Almost sure convergence of LFPP in the positive-Q regime

An upgrade of the higher-dimensional construction

13 November 2026

1 Statement and inputs

Fix $d \geq 3$, $\xi > 0$, and the parameter $Q = Q(\xi) > 0$ of the preceding construction. Let h be the whole-space log-correlated Gaussian field, with its additive constant fixed by $h_1(0) = 0$. Use the same nonnegative, smooth, radial, compactly supported probability mollifier ϱ as in that construction, and set

$$h_\varepsilon = h * \varrho_\varepsilon, \qquad D_h^\varepsilon(x, y; U) = \inf_{P: x \to y,\, P \subset U} \int_P e^{\xi h_\varepsilon(z)} \, |dz|, \qquad a_\varepsilon = \operatorname{Med} D_h^\varepsilon(0, e_1; B_2).$$

Here $B_r(z) = z + (-r, r)^d$, and the median is the lower median. Write $X_\varepsilon = a_\varepsilon^{-1} D_h^\varepsilon$. The field, its approximations, and the limiting metric $D = D_h$ are coupled on their original probability space.

Theorem. Under the conclusions and quantitative estimates of the preceding positive-Q construction and uniqueness proofs, almost surely

$$X_\varepsilon \longrightarrow D \qquad \text{as } \varepsilon \downarrow 0$$

in the topology on lower semicontinuous functions on $\mathbb{R}^d \times \mathbb{R}^d$. The limit is through all positive real cutoffs, with the original deterministic normalizers a_ε.

Explicitly, on one event of probability one, the following hold.

1. For every $\varepsilon_n \to 0$ and $(x_n, y_n) \to (x, y)$,
$$D(x, y) \leq \liminf_n X_{\varepsilon_n}(x_n, y_n).$$
2. For every pair (x, y) and every such cutoff sequence, there are $(x_n, y_n) \to (x, y)$ for which $X_{\varepsilon_n}(x_n, y_n) \to D(x, y)$.

The theorem is for each fixed parameter and dimension. Infinite distances are included. In particular, it applies to the supercritical part of the positive-Q regime. Its hypotheses below are already quantitative inputs to the earlier construction; this note does not independently reprove that construction or its uniqueness theorem.

1.1 The precise imported estimates

We use [T], [C], and [U] in the following forms.

Metric identification. The metrics X_ε converge in probability to the same measurable assignment D_h in the lower-semicontinuous topology, jointly with their evaluations at rational point pairs. The analogous convergence holds after a bounded smooth field perturbation. The metric is a separating extended length metric, finite-distance pairs admit Euclidean-continuous geodesics, and actual internal metrics are local. Continuous Weyl scaling and deterministic translation/dilation covariance hold. Fixed rational points have finite internal distance in every connected open buffer containing them. These are the conclusions of [U, Theorem 1.1 and Sections 2–6, 29].

Scale estimates and actual chains. Put

$$c_\varepsilon(r) = \frac{r a_{\varepsilon/r}}{a_\varepsilon}, \qquad c(r) = r^{\xi Q}, \qquad \kappa_\varepsilon(r) = \frac{c_\varepsilon(r)}{c(r)}.$$

For every $\eta > 0$, the raw relative-power bounds hold uniformly for $\varepsilon \le s \le r \le 1$:

$$C_\eta^{-1}(s/r)^{\xi Q+\eta} \le \frac{c_\varepsilon(s)}{c_\varepsilon(r)} \le C_\eta(s/r)^{\xi Q-\eta}.$$

Normalized annular crossings have uniformly bounded inverse moments of every fixed order. The construction supplies actual upper nice-site chains, their vertices and paths, and simultaneous lower shell charges, including the terminal straight-segment estimate for LFPP. The chain at relative mesh depth ℓ stays within a distance $C(1+\ell)^k r 2^{-\ell}$ of the point to which it is attached. The upper and lower labels use the same coarse field. These are the estimates of [C, Sections 4–5], applied after the scalar normalizer comparison in [T]. The compact path extraction in [U, Section 2] transfers the chains to the limiting metric by the finite-depth retention argument supplied in Appendix A; this is also the chain structure used in [U, Lemma 9.1]. Thus we retain the actual chains, rather than just their final connection inequality.

Connections and cores. In any fixed annular band with a larger annular buffer, for every $\delta > 0$ and $p < 1$, an event of probability at least p gives a common connected finite core to which every finite crossing arm attaches, in its own metric, at cost at most δ times the arm length. The core has tight total length in the appropriate root units. Its vertex menu can be chosen deterministic and finite after fixing the depth. Finitely many prescribed rational points in the buffer may be added to the menu. In particular, for a suitable fixed C, every two crossing arms P_1, P_2 can be connected inside the buffer at cost

$$C(\mathrm{len}(P_1) + \mathrm{len}(P_2)).$$

The assertions are simultaneous over arms chosen after the field, both for the approximations and for D. See [C, Proposition on local relative connectors], [T, Section 18], and [U, Sections 3–4]. The transfer in [U] uses recovery of arbitrary finite metric paths and retained successful-annulus events, not an uncountable intersection of probability-one events.

Annular concentration. Local events invariant under addition of a field constant satisfy the following implication at sufficiently separated concentric scales. Given $a > 0$ and $b \in (0,1)$, their marginal probabilities can be required sufficiently close to one that the probability of fewer than bn successes among n scales is at most Ce^{-an}. The determining annuli are fixed before the marginal probability and core depth are chosen. We use [CHZ, Lemma 2.7], applied as in [U, Section 14].

All universal path statements below are deterministic consequences of these simultaneous local estimates. For annular concentration, good events are formulated intrinsically using the two actual internal metrics in the fixed buffer and, for normalized lengths, the local root average. They are invariant under field constants and measurable in the completed local field sigma algebra, as in

[U]. Auxiliary coarse-field or white-noise chain labels are used to prove their marginal probability estimates; those labels need not themselves be annulus-local. Actual paths staying in the buffer alone would not justify locality of the auxiliary labels.

1.2 Two consequences used throughout

First, for every fixed exponent band $\varepsilon^t \le r \le \varepsilon^s$, $0 < s < t < 1$, there is a deterministic $A_* < \infty$ such that, outside an event of probability $O(\varepsilon^{\beta t})$, every path in either metric which makes a crossing of radius comparable to ε^t in a fixed compact region has length at least

$$\varepsilon^{A_* t}.$$

The exponent A_* can be fixed before the band and before any comparison accuracy. To see this, use a grid with $O(\varepsilon^{-dt})$ centers. The inverse crossing moments bound all normalized crossings below by a small fixed power of the grid scale. The Gaussian coarse field is simultaneously bounded below by $-A\log(1/r)$, with a polynomial error, after choosing $A > \sqrt{2d}$ sufficiently large. The relative-power estimate bounds $c_\varepsilon(r)$ below by $C_\eta^{-1} r^{\xi Q+\eta}$. A fixed-factor grid annulus is crossed whenever two endpoints are separated by the required scale. These observations give the displayed bound in both metrics. The same bound therefore applies when the available root radius is larger than ε^t.

Second,

$$D(\overline{B}_k, \partial B_R) \longrightarrow \infty \quad (R \to \infty)$$

almost surely for every integer k. Indeed, on dyadic large annuli the crossing has the law of $R^{\xi Q} e^{\xi h_R(0)}$ times a positive normalized variable with all inverse moments. Gaussian tails and Borel–Cantelli give $h_R(0) \ge -(Q/4)\log R$ eventually, while inverse moments make the normalized variable at least $R^{-\xi Q/4}$ eventually. Thus these crossings tend to infinity. Every path from $\overline{B}_k$ to a sufficiently distant boundary crosses one of them. This argument uses $Q > 0$.

2 An initial comparison

We first need an estimate with a fixed constant. Convergence in probability alone supplies no summable rate.

Lemma (initial comparison). There is a deterministic $C_0 < \infty$ such that, for every fixed $R, M < \infty$, some $s_0, \beta_0 > 0$ give the following outside an event of probability $O_{R,M}(\varepsilon^{\beta_0})$:

$$\mathsf{e}\big(B_{C_*\varepsilon^{s_0}}(x), B_{C_*\varepsilon^{s_0}}(y); \overline{B}_{R+1}\big) \le C_0 \operatorname{len}_{\mathsf{d}}(P),$$

simultaneously for $(\mathsf{d}, \mathsf{e}) = (D, X_\varepsilon)$ and (X_ε, D), and every source path $P \subset \overline{B}_R$ of source length at most M from x to y. The constants C_0, s_0 may be fixed independently of R, M.

Proof. Appendix A proves first a ratio-weighted comparison on any deterministic positive fraction of the radii in a fixed exponent band. Its constant is

$$C \max_{r \in S_\varepsilon} \kappa_\varepsilon(r) \quad \text{or} \quad C \max_{r \in S_\varepsilon} \kappa_\varepsilon(r)^{-1}.$$

The proof averages target upper-chain costs against source lower shell charges. At the LFPP cutoff it retains a nearby rational vertex when the target is D. It therefore does not assert that a point selected on an arbitrary LFPP path has finite D-distance.

A maximal-prefix argument connects these local paths. An outer seed scale, separated from all subsequent routing scales, places the target starting point strictly inside a ball that the source path has visited for the last time. Thus every later target prefix crosses all routing annuli. Two first-exit balances, obtained from the preceding polynomial lower bound, make the prefix extension close.

Finally, if a positive fraction of radii had $\kappa_\varepsilon < T^{-1}$, or $\kappa_\varepsilon > T$, the corresponding estimate would force a fixed nondegenerate distance to be arbitrarily small when T is large. Uniform point upper bounds and annular lower bounds rule this out. A positive fraction of radii consequently has $T^{-1} \le \kappa_\varepsilon \le T$. The ratio-weighted estimate on this set gives the lemma. Appendix A gives the complete argument, including its parameter order. □

The proof applies to all bounded-length paths; in particular it is available before global LFPP geodesics or confinement have been established.

