independent complementary coordinate. □

Lemma A.6 (A coupled subsequence from tight coordinates)**.** *Suppose (μ_n) is a tight sequence of probability laws on a complete separable metric space S. There are a subsequence (n_k), a random variable Z, and jointly realized variables Z_k of laws μ_{n_k} such that $Z_k \to Z$ almost surely. Coordinatewise tightness suffices on a countable product of complete separable metric spaces.*

Proof. Use the refining small-diameter partitions from the preceding proof. They can be chosen so that, at each fixed level, every compact set is covered by finitely many cells: a compact set is covered by finitely many of the first balls at each of the finitely many refinement levels, and all its points are assigned to one of the corresponding finitely many cells. A diagonal subsequence makes $\mu_n(A)$ converge for every partition cell A. Tightness prevents loss of mass to infinitely many cells. Indeed, a compact set of uniformly large mass is covered by finitely many cells at the given level.

Thus the limits p_A sum to one at each level. The same finite-cover argument proves the parent–child consistency relation.

Sample a nested chain with these limiting cell probabilities by the interval construction, and let Z be the limit of its representatives. Denote its level-k cell label by J_k; the label is retained even if Z lies on that cell's boundary. For fixed k, the cell-probability vectors of μ_n converge in total variation to $(p_A)_{A\in\mathcal{P}_k}$. To verify this, choose finitely many cells carrying limiting mass above $1-\epsilon$, use coordinatewise convergence on them, and bound both tails by 2ϵ. Choose n_k so that this total variation distance is below 2^{-k}.

Two countable probability vectors p, q admit a coupling with disagreement probability $\frac{1}{2}\sum_i |p_i - q_i|$: put mass $\min(p_i, q_i)$ on each diagonal pair, and couple the two remaining probability vectors arbitrarily. Conditional on J_k, use this coupling to choose a label J'_k with the level-k cell law of μ_{n_k}. Then sample Z_k from μ_{n_k} conditioned on that cell. The construction in Lemma A.5 supplies all the required samples using independent uniform variables. When $J'_k = J_k$, both Z_k and Z belong to the closure of the same cell, so their distance is at most 2^{-k}. The disagreement probabilities are summable. Applying the elementary union bound to the tail union of these events proves that only finitely many occur. Hence $Z_k \to Z$ almost surely, with the correct marginal laws.

For a countable product, choose compact sets K_i with coordinate failure probabilities below $\epsilon 2^{-i-1}$, uniformly in n. Their product is compact in the metric $\sum_i 2^{-i}\min\{1, d_i\}$, by a diagonal subsequence argument, and has probability at least $1-\epsilon$. The product is complete and separable in this metric. The first part applies. □

Lemma A.7 (Measurable maps and graph limits)**.** *Let W have a fixed probability law on a complete separable metric space, and let F be Borel with values in another such space. For every $\epsilon > 0$ there is a compact set K with $\mathbb{P}(W \in K) > 1-\epsilon$ on which F is continuous. If (W, X_n) converges in distribution to $(W, F(W))$, on the same original probability space with the same W for every n, then $X_n \to F(W)$ in probability.*

Proof. For a map measurable in the completed sigma field, first choose a Borel version: replace its countable small-cell simple approximations by Borel versions, discard the union of the null exceptional sets, and take their pointwise limit. At each accuracy 2^{-m}, partition the target into countably many Borel sets of diameter at most 2^{-m}. Finitely many of their preimages cover all but $\epsilon 2^{-m-2}$ of the law of W. By Lemma A.3, choose compact subsets of these finitely many disjoint preimages with total additional loss below $\epsilon 2^{-m-2}$. Let K_m be their union and put $K = \bigcap_m K_m$. This is compact and has the asserted probability. On K, the function assigning to a point the representative of its level-m target cell is continuous: its finitely many level sets are disjoint compact sets. These functions converge uniformly to $F|_K$, proving continuity.

Fix a target error $\delta > 0$. The graph $G_K = \{(w, F(w)) : w \in K\}$ is compact. Uniform continuity of $F|_K$ gives $a > 0$ such that, if $w \in K$ and $d(X, F(w)) \geq \delta$, then the product distance from (w, X) to G_K is at least a. For example, choose $a < \delta/3$ so that $w, w' \in K$ and $d(w, w') < a$ imply $d(F(w), F(w')) < \delta/3$. The closed-set upper bound for distributional convergence now gives

$$\limsup_n \mathbb{P}[d(X_n, F(W)) \geq \delta] \leq \mathbb{P}(W \notin K) + \mathbb{P}[d((W, F(W)), G_K) \geq a] \leq 2\epsilon.$$

For completeness, that closed-set bound follows by testing the continuous functions $\max\{0, 1 - jd(\cdot, C)\}$ for a closed set C, and then decreasing them to its indicator. Let $\epsilon \downarrow 0$. □

Lemma A.8 (Analytic projection and measurable selection)**.** *Let S, T be complete separable metric spaces. The projection onto S of a Borel subset of $S \times T$ is measurable in the completion of every*

finite Borel measure on S. The same holds for its complement. More generally this holds for projections of closed sets with an additional complete separable witness coordinate. Every such relation with nonempty sections admits a selector measurable in these completions.

