Almost-sure parameter continuity of the LQG metric throughout the positive-Q regime

Research proof prepared for Leo

13 November 2026

Abstract

For every integer $d \geq 2$, the canonical metrics for a common log-correlated Gaussian field admit an almost-surely continuous modification in the Weyl parameter on $\{\psi > 0 : Q_d(\psi) > 0\}$, with values in the space of lower semicontinuous extended metrics. The fixed-parameter construction and uniqueness in the supplied positive-Q manuscript are used as inputs. The new geometric argument replaces finite boundary diameters by simultaneous endpoint chains, controlled finite trees, and separate source-metric companions of uniformly bounded mass. A discount supported in a thin neighborhood of such a tree supplies either a nearly optimal target shortcut or a cheap source restoration. This recovers the finite Gaussian forcing count in every $d \geq 3$. All precision losses are bounded by a tower of a fixed height; a parameter-grid argument then gives almost-sure, rather than merely stochastic, continuity. The topology and the modification quantifier are part of the statement.

Contents

1 Statement, topology, and previous inputs **3**
1.1 The quantitative statement . . . 3
1.2 One-metric inputs and their quantitative form . . . 4

2 Uniform estimates on a positive-Q compact interval **5**

3 Simultaneous endpoint chains and controlled contacts **7**

4 A rough comparison at finite depth **10**

5 Quantitative protected carriers in dimension at least three **12**
5.1 Minimum trees and local quasigeodesic control . . . 12
5.2 Quantitative holes and two leads . . . 14
5.3 The protected carrier . . . 16

6 Separate companions, local wells, and comparison **18**
6.1 A separate source companion . . . 18
6.2 Repairing endpoints at controlled contacts . . . 19
6.3 The pure tube well . . . 20
6.4 The finite Gaussian forcing count . . . 22
6.5 Local contact sewing and annular amplification . . . 26
6.6 Completion of the reciprocal estimate . . . 29

7 Completion of the parameter-continuity argument **30**
7.1 A complete space for lower semicontinuous convergence 30
7.2 Explicit grids and the continuous modification 31
7.3 Backward recovery, positivity, and triangle inequalities 32
7.4 Continuity at deterministic points 32
7.5 Uniform Euclidean confinement and geodesics 33
7.6 Median normalization and the modification quantifier 34

Sources and status of the inputs **34**

1 Statement, topology, and previous inputs

Fix an integer $d \geq 2$. Let h be the whole-space log-correlated Gaussian field whose logarithmic covariance has coefficient one, with unit-sphere average zero. Write

$$\mathcal{I}_d = \{\psi > 0 : Q_d(\psi) > 0\}, \qquad q_d(\psi) = \psi Q_d(\psi).$$

The fixed-parameter canonical metric and its strong characterization are those in the supplied positive-Q construction [1]; for $d = 2$ use the canonical planar metric. Thus locality, continuous Weyl scaling, exact similarities, Euclidean continuity of finite metric curves, finite-pair geodesics, and deterministic-point finiteness are available. These are inputs, not conclusions inferred from uniqueness alone. The argument below extends the parameter-continuity results [2, 3].

An extended metric is symmetric, separates points, satisfies the triangle inequality, and may take the value $+\infty$ off the diagonal. For nonnegative lower semicontinuous functions on $\mathbb{R}^d \times \mathbb{R}^d$, write $F_n \to F$ in the *lsc topology* if

(i) $F(x, y) \leq \liminf_n F_n(x_n, y_n)$ whenever $(x_n, y_n) \to (x, y)$;
(ii) for each (x, y) there are $(x_n, y_n) \to (x, y)$ with $F_n(x_n, y_n) \to F(x, y)$.

This is the topology used in the previous planar supercritical manuscript. Let $\mathcal{M}_d$ denote the lsc extended metrics with this topology.

Theorem 1.1. *For every $d \geq 2$, the canonical median-normalized family*

$$\operatorname{med} D_h^\psi(0, e_1) = 1, \qquad \psi \in \mathcal{I}_d,$$

has a jointly measurable modification $\{\widetilde{D}_h^\psi\}_{\psi \in \mathcal{I}_d}$ such that, almost surely,

$$\mathcal{I}_d \ni \psi \longmapsto \widetilde{D}_h^\psi \in \mathcal{M}_d \quad \text{is continuous.}$$

Here modification means that $\widetilde{D}_h^\psi = D_h^\psi$ almost surely for every fixed deterministic ψ. On the same event every value is an extended length metric, finite pairs admit Euclidean-continuous geodesics, and finite-radius metric balls are Euclidean compact.