3 Sharp comparison on one annulus

For a nonnegative lower semicontinuous function f on a Euclidean space, write

$$f^{[\eta]}(p) = \inf_{|q-p|<\eta} f(q).$$

For pairs of points use the maximum product norm.

Lemma (compact-height comparison). Suppose $f_n \to f$ in the lower-semicontinuous topology. For a compact K, finite M, and $\eta, \nu > 0$, eventually

$$\begin{aligned} f_n^{[\eta]}(p) &\le f(p) + \nu \quad (p \in K,\ f(p) \le M), \\ f^{[\eta]}(p) &\le f_n(p) + \nu \quad (p \in K,\ f_n(p) \le M). \end{aligned}$$

If the convergence is in probability, the simultaneous inequalities hold with probability tending to one.

Proof. Cover the compact truncated epigraph $\{(p,t) : p \in K,\ f(p) \le t \le M\}$ by finitely many product neighborhoods of spatial radius less than $\eta/3$ and height radius less than $\nu/3$. Choose a recovery sequence at each center. The finite cover proves the first inequality. If the second failed, a subsequence would give $p_n \to p$, $f_n(p_n) \to t \le M$, and $f^{[\eta]}(p_n) > f_n(p_n) + \nu$. The lower-limit inequality gives $f(p) \le t$, contradicting the last inequality for large n. The probability statement follows by taking an almost surely convergent subsequence from every subsequence on which the failure probability is bounded away from zero. □

Lemma (local comparison). Let $V = \{1/2 \le |z|_\infty \le 3\}$ and let U be a fixed larger open annulus, with inner radius smaller than $1/(100\sqrt{d})$. All determining annuli may be enlarged in this fixed way to contain the Euclidean sphere used for the field average. For every $M < \infty$ and $\eta, \nu > 0$, with probability tending to one as $\delta \to 0$,

$$\begin{aligned} X_\delta(B_\eta(u), B_\eta(v); U) &\le D(u, v; V) + \nu, \\ D(B_\eta(u), B_\eta(v); U) &\le X_\delta(u, v; V) + \nu \end{aligned}$$

whenever the source distance is at most M, simultaneously for $|u|_\infty = 1$, $|v|_\infty = 2$. The events are local in a fixed larger annulus modulo additive constants, with lengths measured in the canonical root units.

Proof. Choose a nonnegative smooth function which vanishes on a neighborhood of V, and equals a large constant on two closed collars separating V from the boundary of U. The positive lower crossing bounds make every path from V leaving U cost more than $M+3$ in the perturbed metrics, with probability arbitrarily close to one. Apply the preceding lemma to convergence in probability for the perturbed field. Source paths in V have unchanged limiting length and a multiplicative $1+o(1)$ change in LFPP length. All resulting target paths of bounded length are confined to U. Removing the nonnegative perturbation decreases their lengths. The perturbation height was used only to prove the probability bound, so the probability of the asserted event tends to one.

Closed-constraint source metrics are lower semicontinuous by compact path extraction using the inverse Euclidean modulus. The event uses only actual internal metrics in the larger annulus and the compactly supported mollification there. This proves its locality. □

After translating and dilating by $z + r\cdot$, the two canonical root units are

$$r^{\xi Q} e^{\xi h_r(z)}, \qquad c_\varepsilon(r) e^{\xi h_r(z)}.$$

Thus target lengths acquire the factor $\kappa_\varepsilon(r)$ in the D-source direction and its reciprocal in the other direction. The local comparison is uniform for r in any fixed exponent band, since $\varepsilon/r \to 0$ uniformly there.

4 Replacing planar separation

This section contains the new geometric part of the upgrade. The source path is a **global** geodesic, although its image is contained in a fixed Euclidean ball. The distinction matters: a shortcut through a core need not remain inside that ball.

4.1 Time spent near a root

Let $P : [0, L] \to \mathbb{R}^d$ be a source geodesic parametrized by its length, with $L \le M$. For a root center c visited near its center, let $I_a(c, r)$ be the interval between the first and last visits of P to $\overline{B}_{ar}(c)$. Assume the global endpoints are outside $B_{100r}(c)$.

For geometrically separated radii $r_1 < \cdots < r_n$, with separation greater than $100/20$,

$$\prod_{j=1}^{n} \frac{|I_{100}(c, r_j)|}{|I_{20}(c, r_j)|} \le \frac{M}{|I_{20}(c, r_1)|}.$$

This follows from the nesting $I_{100}(c, r_j) \subset I_{20}(c, r_{j+1})$. The mesoscopic lower crossing estimate bounds the denominator below by $\varepsilon^{A_* t}$ whenever the smallest available radius is at least ε^t. Consequently, in a band $\varepsilon^t < r < \varepsilon^s$ with $s \le t/4$, a sufficiently large fixed A makes

$$|I_{100}(c, r)| \le A |I_{20}(c, r)|$$

fail on fewer than a prescribed small fraction of the available radii. The constant A is independent of the band and of the eventual comparison accuracy.

Choose the root core event with attachment tolerance $\delta < 1/(2A)$. In source root units S, impose a core total length bound HS in a buffer around the band $A_{20r,100r}(c)$, and a crossing lower bound mS for $A_{r,2r}(c)$. These conditions have probability as close to one as required, with fixed $H < \infty$, $m > 0$.

The incoming and outgoing arms between the global first and last visits to these balls attach to the same core. Their selected contacts bracket I_{20}, and their combined length is at most $|I_{100}|$. Geodesic minimality therefore gives

$$|I_{20}| \le \delta|I_{100}| + HS \le \delta A|I_{20}| + HS.$$

It follows that

$$\mu_P(B_{20r}(c)) \le |I_{20}| \le 2HS,$$

where μ_P is source arclength pushed forward by P. In particular, if G is any selected crossing segment in $A_{r,2r}(c)$, then

$$mS \le |G| \le 2HS, \qquad \mu_P(B_{20r}(c)) \le (2H/m)|G|.$$

This controls all visits to the inner ball, including excursions which leave it and return.

Annular concentration supplies a positive density of good roots at every center in a grid of mesh a sufficiently small fixed fraction of ε^t. The product bound excludes only a smaller density. Thus every point of every such geodesic with the required clearance from its two global endpoints has a root which is both good and satisfies the time bound. The exceptional probability is polynomial in ε, after a union bound over the grid. No path is fixed before the good event is constructed.

4.2 A spatial covering captures source length

Write $r_{\max} = \varepsilon^s$. Choose the output constant C_* much larger than all fixed annular radii. If the original endpoint separation is strictly less than $2C_* r_{\max}$, the two open output balls overlap and the desired estimate is immediate. In the remaining case, let b_0 be the first visit to a fixed large $r_{\max}$-ball about the final endpoint, and let a_0 be the last visit to the corresponding ball about the initial endpoint before b_0. Call $[a_0, b_0]$ the interval J. Taking the trimming radius large enough and then C_* larger still, J is nonempty, its endpoints are separated by more than $4r_{\max}$, and every candidate root along $P(J)$ is far from the original endpoints.

Choose one usable root at each point of $P(J)$. Only finitely many grid centers and radii occur. A greedy selection in decreasing radius order gives pairwise disjoint balls $B_{4r_i}(c_i)$ whose fivefold enlargements $B_{20r_i}(c_i)$ cover $P(J)$.

For each selected root, retain a witness time in J. At least one endpoint of $P(J)$ is outside $B_{2r_i}(c_i)$. Starting at the witness in the appropriate time direction, take the first hit of $\partial B_{2r_i}(c_i)$, and the last hit of $\partial B_{r_i}(c_i)$ before that hit. Denote the interval between them by G_i. It lies in J, and its image lies in the closed annulus $A_{r_i,2r_i}(c_i)$. Since the source is a global geodesic, its internal distance in the closed source annulus V, after scaling, equals $|G_i|$. If the time direction is reversed, symmetry swaps the two boundary endpoints. The G_i have disjoint spatial images, hence disjoint time intervals. The preceding bound implies

$$|J| \le \sum_i \mu_P(B_{20r_i}(c_i)) \le \frac{2H}{m} \sum_i |G_i|.$$

Consequently

$$\sum_i |G_i| \ge \lambda|J|, \qquad \lambda = \min\{1/2, m/(2H)\} > 0.$$

The constants determining λ were fixed before the desired accuracy. This order is essential.

4.3 Joining paths with displaced endpoints

Fix an arbitrarily small $e > 0$. Around each endpoint of a selected G_i, take a small ball of radius θr_i, with fixed $\theta > 0$ sufficiently small. The choice of θ includes the fixed enlargement needed for the connector buffers, so these buffers lie in the root determining annulus. The endpoint balls are pairwise disjoint: the two endpoints within one root are separated by at least r_i, and different roots have disjoint B_{4r_i}.

The connection estimate for two arms, together with annular concentration, has the following consequence. Fix the marginal success probability sufficiently close to one that the exponential concentration rate beats the dimension of a relative spatial grid; this fixes a connection constant C. Next choose N with $2C/N < e$, and finally choose $\eta > 0$ sufficiently small. With root probability as close to one as required, the following holds simultaneously for every possible endpoint p in the compact normalized crossing region. There is a grid point c_p, with $|c_p - p| \le \eta r_i/10$, and at least N disjoint successful annuli centered at c_p, with inner radii at least $3\eta r_i$ and all enlarged connector buffers inside $B_{\theta r_i/2}(c_p) \subset B_{\theta r_i}(p)$. Indeed, each grid point has order $\log(1/\eta)$ available separated annuli, while there are $O(\eta^{-d})$ grid points. The prescribed concentration rate makes their union bound tend to zero as $\eta \to 0$. After these fixed choices, the LFPP cutoff is taken sufficiently small. This uses annuli centered at nearby grid points; it does not claim concentric good annuli at every real center.