Proof. We supply the closed-set representation and the measurability argument. A set is called analytic here if it is the projection of a closed subset of its ambient space times a complete separable witness space. Countable unions of analytic sets are analytic: use the disjoint union of their witness spaces, keeping the chosen integer as a discrete first coordinate. Countable intersections are analytic: use the countable product of the witness spaces and impose all the closed conditions. These witness spaces can be given complete separable metrics. A closed set is analytic using a one-point witness; an open set is a countable union of closed sets, using positive distance from its complement. The class of sets for which both the set and its complement are analytic is therefore closed under complement and countable union and contains closed sets. It contains all Borel sets. A projection of a Borel relation is consequently analytic by keeping its second coordinate as part of the witness.

Let A be the projection of a closed set $C \subset S \times Z$. Cover Z by a countable tree of open sets U_s, indexed by finite integer strings, so that children cover their parent, $\overline{U_{si}} \subset U_s$, and diameters at depth m are at most 2^{-m}. Use small balls with centers in a countable dense set; empty nodes are discarded. Put

$$F_s = \overline{\mathrm{proj}_S(C \cap (S \times U_s))}.$$

These closed sets are nested along branches, and

$$A = \bigcup_{\alpha \in \mathbb{N}^{\mathbb{N}}} \bigcap_m F_{\alpha|m}.$$

For the nontrivial inclusion, if x belongs to every set along a branch, choose $(x_m, z_m) \in C$ with $x_m \to x$ and $z_m \in U_{\alpha|m}$. The z_m are Cauchy, hence converge to z, and closedness gives $(x, z) \in C$.

Fix a finite Borel measure μ on S, completed. For a node s, let A_s be the union of branch intersections continuing through s; then $A_s = \bigcup_i A_{si}$ and $A_s \subset F_s$. Every subset H has a measurable hull of measure $\mu^*(H)$: intersect measurable covers whose measures decrease to the outer measure. Choose a hull B_s of A_s, inside F_s and its parent's hull. The measurable set $B_s \setminus \bigcup_i B_{si}$ is disjoint from A_s and has measure zero; otherwise removing it would decrease the measure of a measurable hull. Let N be the union of these countably many null sets. Every $x \in B_\varnothing \setminus N$ belongs successively to a child hull at each level. The resulting branch places x in A. Thus $B_\varnothing \setminus N \subset A \subset B_\varnothing$, which proves completed measurability. Complements are measurable as well. This argument works for every finite Borel μ.

For selection, represent the relation as the projection of a closed $C \subset S \times Z$, where Z includes the desired T coordinate and any extra witnesses. For each x with a nonempty section, choose the first tree cell U_s meeting that section, then the first child still meeting it, and continue. Each viability test is an analytic projection and is measurable in every completion by what was just proved. The cell representatives give measurable functions of x and converge to a limit. For verification at one fixed x, choose any $z_m \in C_x \cap U_{s_m}$. Their distances from the corresponding representatives tend to zero. Closedness of C_x puts the common limit in C_x. These verifying witnesses need not be selected measurably. Project the measurable representative limit to T. This is the required selector. Extend it by a fixed value off the projection if a total map is desired. □

Lemma A.9 (Compact curve codes and lengths)**.** *Let K be a compact subset of Euclidean space, $T < \infty$, and let $\omega(a) \downarrow 0$ as $a \downarrow 0$. The curves $P : [0, T] \to K$ satisfying $|P(s) - P(t)| \le \omega(|s - t|)$ form a compact set for uniform convergence. If D is Euclidean lower semicontinuous, the length of*

such a curve is the supremum of its partition sums over rational partition times, with the endpoints included. In particular length and metric-Lipschitz conditions are countable Borel tests in the distance and curve coordinates. Length is lower semicontinuous under joint epigraph convergence of distances and uniform convergence of curves. For a positive continuous weight w, the weighted length of a rectifiable curve is also a Borel function of these coordinates and is lower semicontinuous under such convergence when w is fixed, or converges uniformly on the common compact range.

Proof. Extract convergent values on a countable dense set of times by successive subsequences. The common modulus makes this convergence uniform, and the limits extend to one continuous curve with the same modulus. This proves sequential compactness, which is compactness in a metric space.

Approximate an arbitrary finite increasing time partition by rational increasing partitions. Euclidean continuity of P and lower semicontinuity of D bound the original sum above by the lower limit of the rational sums. Taking suprema proves the length identity. The same approximation shows that metric Lipschitz inequalities at rational times imply them at all times. Each finite partition sum is lower semicontinuous under the stated joint convergence; its supremum is therefore lower semicontinuous. The countable formulas also prove the asserted measurability.