For any prescribed deterministic countable set $S \subset \mathbb{R}^d$, the modification can additionally be chosen so that every $\psi \mapsto \widetilde{D}_h^\psi(x, y)$, $x, y \in S$, is finite and continuous. In particular, for each prescribed deterministic sequence $\psi_n \to \psi \in \mathcal{I}_d$, the originally specified versions satisfy $D_h^{\psi_n} \to D_h^\psi$ almost surely in the lsc topology.

The countable-set assertion does not permit the endpoints to be chosen arbitrarily after the field. Locally uniform convergence of all distances is not asserted in the singular regime. On a subcritical interval the result is compatible, by agreement on a dense parameter set, with the stronger locally uniform versions already constructed in [3] and its planar counterpart. If the Euclidean endpoint $\psi = 0$ is included, its treatment is the existing subcritical endpoint argument; the present new argument concerns positive ψ. No one-dimensional assertion is being made.

1.1 The quantitative statement

Work first in mean-log normalization, and denote the resulting metrics again by D_h^ψ:

$$D_h^\psi \;\leftarrow\; \exp\{-\mathbb{E} \log D_h^\psi(0, e_1)\} D_h^\psi.$$

The uniform moments below make this a finite positive deterministic rescaling. No continuity of this rescaling is assumed. For a metric D, set

$$D_r(x, y) = \inf\{D(x', y') : x' \in \overline{B}_r(x),\ y' \in \overline{B}_r(y)\}.$$

Using closed instead of open endpoint balls only changes harmless constant margins. Define $\mathsf{T}_0(t) = t$ and $\mathsf{T}_{j+1}(t) = \exp(\mathsf{T}_j(t))$.

Theorem 1.2 (Reciprocal comparison)**.** *Fix $I = [a, b] \Subset \mathcal{I}_d$. There are constants $C \geq 1$ and an integer $H \geq 0$, depending only on d, I, such that the following holds. Put*

$$A(\delta) = \max\{2, \mathsf{T}_H(\delta^{-C})\}.$$

For $\psi, \eta \in I$, $0 < \delta < \delta_0$, $R \geq 1$, and an integer n satisfying

$$|\psi - \eta| \leq A(\delta)^{-1}, \qquad n \geq A(\delta)(1 + \log(R + 1)), \qquad |\psi - \eta| n \leq A(\delta)^{-1}, \tag{1}$$

let $r = \exp[-n/A(\delta)]$. Except on an event of probability at most

$$A(\delta)(R + 1)^{A(\delta)} \exp[-n/A(\delta)], \tag{2}$$

one has simultaneously for all $x, y \in B_R$

$$(D_h^\eta)_r(x, y) \leq (1 + \delta) D_h^\psi(x, y) + r, \qquad (D_h^\psi)_r(x, y) \leq (1 + \delta) D_h^\eta(x, y) + r. \tag{3}$$

This deliberately weak form absorbs all finite-height quantitative losses into one function. Its height H is fixed; it does not depend on δ, the number of field modifications, or the number of comparison improvements. The finite-height assertion is essential for the parameter chaining. A qualitative threshold for each δ would not prove Theorem 1.1.

For $d = 2$, Theorem 1.2 follows from the stronger comparison in [2, Proposition 1.2]. We prove its new $d \geq 3$ extension. The analytic forcing count is adapted from [3, Section 4]; the replacement of its finite sphere diameters is proved here.

1.2 One-metric inputs and their quantitative form

The following facts are used from [1]. All domains have fixed positive buffers. Metric statements about arms are simultaneous over arms chosen after the field. They concern actual intrinsic metrics, not just finite-cutoff certificates.

(P1) For deterministic point pairs in fixed connected buffered geometries, normalized internal distances have a uniformly bounded positive p-moment for some $p > 0$. A normalized annular crossing X obeys

$$\mathbb{P}[X < e^{-t}] \leq C \exp\{-ct^2/\log(t + 2)\}, \qquad t \geq 1. \tag{4}$$

See Sections 6–12 and Theorem 22.2 there.

(P2) For each $\varepsilon > 0$, a finite connected core supplies actual selected contacts on arbitrary crossing arms. Their attachment costs are at most ε times their own arm lengths, plus the core cost when needed. In normalized units the core mass has a positive moment bounded by $C\varepsilon^{-C}$. A prescribed high success probability can be obtained without changing the outer annular geometry. The starting fine depth is $O(\log \varepsilon^{-1})$; see Lemma 9.2 and Sections 17–18, 23. In particular two same-metric arms can be joined at cost at most C times their combined lengths on an arbitrarily likely local event.