Replace each G_i by its sharp local target path, with endpoints in its ηr_i-balls. Replace the complementary source intervals by paths from the previous comparison, whose endpoint error is smaller than $\eta r_i/4$. At each junction the two adjacent target paths cross every one of the N target annuli. Their other endpoints lie in different disjoint junction balls. The one exception is an external endpoint already inside a junction ball; then discard the external piece and start or finish on the adjacent good piece. This moves the external endpoint only by a fixed multiple of $r_{\max}$.

For a given junction, the sums of the lengths of the two crossing arms over the N disjoint annuli are at most the total lengths of the two incident paths. One annulus therefore supplies a connector of length at most

$$\frac{C}{N}(\mathrm{len}(P_-) + \mathrm{len}(P_+)).$$

Subdivide each parametrized rectifiable piece at its finitely many selected connector attachment parameters, and add the connectors as edges of a finite abstract graph. After all junctions have been treated, this graph connects the two external endpoints. A simple graph path maps to a target path of length at most the sum of all edge lengths; incidental intersections of the geometric images cause no problem. Each original piece is charged at most twice, so the final length is at most $1 + e$ times the original sum. Infimum errors are taken arbitrarily small after the finite collection of pieces is selected.

This argument uses neither a separating curve nor a triangle inequality for distances between sets.

5 Contraction and the normalization windows

Choose $d_0 > 0$. The local comparison on each selected crossing, with source lengths between m and $2H$ in root units, gives target cost at most

$$(1 + d_0)K|G_i|,$$

where K bounds $\kappa_\varepsilon(r)$, or its reciprocal, on the deterministic set of available root radii. Let C be the previous comparison constant. The preliminary total target cost on J, writing $G = \sum_i |G_i|$, is at most

$$C(|J| - G) + (1 + d_0)KG.$$

The joining argument therefore replaces C by

$$F(C, K) = (1 + e)\max\{(1 - \lambda)C + \lambda(1 + d_0)K, (1 + d_0)K\}.$$

The new endpoint radius is $C_*\varepsilon^s$. If the previous endpoint exponent is larger than t, its radius is eventually smaller than the required $\eta\varepsilon^t/4$.

Choose e small enough that

$$a = (1 + e)(1 - \lambda) < 1, \qquad \chi = \frac{(1 + e)\lambda(1 + d_0)}{1 - a}.$$

Then $\chi \to 1$ as $e, d_0 \to 0$, and

$$F^{\circ N}(C_0, K) \le \chi K + a^N C_0.$$

Indeed, $F(\chi K, K) = \chi K$, and the excess above this fixed point contracts by at most a.

The one-annulus event consists of the source core bounds, the target annular connections, and the local sharp comparison. Its probability is at least a fixed $p_0 < 1$, as close to one as needed. Its annular determining region is fixed. At every center in a mesh of order ε^t, concentration gives a positive density of successes among any deterministic positive-density subset of the radii in the band. The mesoscopic lower bound and time-counting argument then make the contraction simultaneous over all source geodesics contained in B_R with length at most M, with polynomially small failure probability.

Proposition (comparison arbitrarily close to one). For every $u > 0$ and fixed $R, M < \infty$, there are $s, \beta > 0$ such that, with probability at least $1 - O_{R,M,u}(\varepsilon^\beta)$,

$$X_\varepsilon(B_{C_*\varepsilon^s}(x), B_{C_*\varepsilon^s}(y); \overline{B}_{R+1}) \le (1 + u)D(x, y),$$
$$D(B_{C_*\varepsilon^s}(x), B_{C_*\varepsilon^s}(y); \overline{B}_{R+1}) \le (1 + u)X_\varepsilon(x, y),$$

whenever the relevant global source geodesic lies in $\overline{B}_R$ and has length at most M.

Proof. Choose $v > 0$ small, put $k = (1 + v)^{-1} < 1$, and then choose e, d_0 so that

$$\chi k < 1, \qquad \chi(1 + v) < 1 + u/2.$$

Choose N so large that

$$a^N C_0 < \min\{(1 - \chi k)/2, u/2\}.$$

Take $3N$ ordered, disjoint exponent bands (s_j, t_j), with $s_j = t_j/4$, $t_{j+1} = t_j/8$, and $t_1 < s_0$. The preceding endpoint errors fit into every subsequent band.

If at least N bands had at least one third of their radii satisfying $\kappa_\varepsilon(r) \le k$, iterate the contraction on those radii. Along any sequence of offending cutoffs this would give, with probability tending to one, a comparison with a fixed constant strictly below one from D to X_ε. Apply it to a geodesic between a fixed distinct rational pair. Its distance is finite and positive; restrict to an event on which its length and image have fixed bounds. Lower-semicontinuous convergence in probability implies

$$\liminf_{\varepsilon \to 0} X_\varepsilon(B_{\rho_\varepsilon}(p), B_{\rho_\varepsilon}(q)) \ge D(p, q)$$

in the subsequence sense, for every deterministic $\rho_\varepsilon \to 0$. This rules out the strict contraction. Hence fewer than N bands can have this property for all sufficiently small cutoffs.

The same argument with κ_ε^{-1} rules out N bands having one third of their radii with $\kappa_\varepsilon^{-1} \le k$. In this direction use rational point-distance convergence in probability and confinement in probability from the initial comparison. Indeed, any bounded-length LFPP path leaving a large fixed ball would contradict the initial lower comparison and the limiting large-annulus barrier. Thus the global LFPP geodesic needed for the contradiction is confined with probability arbitrarily close to one. Lower semicontinuity gives $D(B_\rho(p), B_\rho(q)) \to D(p, q)$.

There remain at least N bands in which at least one third of the radii satisfy

$$(1+v)^{-1} \le \kappa_\varepsilon(r) \le 1+v.$$

Apply the contraction in both directions on these radii. Its iterated bound is at most $\chi(1+v)+a^N C_0 < 1+u$. There are only finitely many bands and possible choices of their indices, so the exceptional probabilities remain polynomial. The argument applies for every fixed localization and length cap. This proves the proposition. □

6 Almost sure convergence through every cutoff

6.1 Interpolation and the common event

The normalizers are regularly varying:

$$\frac{a_{r\varepsilon}}{a_\varepsilon} \longrightarrow r^{1-\xi Q} \qquad (r>0).$$

This is [U, Corollary 29.4]. It also follows directly by comparing the exact LFPP dilation identity with convergence in probability and the limiting dilation identity: a deterministic multiplicative factor in two approximations to the same nondegenerate lower-semicontinuous limit must tend to one. A recovery sequence rules out a factor tending to zero or infinity, and a finite limiting factor must fix a strictly positive finite distance. The lower-median normalizer is positive and Borel, so uniform convergence for regularly varying functions applies.

For each $p>0$, with $\delta_n = n^{-p}$ and $I_n = [(n+1)^{-p}, n^{-p}]$, it follows that

$$\sup_{\varepsilon \in I_n} |\log(a_\varepsilon / a_{\delta_n})| \to 0.$$

Also, for every compact K, almost surely

$$\sup_{\varepsilon \in I_n, z \in K} |h_\varepsilon(z) - h_{\delta_n}(z)| \to 0.$$

A proof valid in every dimension is given in Appendix B. Consequently every path $P \subset K$ has its normalized lengths at ε and δ_n within multiplicative factors $e^{\pm \omega_n(K)}$, where $\omega_n(K) \to 0$ uniformly for $\varepsilon \in I_n$.

For each rational accuracy $u>0$ and integer bounds R, M, choose p so large that $p\beta > 1$ for the finitely many comparison statements being used. Borel–Cantelli makes their estimates hold eventually on the mesh $\delta_n = n^{-p}$. The mesh may depend on the accuracy and localization bounds. Intersect these countably many events and the corresponding field-interpolation events. All remaining arguments are on this single event of probability one.

6.2 Confinement for all sufficiently small real cutoffs

Fix integers k, M. Choose a sufficiently large integer R, depending on the realized field, that

$$D(\overline{B}_k, \partial B_R) > 4C_0 M.$$

Apply the initial comparison on a suitable polynomial mesh to paths stopped when they first hit ∂B_R. A mesh path starting in $\overline{B}_k$, of length at most $2M$, which exits B_R, would imply

$$D(B_{\rho_n}(\overline{B}_k), B_{\rho_n}(\partial B_R)) \le 2C_0 M.$$

The left side tends to $D(\overline{B}_k, \partial B_R)$: its endpoint sets lie in a fixed compact region, so this follows from compactness and lower semicontinuity. This is impossible eventually.

If a path at an intermediate cutoff had length at most M and exited B_R, its stopped part lies in $\overline{B}_R$. Interpolation would give a neighboring mesh path of length at most $2M$, again a contradiction. Thus all such paths are confined eventually for every positive real cutoff.

This also supplies global LFPP geodesics whenever the distances under consideration are bounded. Nearly minimizing paths stay in a fixed compact ball. For each fixed cutoff the smooth density has a positive minimum there, so compactness of bounded Euclidean-length paths gives a minimizing path. This establishes the geodesic property needed for the sharp source comparison at these cutoffs.

6.3 The lower-limit inequality

Let $\varepsilon_j \to 0$, $(x_j, y_j) \to (x, y)$, and pass to a subsequence on which $X_{\varepsilon_j}(x_j, y_j)$ tends to its finite lower limit L. The infinite case is immediate. Choose integer endpoint and length bounds. The confinement just proved gives a finite integer radius containing all relevant near-minimizing paths at all sufficiently small cutoffs, including neighboring meshes and a harmlessly larger length cap.