For the weighted assertion, for each finite rational partition $0 = t_0 < \cdots < t_k = T$ form

$$\sum_{i=1}^{k} \left(\min_{t \in [t_{i-1}, t_i]} w(P(t)) \right) D(P(t_{i-1}), P(t_i)).$$

Its supremum over partitions equals $\int w(P(t))\, d\ell_D(t)$. The inequality from left to right follows by comparison on each subinterval. For the reverse inequality first partition so that $w \circ P$ oscillates by less than a prescribed positive error on each piece, and refine inside each piece to approximate its unweighted length. Uniform continuity of the weight and the unweighted partition formula let the error decrease to zero. The same argument handles infinite length, since the weight has a positive minimum on the compact range. Each displayed finite sum is lower semicontinuous under joint convergence: its weight minima converge and its distance terms have the lower-limit inequality. Their countable supremum gives the claimed lower semicontinuity and measurability. □

B Topological foundations

This section proves the separation and extension statements used in the uniqueness argument. An arc means the image of an injective continuous map from a compact interval. No estimate for its Euclidean length, Lebesgue measure, or Hausdorff dimension is assumed. A finite embedded forest means a space homeomorphic to a finite graph with no cycles. The arguments below apply to arbitrary such embeddings. The separation statement also applies to finite unions of compact sets homeomorphic to closed subsets of intervals; an extension of such a set to an ambient embedded arc is unnecessary. The separation calculation is the compact-set form of Alexander duality; see Hatcher [12, Section 3.3] for its classical topological context. We prove the needed cochain and compact-limit statements here, including their use for arbitrary embedded interval subsets.

B.1 A closed partition of an interval

Lemma B.1. *A compact interval cannot be partitioned into countably many pairwise disjoint nonempty closed sets unless the partition has one member.*

Proof. Rescale the interval to $I = [0, 1]$. Suppose that (F_n) is such a partition and has at least two members. Write

$$U = \bigcup_n \operatorname{int}_I F_n, \qquad K = I \setminus U.$$

If K were empty, the disjoint relatively open sets $\operatorname{int}_I F_n$ would disconnect I. Thus K is a nonempty compact set.

We claim that $K \cap F_n$ has empty relative interior in K. Otherwise there is a relatively open interval $J \subset I$ such that $\varnothing \neq K \cap J \subset F_n$. Choose $x \in K \cap J$. Since $x \notin \operatorname{int}_I F_n$, choose $y \in J \setminus F_n$ sufficiently close to x that the interval between x and y lies in J. Necessarily $y \notin K$; let $m \neq n$ be the index for which $y \in \operatorname{int}_I F_m$. The component of $\operatorname{int}_I F_m$ containing y has an endpoint z on the side toward x, with z between x and y. Closedness gives $z \in F_m$. This endpoint is not in $\operatorname{int}_I F_m$, and it cannot be in the interior of a different partition member. Consequently $z \in K \cap J \subset F_n$, a contradiction.

For completeness, a nonempty compact metric space K is not a countable union of closed sets with empty relative interior. If (C_n) were such a cover, choose successively nonempty relatively open sets V_n with

$$\overline{V}_n \subset V_{n-1} \setminus C_n, \qquad \operatorname{diam} V_n \leq 2^{-n}, \qquad V_0 = K.$$

To choose V_n, first choose a point in the nonempty open set $V_{n-1} \setminus C_n$ and then a sufficiently small relative ball whose closure is contained there. The nested compact sets $\overline{V}_n$ have a common point, which belongs to no C_n. Apply this observation to $C_n = K \cap F_n$ to finish the proof. □

B.2 Scalar extension and finite trees

Lemma B.2. *Let A be a closed subset of a metric space X. Every continuous map $f : A \to [0, 1]$ has a continuous extension $F : X \to [0, 1]$.*

Proof. The empty-set case is immediate. By an affine change of variables it is enough to consider $f : A \to [-1, 1]$. We construct continuous $g_n : X \to \mathbb{R}$, starting with residual $r_0 = f$ and bound $M_0 = 1$. Given $|r_n| \leq M_n$, put

$$A_+ = \{a \in A : r_n(a) \geq M_n/3\}, \qquad A_- = \{a \in A : r_n(a) \leq -M_n/3\}.$$

These are disjoint closed subsets of X. If both are nonempty, set

$$g_n(x) = \frac{M_n}{3} \frac{\operatorname{dist}(x, A_-) - \operatorname{dist}(x, A_+)}{\operatorname{dist}(x, A_-) + \operatorname{dist}(x, A_+)}.$$

The denominator is positive at every point: distance zero to a closed set means membership in that set. If only A_+ is nonempty, use $g_n = M_n/3$; if only A_- is nonempty, use $g_n = -M_n/3$; if both are empty, use $g_n = 0$. In every case $|g_n| \leq M_n/3$ and

$$|r_n - g_n| \leq 2M_n/3 \quad \text{on } A.$$

For the displayed formula, check separately $r_n \geq M_n/3$, $r_n \leq -M_n/3$, and $|r_n| < M_n/3$. Define $r_{n+1} = r_n - g_n|_A$ and $M_{n+1} = 2M_n/3$. The series $G = \sum_{n \geq 0} g_n$ converges uniformly, since $\sum_n M_n/3 = 1$. It is continuous, takes values in $[-1, 1]$, and agrees with f on A, because $r_n \to 0$ uniformly. Rescaling gives the asserted extension. □

Lemma B.3. *Let T be a finite tree, with its usual graph topology. A continuous map from a closed subset of a metric space to T extends continuously to the whole space. In particular, a compact embedded finite tree $T \subset \mathbb{R}^d$ has a continuous retraction $\mathbb{R}^d \to T$.*

Proof. Give every abstract edge length one and use the resulting path metric d_T. Its topology is the graph topology. For an embedded tree the homeomorphism from the abstract compact graph identifies this topology with the subspace topology.