Fix rational $u > 0$ and use a mesh for the sharp comparison with these radius and length bounds. Interpolation of an X_{ε_j}-geodesic gives

$$X_{\delta_{n(j)}}(x_j, y_j) \le e^{\omega_{n(j)}} X_{\varepsilon_j}(x_j, y_j).$$

The global mesh geodesic is confined as well. The sharp comparison therefore gives

$$D(B_{\rho_j}(x_j), B_{\rho_j}(y_j)) \le (1+u) e^{\omega_{n(j)}} X_{\varepsilon_j}(x_j, y_j), \qquad \rho_j \to 0.$$

Choose nearly minimizing endpoint pairs in these balls. They tend to (x, y), so lower semicontinuity gives $D(x, y) \le (1+u)L$. Let $u \downarrow 0$ through rationals. This proves the required lower-limit inequality simultaneously for every sequence and every endpoint pair.

6.4 Recovery

Suppose $0 < D(x, y) < \infty$. A D-geodesic from x to y has bounded length and compact Euclidean image. The first sharp comparison, on the appropriate mesh, supplies target paths with endpoints tending to (x, y) and lengths at most $(1+u)D(x, y) + o(1)$. The paths stay in the fixed comparison buffer. Interpolation uses the same endpoints and paths for every cutoff between adjacent mesh values, with a multiplicative error tending to one.

Use $u = 1/m$, and choose decreasing cutoff thresholds so that the endpoint errors and additive errors for the m-th comparison are at most $1/m$. A diagonal choice gives recovery pairs through all positive

cutoffs with upper limit at most $D(x,y)$. The lower-limit inequality gives equality. If $D(x,y)=0$, separation gives $x=y$, and the constant diagonal pair is a recovery pair. If $D(x,y)=\infty$, the lower-limit inequality applied to the constant pair already gives $X_\varepsilon(x,y)\to\infty$. This proves the theorem.

7 Relation to the earlier planar argument

Devlin's theorem [D] proves almost sure convergence in the subcritical planar setting. The earlier supercritical note [P] retains endpoint neighborhoods and uses short separating paths. The present proof keeps the compact-height comparison and normalization-window mechanism, while replacing planar separation by three statements proved above: the bound on the time a geodesic spends in a ball, spatially disjoint root selection, and connections between actual target arms. No finite-diameter deterministic hypersurface and no upper Euclidean modulus for D enters the argument.

The proof therefore uses the full positive-Q chain and core estimates from the previous construction. A claim based only on convergence in probability, with no such estimates, would not be justified.

A Finite-cutoff mixed connections and the initial comparison

We use exactly the compactly supported mollifier and the normalized metric $X_\varepsilon = a_\varepsilon^{-1} D_h^\varepsilon$ of the theorem. This appendix proves the initial comparison from the retained chain-level estimates, simultaneous lower shell charges, and annular concentration.

Fix $d\ge 3$, $Q>0$, and $\xi>0$. Put

$$c_D(r)=r^{\xi Q},\qquad c_X(r)=\frac{r a_{\varepsilon/r}}{a_\varepsilon},\qquad \kappa_\varepsilon(r)=\frac{c_X(r)}{c_D(r)}.$$

If the source and target metrics are denoted d,e, respectively, write

$$R_\varepsilon(r)=\frac{c_{\mathsf{e}}(r)}{c_{\mathsf{d}}(r)}=\kappa_\varepsilon(r)\quad\text{or}\quad \kappa_\varepsilon(r)^{-1}.$$

The common coarse random multiplier at a root (z,r) is $e^{\xi H_r(z)}$; it cancels in this ratio.

A.1 Retained chain estimates and raw relative powers

The needed stronger input underlying the nice-core proposition is the actual fine-edge chain construction, not just its final distance bound. In [C, Section 4], the point-chaining estimate gives the prelimit chain cost

$$C\sum_{\ell=M_0}^{\lfloor t\rfloor-1}(1+\ell)^{2d}2^{-\ell}b_{t-\ell}\exp\left(\xi\sup_{T_\ell(x)}H_\ell\right)+\text{terminal term},$$

where the tube radius is $C(1+\ell)^2 2^{-\ell}$, and the failure probability at any level at least M_0 is bounded by $C\sum_{\ell\ge M_0}2^{d\ell}e^{-c\ell^2}$, uniformly in the cutoff. The finite-depth retention argument below transfers the rational internal edge paths and their charges to the limiting metric.

In a fixed buffered annular geometry, at relative dyadic depth ℓ, this construction supplies a finite set of possible nice vertices and actual internal target paths connecting successive selected vertices.

If x is in the test band, its selected level-ℓ vertex $v_\ell(x)$ satisfies

$$|v_\ell(x) - x| \le C(1+\ell)^k r 2^{-\ell}. \tag{A.1.1}$$

The connection from $v_\ell(x)$ to $v_{\ell+1}(x)$, for $M_0 \le \ell < N$, has an upper charge $U^{\mathsf{e}}_\ell(x)$. A terminal straight-segment charge, when needed, is indexed by N. If the target is limiting D and the source is LFPP, define $U_N = 0$ and stop at v_N; there is no edge beyond that vertex. These charges are constant on, or bounded by labels on, a deterministic mesh, with only polynomially many enlarged supporting cells at each point. Actual paths remain inside a fixed annular buffer. At the starting depth M_0, all vertices are joined by a finite connected core of tight total target mass in root target units. The prescribed off-center point $z + (r/2)e_1$ may be added to this finite core: internal point-distance tightness in the same annular buffer controls the added paths.

For any sufficiently small slack $\eta > 0$, the upper-chain estimates, lower shell estimates, and common coarse oscillation bounds give, with probability tending to one as $M_0 \to \infty$, uniformly in the cutoff and root,

$$U^{\mathsf{e}}_\ell(x) \le C_\eta R_\varepsilon(r) 2^{3\eta\ell} L^{\mathsf{d}}_\ell(x), \qquad M_0 \le \ell \le N, \tag{A.1.2}$$

where $N = \lfloor \log_2(r/\varepsilon) \rfloor$, and L^{d}_ℓ is an actual lower charge for a completed source shell crossing. Polynomial factors are included in the third slack. For own-target attachments, the same statement holds with $R_\varepsilon(r) = 1$, through all depths when the target is D, and through the terminal depth with its terminal segment when the target is LFPP.

Here is the bookkeeping behind (A.1.2), rather than an inference from a completed same-metric connector theorem. Before dividing the target upper label by the source lower label, both contain the same factor $e^{\xi H_{r2^{-\ell}}}$. The target label is bounded by a polynomial in ℓ times its raw unit $c_{\mathsf{e}}(r2^{-\ell})$; the source lower label is at least its raw unit $c_{\mathsf{d}}(r2^{-\ell})$ times $2^{-\eta\ell}$. Coarse oscillation contributes a second arbitrarily small exponential slack. The raw relative-power estimate gives

$$\frac{c_{\mathsf{e}}(r2^{-\ell})}{c_{\mathsf{d}}(r2^{-\ell})} \le C_\eta R_\varepsilon(r) 2^{\eta\ell}. \tag{A.1.3}$$

After allocating the slacks slightly more finely, these factors give (A.1.2). Equation (A.1.3) is available from the retained two-sided relative-power estimates; alternatively it follows from regular variation of a by Potter bounds. At depths for which the relative cutoff lies in a fixed compact interval, bounded-cutoff comparisons supply the same bound. Choose the slacks so that $3\eta < 1$.

The coordinate-section estimate is deterministic. For every source arm P crossing the test band, some coordinate projection spans an interval of length comparable to r. Choose I in its middle buffered portion, at positive coordinate distance from both endpoint coordinate values, and select its first-hit section $x(u)$, $u \in I$. Increase M_0 so that every polynomially enlarged level-ℓ tube, $\ell \ge M_0$, is smaller than this fixed coordinate buffer. This guarantees a completed source shell crossing in at least one direction at each selected contact. A level-ℓ cell can be selected for levels in a set of measure at most $Cr2^{-\ell}$. Choose one completed source shell subpath for each visited cell. Bounded enlarged-support overlap, with polynomial factors already allocated to the slack, gives

$$\frac{1}{|I|}\int_I \sum_{\ell=M_0}^{N} U^{\mathsf{e}}_\ell(x(u))\,du \le C_\eta R_\varepsilon(r) \sum_{\ell \ge M_0} 2^{-(1-3\eta)\ell}\,\mathrm{len}_{\mathsf{d}}(P). \tag{A.1.4}$$

The argument is simultaneous over paths chosen after the field. No independence of the arm and the labels is used.

A.1.1 Retaining actual finite-depth chains

The transfer of universal arm connectors does not by itself give a truncated upper chain ending at a prescribed fine mesh. We record the additional compactness argument needed here. It uses the actual finite-edge construction underlying [C, "Point chaining and the derivation of scale comparisons," equation (chainraw)], rather than inferring intermediate vertices from its sum-of-costs inequality. The successive-level compatibility is also explicit in the proof of [C, Proposition "Moment exponent"]: its first paragraph imposes the successive-level nice-site events at every mesh, warns that separate good components would not suffice, and joins the endpoint chains through a finite mesh core.