We use induction on the number of edges. A one-vertex tree is immediate. Remove a leaf edge with endpoints v, e, where e is the leaf, leaving a smaller tree T'. Let $c : T \to T'$ collapse this edge to v, and let $u : T \to [0,1]$ be zero on T' and equal to distance from v on the leaf edge. Both maps are continuous. For $f : A \to T$, extend $c \circ f$ to $F' : X \to T'$ by induction, and extend $u \circ f$ to $U : X \to [0,1]$ by Lemma B.2.

For $q \in T$ and $t \in [0,1]$, let $H(q,t)$ be the point a fraction t of the way from q to e along the unique simple path. This interpolation is continuous. Here is a metric check: the paths from e to two points share an initial segment, and comparison of their three possible positions relative to its final branch point gives

$$d_T(H(q,t), H(q',t)) \le d_T(q,q').$$

Also $d_T(H(q,t), H(q,t')) \le \operatorname{diam}(T)|t-t'|$. These two inequalities prove joint continuity. The map $F(x) = H(F'(x), U(x))$ agrees with f: on $f^{-1}(T')$ the second coordinate is zero, and on the leaf edge the first coordinate is v and the second is its edge coordinate. This proves the extension claim. Apply it to the identity on the closed set $T \subset \mathbb{R}^d$ to obtain the retraction. □

B.3 Finite cochains and compact-cover limits

We give the algebraic details needed for separation. All groups in this subsection have integer coefficients.

For a finite simplicial complex L, the group $C^q(L)$ consists of integer-valued functions on oriented q-simplices, with reversal of orientation changing the sign. Define

$$(\delta a)[v_0, \dots, v_{q+1}] = \sum_{i=0}^{q+1} (-1)^i a[v_0, \dots, \widehat{v_i}, \dots, v_{q+1}].$$

Every codimension-two face occurs twice with opposite signs, so $\delta^2 = 0$. Write $H^q(L) = \ker(\delta : C^q \to C^{q+1})/\delta C^{q-1}$. If $P \subset L$ is a subcomplex, $C^q(L,P)$ is the subgroup of cochains zero on P, and $H^q(L,P)$ is defined in the same way. In degree zero, H^0 consists precisely of the integer-valued functions constant on each connected component. A simplicial vertex map acts on chains by sending the ordered vertices of a simplex to their images, with value zero if an image vertex is repeated, and acts on cochains by composition with this chain map. Alternating face cancellation verifies that these maps commute with boundary and coboundary. This specifies all simplicial pullbacks used below.

We will use the following elementary cochain constructions.

Lemma B.4. *The following statements follow from the preceding definitions.*

(i) A short exact sequence of cochain complexes induces the usual long exact sequence of cohomology groups.

(ii) *For subcomplexes P, Q of a finite complex, there is an exact sequence*

$$\cdots \longrightarrow H^{q-1}(P) \oplus H^{q-1}(Q) \longrightarrow H^{q-1}(P \cap Q) \longrightarrow H^q(P \cup Q) \longrightarrow H^q(P) \oplus H^q(Q) \longrightarrow \cdots .$$

The first displayed arrow is the difference of restrictions.

(iii) *Barycentric subdivision does not change cohomology, including relative cohomology. A finite forest has zero cohomology in every positive degree. A triangulated d-sphere has $H^{d-1} = 0$ for $d \geq 3$ and $H^d \cong \mathbb{Z}$.*

Proof. For (i), if $0 \to A^* \xrightarrow{i} B^* \xrightarrow{p} C^* \to 0$ is exact and c is a cocycle in C^q, choose a lift $b \in B^q$. Since $p(\delta b) = 0$, there is a unique $a \in A^{q+1}$ with $i(a) = \delta b$. It is a cocycle because $i(\delta a) = 0$. Changing the lift adds a coboundary to a, and changing c by a coboundary has the same effect. This defines the connecting map. For example, its vanishing says $a = \delta a_0$, so $b - i(a_0)$ is a cocycle lifting c; the other exactness assertions are the same lift-and-subtract calculation.

For (ii), use the short exact sequence

$$0 \longrightarrow C^*(P \cup Q) \longrightarrow C^*(P) \oplus C^*(Q) \xrightarrow{(a,b) \mapsto a|_{P\cap Q} - b|_{P\cap Q}} C^*(P \cap Q) \longrightarrow 0.$$

The last map is onto because a cochain on the intersection can be extended by zero as a graded-group element. The kernel consists exactly of agreeing cochains, which paste to the union.

Here are the chain calculations behind (iii). Let $C_q(L)$ be the free group on oriented simplices, with boundary the alternating face sum. For a cone with vertex v, the operation $v * c$ satisfies

$$\partial(v * c) = c - v * (\partial c)$$

in positive degrees, with the corresponding augmented identity in degree zero. Thus every augmented cycle in a simplex or a cone is a boundary. Subdivision sends an oriented simplex to the sum of its oriented barycentric subsimplices; internal faces cancel, so this is a chain map s. Choose for each barycenter a vertex of the original face; nested faces ensure that these choices give a simplicial map r back to the original complex. The maps rs and the identity are chain homotopic: define the homotopy successively on the dimensions of a simplex. Once it has been defined on its boundary, the required difference on the simplex is a cycle supported in that simplex, and the cone identity fills it. The same construction for sr takes place in the subdivided carrier simplex, which is a cone from its barycenter over its subdivided boundary. This gives the other chain homotopy. All carriers of simplices of a designated subcomplex remain in that subcomplex, so the construction also gives the relative assertion. Dualizing these identities proves the cohomology isomorphisms. The same carrier construction shows that two simplicial maps whose images of each simplex together lie in one simplex induce the same map on cohomology.

For a finite tree, every 1-cochain is the coboundary of a vertex function: choose a root, give it value zero, and sum the signed edge values along the unique path to each vertex. There are no simplices of dimension at least two. This proves the forest assertion.

For the boundary of a $(d+1)$-simplex, a cocycle of degree $1 \leq q < d$ is also a cocycle on the whole simplex, since all its $(q+1)$-faces lie on the boundary. The cone identity gives a primitive. In degree d, evaluation on the coherently oriented sum of boundary faces gives an onto map to $\mathbb{Z}$; its kernel consists of cochains that extend to cocycles on the whole simplex, and these again have primitives. Thus its degree-d cohomology is $\mathbb{Z}$. These conclusions hold for subdivisions. Equivalently, for any finite triangulation of a space homeomorphic to this sphere they follow from the compact-cover identification proved next, which depends on subdivision alone and is invariant under homeomorphisms. □

If $\mathcal{U} = (U_i)$ is a finite open cover of a compact metric space X, its *nerve* $N(\mathcal{U})$ has a vertex for each member, and a simplex for every subfamily with nonempty intersection. A refinement induces a simplicial map between nerves by assigning each new member to an old member containing it. Two choices are contiguous: for any simplex, all the chosen old members contain a common point. The last carrier calculation in the preceding proof shows that they induce the same cohomology map. Define

$$\check{H}^q(X) = \varinjlim_{\mathcal{U}} H^q(N(\mathcal{U})), \tag{B.1}$$

where the limit is over refinement. Concretely, an element is a cocycle on a finite-cover nerve; two such elements are equal if their pullbacks become cohomologous on a common refinement. A continuous map induces pullbacks by inverse images of covers. In particular these groups are unchanged by a homeomorphism.

Lemma B.5. *The compact-cover groups have the following properties.*

(i) For a finite polyhedron L, $\check{H}^q(|L|) = H^q(L)$.

(ii) If A is compact in a finite polyhedron and (P_j) is a decreasing, cofinal sequence of compact polyhedral neighborhoods of A, then

$$\check{H}^q(A) = \varinjlim_{j} H^q(P_j),$$

with restriction maps as j increases.

(iii) For compact sets A, B in a finite polyhedron there is an exact sequence

$$\begin{aligned} \cdots \longrightarrow \check{H}^{q-1}(A) \oplus \check{H}^{q-1}(B) &\longrightarrow \check{H}^{q-1}(A \cap B) \\ \longrightarrow \check{H}^q(A \cup B) &\longrightarrow \check{H}^q(A) \oplus \check{H}^q(B) \longrightarrow \cdots . \end{aligned}$$

Proof. For (i), the open stars of the vertices of a triangulation cover its polyhedron and have nerve equal to that triangulation: a collection of these stars meets exactly when their vertices span a simplex. The meshes of iterated barycentric subdivisions tend to zero. Indeed a simplex of dimension at most m has its subdivision mesh at most $m/(m+1)$ times the original mesh; this follows by writing the barycenters of two nested faces as convex combinations of their vertices. Star diameters are at most twice the mesh. By the Lebesgue-number property of a finite open cover of a compact metric space, these star covers are cofinal among all covers. A refinement map between sufficiently fine star covers is carried by the old simplices and hence induces the subdivision isomorphism of Lemma B.4. This proves (i). The Lebesgue-number property itself follows by choosing at each point a ball whose doubled ball lies in a cover member, taking finitely many of the smaller balls, and taking the minimum of their radii.

We detail the finite-cover continuity used in (ii). A finite open cover of A has a finite closed refinement (C_i) whose relative interiors cover A: take sufficiently small relative closed balls inside the prescribed members and then a finite subcover by their interiors. For a finite family of compact sets (C_i), small open ambient neighborhoods can be chosen with exactly the same nerve. To see this, each subfamily with empty intersection has neighborhoods with empty intersection. Otherwise points in intersections of $1/n$-neighborhoods have a convergent subsequence, whose limit belongs to every set in that subfamily. There are only finitely many subfamilies, so one neighborhood radius works for all. Shrinking further makes each neighborhood lie in any prescribed ambient open set containing its C_i.

These observations also justify using closed covers with interiors covering A in the definition: they refine arbitrary open covers, and a sufficiently small open enlargement has the identical nerve. They prove the following two statements about finite nerve data. Every cover, refinement map, and cochain on A can, after a finite refinement, be represented on a neighborhood of A. Every equality between such cochains, including a coboundary identity, that holds after a refinement on A holds on a sufficiently small neighborhood after an extending refinement. To verify the latter assertion, use a closed refinement subordinate to all the finitely many covers and refinement choices in the equality; choose its open enlargements with the same nerve and inside their designated old members. The integer cochains and all their face equations are then literally unchanged. The enlarged sets cover an open neighborhood of A, which contains P_j for all sufficiently large j. The first statement gives surjectivity of the map in (ii), and the second gives injectivity. Part (i) identifies the groups on P_j with finite simplicial cohomology.