Fix a compact test band A inside an open annular buffer V, with positive clearance. Fix a starting mesh depth m. On the successful event of the fine-edge construction, the following are finite construction data at each depth ℓ: the mesh-cell labels; the selected nice vertices for those cells; and the actual buffered edge paths joining successive selected vertices. For a point $x \in A$, denote its selected vertex by $v_{n,\ell}(x)$, where n is the cutoff index. Fix a half-open convention for mesh cells. The cell maps and adjacent-level choices can then be recorded by finitely many labels at every fixed finite depth. The bad-component diameter bound places all possible vertices in a deterministic finite menu and gives

$$|v_{n,\ell}(x) - x| \le C(1+\ell)^k 2^{-\ell}. \tag{A.T1}$$

All paths lie in fixed compact subsets of V. The path joining levels ℓ and $\ell+1$ has the upper charge furnished by the construction. In normalized units it is bounded by

$$U_{n,\ell}(z) = C(1+\ell)^{2d} 2^{-\ell} \frac{b_{n-\ell}}{b_n} \exp\{\xi H_{n,\ell}(z) + \xi O_{n,\ell}\}, \tag{A.T2}$$

with harmless changes for a fractional cutoff and for normalization by a comparable median. One may use an upper bound for the right side; equality is unnecessary. Here z is the relevant mesh-cell label, and $O_{n,\ell}$ is the coarse oscillation on the polynomially enlarged tube. The fixed-depth core joins the possible starting vertices. A fixed upper quantile bounds its total length and, if necessary, its finite depth. The fine-edge contribution to the failure probability is bounded as in [C]:

$$p_m^{\mathrm{chain}} \le C \sum_{\ell \ge m} 2^{d\ell} e^{-c\ell^2}. \tag{A.T3}$$

Write p_m for the total failure probability after adding the chosen core-cap and oscillation failures. Those additional failures can be made arbitrarily small by choosing their deterministic caps; with the caps chosen accordingly, $p_m \to 0$ as $m \to \infty$. The bound (A.T3) is asserted only for the fine-edge contribution.

These are the vertex-and-path data of the same construction used to prove (chainraw): its level costs are added along actual fine-edge paths. The vertices and edge paths must be retained; a bound only on the final attachment distance would not imply the statement below.

Lemma (finite-chain transfer). On the common-field coupling with the canonical limiting metric $D = D_h$, the preceding finite-edge construction has a limiting version with probability at least $1 - p_m$. On this event there are compatible cell maps v_ℓ and actual D-rectifiable edge paths, for every $\ell \ge m$, such that:

1. Each cell map has a deterministic finite vertex menu, and (A.T1) holds with v_ℓ in place of $v_{n,\ell}$.
2. Every finite chain from the depth-m core to $v_N(x)$ is an actual path in V, with length at most the sum of its limiting upper charges, for all $x \in A$ simultaneously.

3. After an arbitrarily small allocation of exponential slack, those charges have the same coarse-field form as the canonical limiting upper-chain charges:
$$U_\ell(z) \le C_\eta(1+\ell)^{2d}2^{-(\xi Q-\eta)\ell}\exp\{\xi H_\ell(z)+\xi O_\ell\}. \tag{A.T4}$$
In physical root units, multiply this bound by the root unit. The raw relative-power bound supplies the deterministic exponent in (A.T4).
4. Any fixed finite connected core-mass bound passes to a limiting connected core with the same bound. For a point at which the upper-charge series is finite, the infinite chain also gives an actual attachment to that point.

Proof. Work on the original probability space of the field and all its mollifications. Enlarge this space by the previously used white-noise decomposition if necessary; this does not change the joint law of $(h, D_h, (D_n)_n)$. Convergence in probability to D_h allows a deterministic subsequence with almost-sure epigraph convergence. The compact path extraction property [U, input 2(iii)] applies on this convergence event. For paths in a fixed compact buffer, its needed lower barriers also follow directly from the moving-endpoint lower bound and separation of D_h: otherwise a further convergent endpoint subsequence with a fixed positive Euclidean separation would have zero limiting distance. Consequently all the extraction below can be made on this original common-field realization.

For each fixed ℓ, retain the whole driving field and its smooth coarse potential. In the common-field realization these are the same coordinates throughout; if a regularized coarse potential is used, it converges locally uniformly to that coordinate. Their suprema on fixed compact tubes therefore converge as well. The raw relative-power estimates give a deterministic bound for the normalized coefficient in (A.T2). Thus, for each fixed cell and level,
$$\limsup_n U_{n,\ell}(z) \le U_\ell(z) < \infty,$$
with U_ℓ satisfying (A.T4). If exact coefficient limits are desired, take a further deterministic subsequence for their countably many bounded values. Only the displayed upper bound is needed.

Let $G_{n,m}$ be the event that the prelimit construction succeeds at every available level from m to its terminal cutoff, including the chosen finite core cap. The uniform estimate gives
$$\mathbb{P}(\limsup_n G_{n,m}) \ge \limsup_n \mathbb{P}(G_{n,m}) \ge 1-p_m.$$
Work on the limsup event and select an infinite sequence of successful cutoffs. This sequence may depend on the realization: the argument proves existence of limiting witnesses and does not require a measurable choice of them.

First fix a finite depth N. Only finitely many mesh cells, vertex choices, and adjacent-level combinatorial choices occur up to that depth. Pass to a subsequence on which all these discrete choices are constant. For every chosen edge, select the actual prelimit path supplied by the construction, with an additive length error tending to zero if its cost was stated as an infimum. Its endpoints are now fixed vertices; its image is confined to a fixed compact buffer; and its lengths are bounded by the convergent coarse-charge bounds. Parametrize it proportionally to its prelimit metric length.

The compact path extraction property [U, input 2(iii)] gives a uniformly Euclidean-convergent subsequence. Its limiting path joins the same vertices. The moving-endpoint lower bound, applied on finite time partitions, gives
$$\mathrm{len}(P; D) \le \liminf_n \mathrm{len}(P_n; D_n) \le U_\ell(z).$$

The proportional-to-length parametrization also gives $D(P(s), P(t)) \le L|s-t|$ for a uniform length bound L, so the limit is a metric-continuous rectifiable path. This step uses compact extraction of actual paths, not convergence of endpoint distances alone.

There are finitely many paths through depth N, so they can be extracted together. Repeat for $N = m+1, m+2, \ldots$, always refining the previous subsequence, and take a diagonal subsequence. The stabilized discrete maps and paths are compatible at every finite depth. Consequently the cell rule gives $v_\ell(x)$ for every $x \in A$, and the selected edges concatenate to the required finite chain. No intersection over uncountably many points is taken: all points use the same countable collection of finite cell maps and edge paths. Bound (A.T1) passes unchanged.

The core has finitely many possible vertices and edges after its depth cap. Stabilize its abstract graph at the same time and extract its paths. Connectedness is preserved because the endpoint incidences are fixed. For its finitely many edges, the lower bound for limiting lengths and the finite-sum inequality give

$$\sum_e \operatorname{len}(P_e; D) \le \liminf_n \sum_e \operatorname{len}(P_{n,e}; D_n).$$

Thus its total-mass bound is preserved.

Finally, if $\sum_{\ell \ge m} U_\ell(z_\ell(x)) < \infty$, the compatible chain vertices are Cauchy for D and converge Euclideanly to x by (A.T1). Completeness of the finite component and Euclidean continuity identify their metric limit with x. Equivalently, lower semicontinuity applied to the tails gives $D(v_N(x), x) \le \sum_{\ell \ge N} U_\ell \to 0$. Concatenation and endpoint completion give an actual attachment. This proves the lemma. □

Locality of the event used in annular concentration. The auxiliary coarse potentials in (A.T2) need not themselves be measurable from the field restricted to the annulus. They are used to prove the marginal probability, and are not included in the event to which annular concentration is applied. After coordinate averaging, define that event intrinsically: every source and target arm in the test band has the asserted target connection in the fixed buffer, with the stated relative length bound and terminal endpoint tolerance; the prescribed finite core and its internal connections satisfy the asserted cost bounds. All these quantities use actual internal metrics in the buffer. If a root unit is used, its averaging sphere is included in the determining annulus. Locality of the metric assignments and the support of the compact mollifier therefore make this intrinsic event measurable from the field in the determining annulus. Adding a constant scales every metric length and the root unit by the same factor, so the event is invariant under that addition. Universal arm statements are measurable in the completed local sigma-field by the compact-path/analytic-projection argument of [U, Section 14]. The transferred chain witnesses prove that this intrinsic event has the required high marginal probability. This separation between the probability witnesses and the intrinsic local event is essential.

Application to the mixed cutoff lemma. Intersect this limiting-chain event with the prelimit upper-chain and the two source lower-shell events. At each depth, their coarse potentials are the same field coordinates. Divide the target upper charge by the source lower charge and use the raw relative-power estimates. The prescribed small oscillation and label slacks give $U_\ell^{\mathsf{e}} \le C_\eta R_r 2^{3\eta\ell} L_\ell^{\mathsf{d}}$, after dividing the slack budget among these factors. For an LFPP source and limiting target, use the finite chain ending at v_N, so no unbounded tail is introduced. For a limiting source and LFPP target, use the LFPP terminal straight segment and the limiting lower shell charge at the same depth. The transfer lemma introduces no new probability estimate beyond the uniform prelimit

chain failure bound, the raw relative powers, the Gaussian coarse-coordinate convergence, and the compact path extraction already available in [C], [T], and [U].

A.1.2 The two terminal cases

Source LFPP, target D. Stop the target chain at its last nice vertex $v_N(x(u))$. Do not append an infinite target tail to the selected source contact. Equations (A.1.1) and (A.1.4) give a target path from a core vertex to a point v with

$$\mathrm{dist}_{\mathrm{Euc}}(v, P) \le C\varepsilon(1 + \log(r/\varepsilon))^k, \tag{A.1.5}$$

at the cost in (A.1.4). The point v has finite target distance because it is an actual nice vertex. This step makes no assertion that an arbitrary LFPP-arm point is nonsingular for D.