For clarity, polyhedral neighborhoods with these properties exist by subdivision. In a sufficiently fine subdivision take all closed simplices meeting the compact set and their closed stars. Their union contains a neighborhood of the set and is contained in any prescribed neighborhood once the mesh is small enough. Successive meshes can therefore be chosen so that the neighborhoods decrease inside one another's interiors. Finitely many such choices can be made in one common subdivision.

For (iii), choose decreasing polyhedral neighborhoods A_j, B_j in common subdivisions. Their unions are cofinal neighborhoods of $A \cup B$. Their intersections are cofinal neighborhoods of $A \cap B$: if not, a sequence of points in $A_j \cap B_j$ outside a fixed neighborhood of $A \cap B$ has a limit in $A \cap B$, a contradiction. If the intersection is empty, the neighborhoods are eventually disjoint. Apply the finite exact sequence of Lemma B.4 and take limits. Exactness is preserved here for an elementary reason. An element is zero in a direct limit precisely when it is zero at a later stage; move to that stage and use exactness there to lift it. This proves (iii) and the claimed sequence without any separation theorem. □

B.4 Separation computed by finite cochains

For any space Y, let $H_0(Y)$ denote the free abelian group on its path components, and let $\widetilde{H}_0(Y)$ be the kernel of the map which sums the coefficients. For a nonempty open subset of a Euclidean space or sphere, $\widetilde{H}_0(Y) = 0$ exactly when Y is connected. Indeed the points reachable from a given point by polygonal paths in local coordinate balls form an open subset with open complement, so a connected open set is path connected.

Lemma B.6. *For a compact set $A \subset S^d$, $d \geq 3$, there is a natural isomorphism*

$$\check{H}^{d-1}(A) \cong \widetilde{H}_0(S^d \setminus A).$$

Consequently, if the complement is nonempty, it is connected if and only if $\check{H}^{d-1}(A) = 0$.

Proof. Use the boundary of a $(d+1)$-simplex as a fixed model of the sphere; radial projection from its interior to the round sphere is a homeomorphism. All neighborhoods below use subdivisions of this one fixed triangulation. First let $A = P$ be a subcomplex of one such subdivision L. Orient all d-simplices coherently, using the boundary orientation and its subdivisions, so that their induced orientations on a shared facet are opposite. Form a graph G_P whose vertices are the d-simplices not in P and whose edges correspond to the $(d-1)$-faces not in P, joining their two adjacent d-simplices. Both adjacent simplices are outside P, since P is a subcomplex.

The graph components are precisely the components of $S^d \setminus P$. An arbitrary point of the complement lies in the relative interior of a simplex not in P and can be joined to the barycenter

of any containing top simplex without meeting P. An edge of G_P is realized by the two straight segments from its top-simplex barycenters to its common facet barycenter. Conversely, a path in the complement has positive distance from P. In finitely many coordinate patches it can be approximated by a piecewise linear path still in the complement and in general position with respect to the finite triangulation. The coordinate patches here need no manifold triangulation theorem: project a small patch of the boundary of the convex simplex from an interior point to a transverse affine hyperplane. This is a coordinate map on a sufficiently small patch and sends each face piece to a polyhedral piece of the same dimension. Here the general-position step is elementary and finite: move each of the finitely many vertices inside a small open ball. The choices for which a segment meets a fixed face of codimension at least two lie in a finite union of sets of empty interior, as is seen by taking the cone from its fixed preceding vertex over that face; this cone has affine dimension at most $d-1$. Avoid also the finitely many facet hyperplanes for vertices and nongeneric segments. A choice outside these sets exists in every open ball, by successively avoiding finitely many proper affine subspaces or their finite polyhedral pieces. Using a subdivision into the coordinate patches gives the same conclusion on the sphere. The resulting path crosses only interiors of facets, one at a time, so records a route in G_P. Endpoints can first be moved to top-simplex barycenters by the preceding segments. This proves the component assertion.

The group $C^d(L,P)$ has a basis e_σ indexed by the vertices of G_P. The coboundary of the basis cochain of a facet outside P is, up to a sign, $e_\sigma - e_{\sigma'}$ for its two adjacent top simplices. Thus

$$H^d(L,P) = C^d(L,P)/\delta C^{d-1}(L,P) \cong H_0(G_P) \cong H_0(S^d \setminus P).$$

In this identification the map $H^d(L,P) \to H^d(L) = \mathbb{Z}$ sums coefficients, because every coherently oriented top simplex evaluates as one on the fundamental cycle. The relative exact sequence contains

$$0 = H^{d-1}(L) \longrightarrow H^{d-1}(P) \longrightarrow H^d(L,P) \longrightarrow H^d(L).$$

Exactness proves the required isomorphism for P.

The isomorphism is natural under decreasing subcomplexes. Inclusion of the relative cochains sends a top-simplex basis vector to the same basis vector, and on graph components it is the inclusion of complements. It is also compatible with subdivision. Under the subdivision chain map, a cochain equal to one on a single oriented child top simplex and zero on the others evaluates as one on the parent; all children in that parent's interior belong to the same complement component. This checks compatibility on the generators of the displayed top-cochain quotient, and hence on its kernel.

For general compact A, choose decreasing cofinal polyhedral neighborhoods P_j. Their complements increase to $S^d \setminus A$. The groups H_0, and also their reduced groups, commute with this union. Every finite collection of points lies in one sufficiently late complement, and a path giving a relation between two of them has compact image disjoint from A, hence also lies in one such complement. This proves both surjectivity and injectivity of the component-group limit map. The finite isomorphisms are compatible by the preceding paragraph. Taking their direct limit and using Lemma B.5(ii) proves the result. □

The one-point compactification of $\mathbb{R}^d$ may be used as S^d in this lemma. Stereographic projection followed by the preceding radial homeomorphism identifies it with the fixed sphere model. An arbitrary compact set is transported by this homeomorphism; there is no requirement that it become a polyhedron. Its polyhedral neighborhoods are then obtained from subdivisions of the fixed triangulation as in Lemma B.5.

B.5 Finite unions of embedded closed interval subsets

Lemma B.7. *Let $A = A_1 \cup \cdots \cup A_m$, where every A_i is compact and homeomorphic to a closed subset of a compact real interval, and let C be a closed subset of A.*

(i) $\check{H}^q(C) = 0$ for every $q \geq 2$.

(ii) The restriction $\check{H}^1(A) \to \check{H}^1(C)$ is onto.

Proof. A closed subset of an interval has zero cohomology in every positive degree. Indeed it has cofinal neighborhoods which are finite unions of disjoint compact intervals: cover it by finitely many short intervals inside a prescribed neighborhood and merge overlaps. The interval calculation in Lemma B.4 and compact-cover continuity prove the assertion. Homeomorphism invariance gives the same assertion for a closed subset of an embedded arc, or of any compact space homeomorphic to a closed subset of an interval.

For (i), induct on m and apply compact exactness to

$$C = (C \cap (A_1 \cup \cdots \cup A_{m-1})) \cup (C \cap A_m).$$

The new piece and its intersection with the old piece are closed subsets of A_m, hence homeomorphic to closed interval subsets, so their positive-degree cohomology vanishes. For $q \geq 2$ the three groups surrounding $\check{H}^q(C)$ in the exact sequence vanish, proving (i).

We supply the circle-valued description of degree one needed for (ii). A continuous map $f : X \to S^1$ from a compact metric space has an associated class in $\check{H}^1(X)$. Cover the circle by short open arcs with real argument functions, chosen so that each nonempty pairwise intersection is connected. Four equally spaced arcs of angular length strictly between $\pi/2$ and π suffice. The inverse-image cover has local real lifts u_i with $f = \exp(2\pi i u_i)$. On overlaps the differences $u_i - u_j$ are constant integers, giving a nerve cocycle. Refining the cover or changing a lift changes this cocycle by a coboundary. The same conclusion holds for any finite cover with continuous local lifts: refine it into small neighborhoods where differences from the chosen short-arc lifts are constant integers. Thus the class is well defined and is natural under restrictions and continuous pullback.

Every class of $\check{H}^1(X)$ arises in this way. Represent it by an integer cocycle n_{ij} on the nerve of a finite cover (U_i), with $n_{ji} = -n_{ij}$, $n_{ii} = 0$, and $n_{ij} + n_{jk} = n_{ik}$ on triple intersections. Choose a continuous partition of unity (ϕ_j) subordinate in the sense $\phi_j(x) > 0 \Rightarrow x \in U_j$. One explicit choice is

$$\phi_j(x) = \frac{\rho_j(x)}{\sum_k \rho_k(x)}, \qquad \rho_j(x) = \min\{1, \operatorname{dist}(x, X \setminus U_j)\},$$

where $\rho_j = 1$ if $U_j = X$. On U_i set $u_i(x) = \sum_j \phi_j(x) n_{ij}$, assigning arbitrary values to n_{ij} when $U_i \cap U_j$ is empty, since that term is zero on U_i. The triple-intersection identity gives $u_i - u_k = n_{ik}$ on $U_i \cap U_k$. Consequently $\exp(2\pi i u_i)$ paste to a continuous map $f : X \to S^1$, and the preceding local-lift calculation shows that its class is the prescribed cocycle.

It remains to extend a continuous $f : C \to S^1$ to A. A continuous map from a closed subset J of $[0, 1]$ to S^1 extends to the interval: on a complementary interval with both endpoints in J join their images by a shorter circular arc at constant speed; if they are antipodal choose either of the two semicircles. On a component with only one endpoint in J use its constant value. If J is empty choose a constant map. This extension is continuous at J. To check this, complementary intervals of length tending to zero and approaching a point of J have both endpoint values approaching the value there, and their shorter arcs remain in an arbitrarily small circular neighborhood of that value. There are only finitely many complementary intervals of length at least any fixed positive

number; on these, ordinary endpoint continuity applies. This proves the continuity assertion for every sequence.