Source D, target LFPP. Append the straight segment from $v_N(x(u))$ to the selected contact $x(u)$. Its LFPP upper charge is bounded by the length of the terminal tube times the maximum smoothed density there. In root coordinates, with $\delta = \varepsilon/r$, the deterministic factor obeys

$$\frac{\delta}{a_\delta} \le C_\eta \delta^{\xi Q - \eta}. \tag{A.1.6}$$

Here the difference between the smoothed field and the common coarse field has a direct Gaussian estimate. In the retained white-noise representation, write the scale kernel as $k_u(x) = u^{-d/2}k(x/u)$, with scale measure du/u, and let H_N contain the scales $u \ge 2^{-N}$. Since $\delta \asymp 2^{-N}$, smoothness and Young's inequality give

$$\|k_u * \rho_\delta - k_u\|_2^2 \le C(\delta/u)^2 \quad (u \ge \delta), \qquad \|k_u * \rho_\delta\|_2^2 \le C(u/\delta)^d \quad (u \le \delta).$$

For the second bound use $\|k_u\|_1 = O(u^{d/2})$ and $\|\rho_\delta\|_2 = O(\delta^{-d/2})$; the first follows by integrating the gradient along translates. Consequently the coarse smoothing error and the smoothed fine tail satisfy

$$\sup_z \mathrm{Var}\big((h * \rho_\delta)(z) - H_N(z)\big) \le C\int_\delta^\infty (\delta/u)^2 \frac{du}{u} + C\int_0^\delta (u/\delta)^d \frac{du}{u} \le C.$$

Include the smooth pinning correction in the common coarse field; its mollification error obeys the same bounded-variance estimate in the fixed root buffer. The corresponding rescaled derivative bounds give a uniform Gaussian supremum tail on each cutoff cell. A union over $O(2^{dN})$ cells therefore bounds the probability that the displayed field difference exceeds ηN anywhere in the root buffer by $C_\eta e^{-c_\eta N^2}$, for sufficiently large N.

This difference and the coarse oscillation on the polynomially enlarged terminal tube are absorbed in an additional arbitrarily small ηN slack. The retained terminal oscillation estimate is summable after the union over all terminal cells: its tail has the form

$$C\exp\{-c\eta^2 N^2/\log(N+2)\}.$$

The source limiting-metric lower shell charge at that scale has all inverse moments, so it supplies the same lower-label slack as in (A.1.2). Thus the terminal straight segment is included as a final upper charge in (A.1.4), and the connector ends exactly on P.

For the target's *own* arm, the attachment is exact in either case. In the limiting metric, the infinite chain has summable length and its vertices converge Euclideanly to the selected contact; lower semicontinuity and completeness append that contact as an actual metric endpoint. In LFPP, the terminal straight segment does so.

A.2 The finite-cutoff mixed-core lemma

Fix $0 < s < t < 1$ and root radii $\varepsilon^t \le r \le \varepsilon^s$. Put

$$\rho_0 = \varepsilon^{(1+t)/2}.$$

For sufficiently small ε, the terminal error (A.1.5) is at most $\rho_0/2$, uniformly in these roots, while $\rho_0 = o(\varepsilon^t)$.

For every attachment tolerance $\tau > 0$ and every $p < 1$, there is a finite K_0, independent of the root and cutoff, such that a local annular event of probability at least p gives both assertions below, simultaneously over all finite crossing arms.

1. Given a target arm P_{e} and a source arm P_{d}, there is a target connector starting on P_{e}, ending at a point v within $\rho_0/2$ of some selected point on P_{d}, and satisfying
$$\mathrm{len}_{\mathsf{e}}(\text{connector}) \le \tau\, \mathrm{len}_{\mathsf{e}}(P_{\mathsf{e}}) + K_0 R_\varepsilon(r)\, \mathrm{len}_{\mathsf{d}}(P_{\mathsf{d}}). \tag{A.2.1}$$
2. The prescribed off-center point $q = z + (r/2)e_1$ can be joined in the target metric to such a v at cost at most
$$K_0 R_\varepsilon(r)\, \mathrm{len}_{\mathsf{d}}(P_{\mathsf{d}}). \tag{A.2.2}$$

When the target is LFPP, the terminal point is an exact source-arm contact. Every path stays in a fixed buffered annulus, such as $B_{32r}(z) \setminus B_{a_d r}(z)$, where $0 < a_d < 1/(100\sqrt{d})$ is fixed. The dimension-dependent inner margin also allows the determining annulus to contain an entire spherical-average surface if the annular concentration argument uses that normalization. In particular the seed point is off-center to preserve annular locality.

Proof. First choose the target's own attachment depth so that its cost is at most τ times its arm length. Increase the common depth if needed so that the mixed attachment from (A.1.4) costs at most $R_\varepsilon(r)$ times the source arm length. Impose an upper quantile H on the total target core mass, including the finitely many paths to the prescribed off-center point, in units $c_{\mathsf{e}}(r)e^{\xi H_r(z)}$. Impose a lower quantile $b > 0$ on the source crossing in units $c_{\mathsf{d}}(r)e^{\xi H_r(z)}$. These events have arbitrarily high probability. The core cost is then at most

$$\frac{H}{b} R_\varepsilon(r)\, \mathrm{len}_{\mathsf{d}}(P_{\mathsf{d}}).$$

Join the attachment vertices through the one connected core. Take $K_0 \ge 1 + H/b$. The terminal descriptions above prove the asserted endpoint property. Locality follows from the actual chain paths and labels lying in the buffered annulus; no source-curve regularity in the target metric is assumed. □

A.3 Finite-band lower barriers and first-exit balances

Fix a compact working ball B_{R+1}, a source cost cap $M < \infty$, and a positive fraction $\delta_* > 0$ of the available geometric radii. We will use exponent bands with $s \le t/4$. The retained lower shell tails and coarse Gaussian bounds imply that, except on a polynomial-probability error, every path in either metric that starts within $\varepsilon^t/20$ of a grid center and leaves its innermost test region has length at least

$$\varepsilon^{At} \tag{A.3.1}$$

for a fixed deterministic A. This is obtained by a union over $O(\varepsilon^{-dt})$ centers and sufficiently high inverse moments; the common coarse field minimum is at least $-Ct \log \varepsilon^{-1}$ on the same event.

The normalizer exponent alone gives, uniformly over each fixed exponent band,

$$\sup_r \big(\kappa_\varepsilon(r) \vee \kappa_\varepsilon(r)^{-1}\big) = \varepsilon^{-o(1)}. \tag{A.3.2}$$

Indeed, $\log a_u = (1-\xi Q)\log u + o(|\log u|)$, uniformly for u in the interval between ε^{1-s} and ε^{1-t}. In particular the supremum in (A.3.2) is at most ε^{-t} for sufficiently small ε.

Suppose there are at least $\nu t \log \varepsilon^{-1}$ successful nested routing annuli at a center, where $\nu > 0$ is fixed. Consider a source forward germ of length at most M, and a reversed target prefix of length at most $CKM+1$, with $K \le \varepsilon^{-t}$. Assume both leave every routing outer region. Write $b_i^{\mathsf{d}} \le a_i^{\mathsf{d}}$ and $b_i^{\mathsf{e}} \le a_i^{\mathsf{e}}$ for their first inner and outer exit lengths, ordering successful annuli from outside to inside. Nesting gives $a_{i+1} \le b_i$, hence

$$\sum_i \log \frac{a_i}{b_i} \le \log \frac{a_1}{b_N}.$$

For sufficiently small ε, each numerator is at most ε^{-2t}, and (A.3.1) bounds the denominator below by ε^{At}. Thus both sums are at most

$$(A+2)t \log \varepsilon^{-1}.$$

Choose $\theta \in (0,1)$ with

$$\log(1/\theta) > 8(A+2)/\nu. \tag{A.3.3}$$

Fewer than a quarter of the successful annuli can violate each of

$$b^{\mathsf{d}} \ge \theta a^{\mathsf{d}}, \qquad b^{\mathsf{e}} \ge \theta a^{\mathsf{e}}. \tag{A.3.4}$$

At least one annulus satisfies both. The same argument with only the source germ is available for seed selection, although the seed below does not require a balance.

Crucially, θ is fixed before τ and K_0, and is independent of the comparison accuracy. The fixed constants C, M are absorbed by decreasing the cutoff threshold, not by changing θ.

A.4 Ratio-weighted coarse comparison

Choose a fixed geometric separation B large enough for the local annular concentration lemma. Let

$$\mathcal{R}_\varepsilon(s,t) = \{B^j : \varepsilon^t < B^j < \varepsilon^s\}.$$

Let S_ε be any deterministic subset with at least $\delta_* \# \mathcal{R}_\varepsilon(s,t)$ elements. Set

$$K = \max_{r \in S_\varepsilon} R_\varepsilon(r).$$

There are constants $C, C_{\mathrm{out}} < \infty$ such that, with polynomially high probability as $\varepsilon \to 0$, every source path $\Gamma \subset B_R$ from x to y of length $L \le M$ satisfies

$$\mathsf{e}\big(B_{C_{\mathrm{out}}\varepsilon^s}(x), B_{C_{\mathrm{out}}\varepsilon^s}(y); \overline{B}_{R+1}\big) \le CKL. \tag{A.4.1}$$

The constants can be fixed for bands with $s \le t/4$ and a fixed δ_*. The probability exponent and sufficiently small cutoff threshold may depend on the band, R, and M.

A.4.1 High-probability geometry

Order the radii of S_ε from largest to smallest. Reserve the largest quarter for seeds, the smallest half for routing, and discard the middle quarter. Rounding changes only boundedly many radii. The smallest seed radius and largest routing radius differ by a factor at least $B^{c\#S_\varepsilon}$, so

$$\frac{r_{\text{route,max}}}{r_{\text{seed,min}}} \le \varepsilon^{c_1} \tag{A.4.2}$$

for some $c_1 > 0$.