Identify A_1 with a closed set $J_1 \subset [0,1]$. The prescribed set $C \cap A_1$ becomes a closed subset of $[0,1]$; extend to that whole interval by the construction just given, then restrict to J_1 and transport back to A_1. Paste this extension to the original map on C. Next extend across A_2, where the prescribed set is $A_2 \cap (C \cup A_1)$, and continue using a closed interval model for each A_i. Every prescribed set is closed, and the finite closed-set pasting property makes each step continuous. We obtain a map $A \to S^1$ extending f. Represent an arbitrary class of $\check{H}^1(C)$ by such an f and take the class of its extension. Naturality gives the desired preimage, proving (ii). □

B.6 Nonseparation by embedded interval subsets

Theorem B.8. *Let $d \geq 3$, let $O \subset \mathbb{R}^d$ be connected and open, and let A be a finite union of compact sets each homeomorphic to a closed subset of a compact real interval. Then $O \setminus A$ is connected and open. In particular this holds for a finite union of compact embedded arcs, whether or not the arcs are contained in O. One may replace $A \cap O$ by any relatively closed subset of itself; the resulting complement is again connected and open.*

Proof. Regard $\mathbb{R}^d$ as S^d minus the point at infinity and put $B = S^d \setminus O$. It is compact. By Lemma B.6,

$$\check{H}^{d-1}(B) = 0,$$

since its complement O is connected. Put $C = A \cap B$. The compact exact sequence includes

$$\check{H}^{d-2}(A) \oplus \check{H}^{d-2}(B) \longrightarrow \check{H}^{d-2}(C) \longrightarrow \check{H}^{d-1}(A \cup B) \longrightarrow \check{H}^{d-1}(A) \oplus \check{H}^{d-1}(B).$$

The last group is zero by Lemma B.7(i). The first arrow is onto: when $d = 3$ use part (ii) of that lemma on its A summand; when $d \geq 4$ the target is zero by part (i). Exactness gives $\check{H}^{d-1}(A \cup B) = 0$, and separation gives connectedness of $S^d \setminus (A \cup B) = O \setminus A$.

We finally verify density and the subset assertion. A compact set homeomorphic to a closed subset of a real interval has empty interior in $\mathbb{R}^d$ for $d \geq 3$. Otherwise an open ball contained in it has, under that homeomorphism, a connected subset of the real line and hence an interval. Remove a point whose parameter is interior to that interval. The interval becomes disconnected, whereas a ball minus one point in dimension at least two is polygonally connected, a contradiction. Finitely many closed sets with empty interior have union with empty interior: inside any given open ball choose a smaller open ball avoiding the first closed set, then one avoiding the second, and so on. Thus the full obstacle A has dense complement in O, and this complement is nonempty.

If K is a relatively closed subset of that obstacle in O, then

$$O \setminus (\text{full obstacle}) \;\subset\; O \setminus K \;\subset\; \overline{O \setminus (\text{full obstacle})}^{\,O}.$$

A set between a connected set and its closure is connected. To check this last elementary assertion, a separation of the larger set would place the connected subset in one part; a point in the other relatively open part would then have a neighborhood missing that dense subset. Relative closedness of K makes its complement open. This completes the proof. □

References

[1] R. H. Cameron and W. T. Martin. Transformations of Wiener integrals under translations. *Annals of Mathematics (2)* **45** (1944), no. 2, 386–396. doi:10.2307/1969276.

[2] W.-K. Chen, N. Dafnis, and G. Paouris. Improved Hölder and reverse Hölder inequalities for Gaussian random vectors. *Advances in Mathematics* **280** (2015), 643–689. doi:10.1016/j.aim.2014.09.029.

[3] A. A. Contreras Hip and Z. Zhuang. Weak exponential metrics for high-dimensional log-correlated Gaussian fields. Preprint, 2025. arXiv:2512.06292.

[4] J. Ding and E. Gwynne. Tightness of supercritical Liouville first passage percolation. *Journal of the European Mathematical Society* **25** (2023), no. 10, 3833–3911. doi:10.4171/JEMS/1273.

[5] J. Ding and E. Gwynne. Uniqueness of the critical and supercritical Liouville quantum gravity metrics. *Proceedings of the London Mathematical Society (3)* **126** (2023), 216–333. doi:10.1112/plms.12492.

[6] J. Ding, E. Gwynne, and Z. Zhuang. Tightness of exponential metrics for log-correlated Gaussian fields in arbitrary dimension. Preprint, 2023, revised 2026. arXiv:2310.03996.

[7] J. Ding, E. Gwynne, and Z. Zhuang. Percolation of thick points of the log-correlated Gaussian field in high dimensions. Preprint, 2024, revised 2026. arXiv:2410.22075.

[8] B. Efron and C. Stein. The jackknife estimate of variance. *The Annals of Statistics* **9** (1981), no. 3, 586–596. doi:10.1214/aos/1176345462.

[9] L. Gross. Logarithmic Sobolev inequalities. *American Journal of Mathematics* **97** (1975), no. 4, 1061–1083. doi:10.2307/2373688.

[10] E. Gwynne and J. Miller. Local metrics of the Gaussian free field. *Annales de l'Institut Fourier* **70** (2020), no. 5, 2049–2075. doi:10.5802/aif.3398.

[11] E. Gwynne and J. 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.

[12] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. Author's electronic edition.

[13] O. Kallenberg. *Foundations of Modern Probability*. Third edition. Probability Theory and Stochastic Modelling, vol. 99. Springer, 2021. doi:10.1007/978-3-030-61871-1.

[14] A. S. Kechris. *Classical Descriptive Set Theory*. Graduate Texts in Mathematics, vol. 156. Springer, 1995. doi:10.1007/978-1-4612-4190-4.