Parameter order: fix the annular geometry and retained scale fractions first. These fix a positive success-density target ν, independently of the attachment tolerance. Choose the barrier exponent A, then θ by (A.3.3), then $\tau < \theta/2$. Choose the one-annulus success probability sufficiently close to one for the desired spatial union bound, and obtain the mixed-core constant K_0 for these fixed choices.

Now apply the retained annular concentration lemma to each deterministic subfamily, and then a union bound over a mesh of size $\varepsilon^t/100$ covering B_{R+1}. Except on a polynomial-probability error, each mesh center has a successful seed radius and at least $\nu t \log \varepsilon^{-1}$ successful routing radii. The same common center, within $O(\varepsilon^t)$ of a given point, is used for all its routing radii. Intersect with the lower-barrier event of Section A.3.

Fix this realization. All local connectors lie in $\overline{B}_{R+1}$ for small enough cutoff.

A.4.2 A seed after the last exit

If $|x - y| \le C_{\text{out}}\varepsilon^s$, the two output balls overlap and (A.4.1) is immediate. Otherwise take a successful seed root (c, r) with $|c - x| \le \varepsilon^t/100$. Choose the fixed output constant larger than all geometric buffer constants. Both endpoints are then respectively inside the inner seed ball and outside its outer ball.

Let σ be the last time that Γ visits $\overline{B}_r(c)$, and let υ be its next hit of $\partial B_{2r}(c)$. The segment $\Gamma|_{[\sigma,\upsilon]}$ is a source crossing arm. Put

$$q = c + (r/2)e_1.$$

The source-only mixed attachment provides a target path from q to v_0, where

$$|v_0 - \Gamma(t_0)| \le \rho_0/2, \qquad t_0 \in [\sigma, \upsilon], \qquad \text{len}_{\mathsf{e}}(q \rightsquigarrow v_0) \le K_0 K L. \tag{A.4.3}$$

The point q is in the initial output ball. For every $t \ge t_0$,

$$|\Gamma(t) - q| \ge r/2. \tag{A.4.4}$$

Equations (A.4.2) and (A.4.4) guarantee that every target prefix starting at this fixed q and ending within ρ_0 of $\Gamma(t)$ crosses every routing annulus when read backward. There is no repeated reseeding and no accumulation of endpoint displacement.

A.4.3 Maximal reachable prefix

Parametrize Γ by source length on $[0, L]$. Write E for the target internal metric on $\overline{B}_{R+1}$. It is lower semicontinuous; compact path extraction using the inverse Euclidean bounds proves this for D, and it is elementary for smooth LFPP. Its finite distances are attained, although an arbitrarily close minimizing path also suffices below.

Fix $\alpha \in (0,1)$ and define

$$F(t) = K_{\text{route}} K t + K_0 K L + \alpha, \qquad K_{\text{route}} = 2K_0/\theta,$$
$$T = \left\{ t \in [t_0, L] : E(q, \overline{B}_{\rho_0}(\Gamma(t))) \le F(t) \right\}. \tag{A.4.5}$$

The set is nonempty by (A.4.3). It is closed: if $t_n \to t$, choose almost minimizing endpoints in the compact moving balls, take a convergent subsequence, and use lower semicontinuity of E. Thus T has a largest point t.

If $\Gamma(t)$ lies within a fixed multiple of $r_{\text{route,max}}$ of y, a target endpoint witnessing (A.4.5) lies in the final output ball, and (A.4.1) follows with $C = K_{\text{route}} + K_0$, after sending $\alpha \downarrow 0$.

Otherwise the source forward germ from $\Gamma(t)$ leaves every routing outer region. Take a target prefix from q to a minimizing endpoint $v \in \overline{B}_{\rho_0}(\Gamma(t))$, and reverse it. By (A.4.4), this target germ also leaves every routing outer region. Its length is at most $F(t) \le CKM + 1$. Section A.3 yields a successful routing annulus with both balances (A.3.4).

In each germ, use the crossing arm between its last inner-boundary contact before its first outer exit and that outer exit. The selected target contact removes a prefix tail of length

$$s_* \ge b^{\mathsf{e}} \ge \theta a^{\mathsf{e}}.$$

The source contact is $\Gamma(t')$ with

$$t' - t \ge b^{\mathsf{d}} \ge \theta a^{\mathsf{d}} > 0.$$

Delete the target tail and append (A.2.1). The new endpoint is within $\rho_0/2$ of $\Gamma(t')$, and the new path has length at most

$$\begin{aligned} F(t) - s_* + \tau a^{\mathsf{e}} + K_0 K a^{\mathsf{d}} &\le F(t) - (1 - \tau/\theta) s_* + (K_0 K/\theta)(t' - t) \\ &< F(t'). \end{aligned} \tag{A.4.6}$$

Thus $t' \in T$, contradicting maximality. If the prefix is only nearly minimizing, choose its additional error before selecting the balanced annulus, smaller than $\varepsilon^{At}/8$. The uniform barrier (A.3.1) makes this less than $s_*/8$. The connector approximation error can then be chosen below $s_*/8$; the strict margin in (A.4.6) absorbs both. Therefore the endpoint-near-y alternative must hold, proving (A.4.1).

This proof requires no source geodesic property: it works for every finite source path of length at most M. It also does not assume global LFPP confinement, because the target metric is constrained to the compact working ball.

A.5 Majority bounds for the deterministic ratios

Apply (A.4.1) in the two directions. For any deterministic positive-fraction radius subset,

$$X_\varepsilon(B_\rho x, B_\rho y; \overline{B}_{R+1}) \le C \left(\max_{r \in S} \kappa_\varepsilon(r) \right) \text{len}_D(\Gamma), \tag{A.5.1}$$

and

$$D(B_\rho x, B_\rho y; \overline{B}_{R+1}) \le C \left(\max_{r \in S} \kappa_\varepsilon(r)^{-1} \right) \text{len}_{X_\varepsilon}(\Gamma), \tag{A.5.2}$$

where $\rho = C_{\text{out}} \varepsilon^s$, source paths stay in B_R, and their cost is at most M.

Fix two distinct deterministic rational points p, q in a buffered ball. The retained point upper tightness provides deterministic R, M such that source paths between them in either metric have cost at most M, with probability at least $9/10$, uniformly for small cutoff. For the limiting metric this uses finiteness of rational internal distances. The annular lower tails provide $m > 0$ such that the target distances between sufficiently small endpoint neighborhoods are at least m, with probability at least $9/10$, in either metric. Alternatively, the LFPP lower bound follows from l.s.c. convergence in probability and the limiting lower bound from lower semicontinuity.

If a positive fraction of the available radii had $\kappa_\varepsilon(r) < T^{-1}$, use that subset in (A.5.1). With probability bounded away from zero for sufficiently small cutoff, its left side is at least m and its right side is at most CM/T. Taking $T > 2CM/m$ is impossible. If a positive fraction had $\kappa_\varepsilon(r) > T$, apply (A.5.2) in the same way. All constants are deterministic and independent of the chosen bad subset.

Consequently there is a deterministic $T < \infty$ such that, in every sufficiently small cutoff, all but fewer than twice the selected fraction of available radii satisfy

$$T^{-1} \le \kappa_\varepsilon(r) \le T. \tag{A.5.3}$$

For instance choose the bad-subset threshold to be one quarter; then at least half the radii satisfy (A.5.3). Using those radii in (A.5.1) and (A.5.2) proves the desired initial two-sided comparison with a deterministic constant $C_0 = CT$, polynomial failure probability, vanishing endpoint radius, and a fixed source cost cap.

The normalizer bootstrap itself is exactly a comparison of nondegenerate crossing or point observables. It does not require an a priori bounded κ, nor pointwise convergence of LFPP at every random endpoint.

B Cutoff interpolation in every dimension

This appendix concerns the compactly supported smooth probability mollifier used in the theorem. It uses only Gaussian estimates for the field and regular variation of the deterministic normalizers. In particular, it requires no upper Euclidean modulus for the limiting metric.

B.1 Gaussian control of the logarithmic cutoff derivative

Let $\varrho \in C_c^\infty(\mathbb{R}^d)$ be the fixed mollifier, with $\int \varrho = 1$, and put

$$\varrho_\varepsilon(x) = \varepsilon^{-d}\varrho(x/\varepsilon), \qquad h_\varepsilon = h * \varrho_\varepsilon.$$

Define

$$\psi(x) = -d\varrho(x) - x \cdot \nabla\varrho(x), \qquad G_\varepsilon(x) = \varepsilon\partial_\varepsilon h_\varepsilon(x) = (h * \psi_\varepsilon)(x).$$

Integration by parts gives $\int \psi = 0$. Each realization of the generalized function h gives a jointly smooth function of (x, ε) for $\varepsilon > 0$.

Lemma. For every compact $K \subset \mathbb{R}^d$, almost surely

$$\sup_{x\in K} |G_\varepsilon(x)| = O\Big(\sqrt{\log \varepsilon^{-1}}\Big) \qquad (\varepsilon \downarrow 0). \tag{B.1}$$

The estimate holds uniformly through all real cutoffs, on one event for all compact sets.

Proof. The covariance of a pinned whole-space log-correlated field, when paired against two zero-integral test functions, is the bilinear form with kernel $\log|u-v|^{-1}$. All terms introduced by the additive normalization vanish in such pairings. Consequently

$$\operatorname{Var} G_\varepsilon(x) = \iint \psi(u)\psi(v)\log|u-v|^{-1}\,du\,dv =: \sigma_\psi^2 < \infty, \tag{B.2}$$

independently of x and ε. In the change of variables giving (B.2), the additional term $\log\varepsilon^{-1}$ vanishes because $\int\psi = 0$.

The spatial derivatives of G_ε have variances at most $C\varepsilon^{-2}$. Its derivative with respect to $\log\varepsilon$ has bounded variance: its test kernel is another fixed smooth zero-integral function, dilated by ε. These statements follow from the same calculation as (B.2).

Write $Z(x,t) = G_{2^{-t}}(x)$. On

$$\mathcal{T}_m = [-R,R]^d \times [m, m+1], \qquad m \geq 1,$$

the Gaussian canonical distance therefore satisfies

$$\big(\mathbb{E}|Z(x,t) - Z(y,s)|^2\big)^{1/2} \leq C\big(2^m|x-y| + |t-s|\big). \tag{B.3}$$

Indeed, integrate the spatial and time derivatives along coordinate segments and use the triangle inequality in L^2.

Here is a direct finite-net argument for the supremum. Take R to be a positive integer. At level $j \geq 0$, use a grid in $\mathcal{T}_m$ of spatial mesh 2^{-m-j} and time mesh 2^{-j}. Its size is at most

$$C_R 2^{dm+(d+1)j}.$$

The grids can be nested, and each point of a level-j grid has a parent at level $j-1$ whose canonical distance from it is at most $C2^{-j}$.

By the Gaussian tail bound and a union bound, after fixing a sufficiently large deterministic A, the following hold except on an event of probability at most $C_R e^{-cm}$: every level-zero value has absolute value at most $A\sqrt{m}$, and every parent increment at level $j \geq 1$ has absolute value at most

$$A2^{-j}\sqrt{m+j}.$$

For completeness, the level-j failure probability is bounded by

$$C_R 2^{dm+(d+1)j} \exp\{-c_0 A^2(m+j)\}.$$

Choose A so that the negative exponent dominates both cardinality factors, and then sum over j. The level-zero estimate is obtained in the same way from (B.2).

Telescoping along parents and using

$$\sqrt{m} + \sum_{j\geq 1} 2^{-j}\sqrt{m+j} \leq C\sqrt{m}$$

shows that $\sup_{\mathcal{T}_m} |Z| \leq C\sqrt{m}$ on this event. Joint smoothness passes the estimate from the union of the grids to all of $\mathcal{T}_m$. Borel–Cantelli now gives this bound for all sufficiently large m, almost surely. Intersect over integer R and use $m \asymp \log \varepsilon^{-1}$. This proves (B.1). □

Corollary. Fix $p > 0$, set $\delta_n = n^{-p}$, and let $I_n = [(n+1)^{-p}, n^{-p}]$. For every compact K, almost surely

$$\sup_{\varepsilon,\delta\in I_n} \sup_{x\in K} |h_\varepsilon(x) - h_\delta(x)| = O_{p,K}\left(\frac{\sqrt{\log n}}{n}\right) \longrightarrow 0. \tag{B.4}$$

Proof. Integrate $G_u(x)$ against $d\log u$. The logarithmic width of I_n is

$$p\log(1+1/n) \leq p/n.$$

On this interval, (B.1) bounds the integrand by $C\sqrt{1+p\log(n+1)}$, uniformly on K. Their product proves (B.4). The event furnished by (B.1) works for every fixed p; no probability estimate is needed for the choice of p. □

B.2 The normalizer is Borel and interpolates on the same mesh

Let

$$Y_\varepsilon = D_h^\varepsilon(0, e_1; B_2), \qquad a_\varepsilon = \inf\{a > 0 : \mathbb{P}[Y_\varepsilon \leq a] \geq 1/2\}.$$

For each positive cutoff, Y_ε is finite and strictly positive almost surely: the smoothed density has a finite positive minimum and maximum on $\overline{B}_2$. It follows that $0 < a_\varepsilon < \infty$.

The family $(\varepsilon, \omega) \mapsto Y_\varepsilon(\omega)$ is jointly measurable. One can compute the distance as an infimum over a countable family of polygonal paths with rational vertices, including the two fixed endpoints and contained in B_2. Piecewise smooth competitors are approximated by such paths inside the open domain, and the density is continuous. Every path integral is jointly measurable. In fact, local uniform continuity of the smoothed density in the cutoff gives continuity of $Y_\varepsilon(\omega)$ for each realization and every positive cutoff.

Thus $\varepsilon \mapsto \mathbb{P}[Y_\varepsilon \le q]$ is Borel for every $q > 0$. For any $b > 0$,

$$\{\varepsilon : a_\varepsilon < b\} = \bigcup_{q \in \mathbb{Q} \cap (0,b)} \{\varepsilon : \mathbb{P}[Y_\varepsilon \le q] \ge 1/2\},$$

so a_ε is Borel. No uniqueness or continuity of the crossing-distance median is needed here.

Recall the established regular variation statement

$$\frac{a_{r\varepsilon}}{a_\varepsilon} \longrightarrow r^{1-\xi Q} \qquad \text{for each fixed } r > 0. \tag{B.5}$$

The uniform convergence theorem for positive Borel regularly varying functions [BGT, Theorem 1.2.1] makes the convergence in (B.5) uniform for r in each compact subinterval of $(0, \infty)$. Since

$$\frac{\varepsilon}{\delta_n} \in \left[\left(\frac{n}{n+1}\right)^p, 1\right] \qquad (\varepsilon \in I_n),$$

and this interval shrinks to $\{1\}$, we obtain

$$\sup_{\varepsilon \in I_n} \left|\log \frac{a_\varepsilon}{a_{\delta_n}}\right| \longrightarrow 0. \tag{B.6}$$

The use of regular variation in (B.5) is compatible with the lower-semicontinuous topology. The following elementary observation explains this point. Suppose random lower semicontinuous functions F_n, G_n on $\mathbb{R}^{2d}$ both converge in probability to F, with $F_n = b_n G_n$ for deterministic $b_n > 0$, and almost surely F has a finite strictly positive value. Then $b_n \to 1$. From any subsequence, pass to a further subsequence on which b_n converges in $[0, \infty]$ and both function sequences converge almost surely. A recovery sequence for G_n excludes the limit zero; a recovery sequence for F_n excludes infinity. At a finite positive limiting factor b, these two recovery sequences and the two lower-limit inequalities give, at a point z with $0 < F(z) < \infty$,

$$F(z) \le bF(z), \qquad F(z) \le b^{-1}F(z).$$

Hence $b = 1$. In the present setting the exact LFPP identity is

$$X_\varepsilon^h(r\cdot, r\cdot) = \kappa_\varepsilon(r) X_{\varepsilon/r}^{h(r\cdot)+Q\log r}(\cdot,\cdot).$$

Both sides after removing the deterministic factor converge in probability to $D_h(r\cdot, r\cdot)$, by limiting dilation covariance and scale invariance of the field modulo constants. Thus $\kappa_\varepsilon(r) \to 1$, which is (B.5) after replacing the fixed dilation by its reciprocal. This also proves directly the normalizer covariance statement cited from [U].

B.3 Uniform comparison of path lengths

For a rectifiable path $P \subset K$, write

$$\mathcal{L}_\varepsilon(P) = a_\varepsilon^{-1} \int_P e^{\xi h_\varepsilon(z)} \, |dz|.$$

Define

$$\omega_n(K) = \xi \sup_{\varepsilon \in I_n, z \in K} |h_\varepsilon(z) - h_{\delta_n}(z)| + \sup_{\varepsilon \in I_n} \left|\log \frac{a_\varepsilon}{a_{\delta_n}}\right|.$$

Equations (B.4) and (B.6) imply $\omega_n(K) \to 0$ almost surely. Pointwise comparison of the two normalized densities gives

$$e^{-\omega_n(K)}\mathcal{L}_{\delta_n}(P) \le \mathcal{L}_\varepsilon(P) \le e^{\omega_n(K)}\mathcal{L}_{\delta_n}(P), \qquad \varepsilon \in I_n. \tag{B.7}$$

The estimate is simultaneous over every rectifiable path contained in K. It therefore applies to infima of constrained path lengths, and to stopped paths before their first exit from a fixed compact ball. This is the interpolation used in the proof of almost sure convergence.

References

[T] *Positive-Q supercritical tightness in higher dimensions: Completion by averaged endpoint penalties.* Earlier research manuscript, 6 November 2026; file `Q>0tightness.pdf`. The common point/annulus normalization, relative-power bounds, actual chains, and limiting construction are inputs here. Manuscript.

[C] *Supercritical tightness in dimension d at least three: A conditional extension, quantitative reductions, and the remaining estimate.* Earlier research manuscript, 25 October 2026; Sections 4–6. Its scalar comparison hypothesis is supplied by [T]. Manuscript.

[U] *Strong uniqueness for every positive Q in dimension d at least three.* Earlier research manuscript, 7 November 2026, revised version; in particular Theorem 1.1, Sections 2–6, Lemma 9.1, Section 14, and Corollary 29.4. Manuscript.

[P] *Almost sure convergence of supercritical Liouville first passage percolation.* Earlier research manuscript, 1 October 2026. Manuscript.

[D] C. Devlin VI. *Almost Sure Convergence of Liouville First Passage Percolation.* Probability Theory and Related Fields **193** (2025), 247–288. arXiv:2309.08001.

[DG] J. Ding and E. Gwynne. *Up-to-constants comparison of Liouville first passage percolation and Liouville quantum gravity.* arXiv:2108.12060.

[CHZ] A. A. Contreras Hip and Z. Zhuang. *Weak exponential metrics for high-dimensional log-correlated Gaussian fields.* arXiv:2512.06292.

[BGT] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular Variation.* Cambridge University Press, 1987; Theorem 1.2.1. Book.