(P3) The relative-core contraction and the thin-shell envelope apply to geodesics, bounded-weight geodesics, and K-quasigeodesics, including local versions with arbitrary nonnegative chronological mass outside the inspection. For fixed K, their normalized local span has a positive moment. Its quantiles have power bounds in K and in the reciprocal failure probability. Interior occupation exponents can be taken at least c/K^C. These conclusions follow from the recurrences displayed below, using Sections 21, 31 and 46–47 there.

(P4) Separated annular Gaussian restrictions modulo constants satisfy a block Hölder inequality with a fixed exponent P_0. If the marginal success probability is $1 - p$, all-point positive-density annular covers at logarithmic depth m fail with probability Ce^{-cm}; c can be made arbitrarily large by decreasing p. The spatial cover uses a deterministic fine grid. See Section 4 and the off-center promotion in Sections 53–54 there.

Every probability statement above initially concerns a deterministic parameter. Compact-parameter uniformity is established next. All later statements over an uncountable family of paths are deterministic consequences of these simultaneous certificates; they are not uncountable intersections of probability-one events.

2 Uniform estimates on a positive-Q compact interval

Proposition 2.1. *The function $q_d(\psi) = \psi Q_d(\psi)$ is locally Lipschitz on $(0,\infty)$. On $I = [a,b] \Subset \mathcal{I}_d$, the inputs (P1)–(P4), fixed high point quantiles, compact inverse spatial bounds, and small-inverse-moment large-scale barriers can be chosen uniformly in deterministic $\psi \in I$. When an inverse-crossing moment of order $s \geq 1$ is needed, one can use the effective bound*

$$\log \mathbb{E}[X^{-s}] \leq C(1 + s^2 \log(s+2)). \tag{5}$$

In particular the following useful forms hold, with positive fixed exponents and constants uniform on I:

- *(i) on each fixed compact region there is a random A_ψ with $\mathbb{P}[A_\psi > L] \leq CL^{-p}$ and $D_h^\psi(x,y) \geq A_\psi^{-1}(|x-y| \wedge 1)^b$;*
- *(ii) for a sufficiently small fixed $\alpha > 0$, except on an event of probability Cu^q, every center in a fixed compact region has at least $c\log(1/u)$ disjoint good same-metric annuli between $2u$ and u^α;*
- *(iii) for deterministic z and small r, $\mathbb{E}[D_h^\psi(z, z+re_1)^p] \leq C_z r^\kappa$, with the internal version in a corresponding buffered ball;*
- *(iv) if $X_\psi(R)$ is the crossing charge of a fixed-shape annulus at radius $R \geq 2$, then $\mathbb{P}[X_\psi(R) < L] \leq CL^q R^{-p}$.*

Any fixed desired polynomial failure exponent in the second assertion is obtained by increasing the local core references before fixing α. More generally an all-point cover providing J good annuli between radii r, R on a guard of diameter L is available once

$$\log(R/r) \geq C(J + \log(1/R) + \log(L+2) + \log(p_0^{-1})),$$

with failure at most p_0, after the marginal success probability has been fixed sufficiently high. This displays the spatial-cover entropy as well as the required number of good annuli.

Proof. Let $L_n(\psi)$ be one fixed buffered LFPP point or crossing distance at cutoff n. Coupling all parameters with the same field gives

$$|\log L_n(\psi) - \log L_n(\eta)| \leq |\psi - \eta| \sup_V |h_n|, \qquad \mathbb{E}\sup_V |h_n| \leq C(n+1).$$

The functions $-(n\log 2)^{-1}\mathbb{E}\log L_n(\psi)$ are therefore equi-Lipschitz on bounded intervals. The Cameron–Martin Lipschitz constant of $\log L_n(\psi)$ is at most $C\psi\sqrt{n+1}$, because each cutoff evaluation has Gaussian norm at most $C\sqrt{n+1}$. Gaussian concentration implies that its mean and any fixed nontrivial quantile differ by $O(\sqrt{n})$, uniformly on I. The exponent theorem identifies the pointwise limit as $1 - q_d$. An equi-Lipschitz finite-net argument gives uniform convergence on compact intervals and the asserted Lipschitz property. In particular $\mathcal{I}_d$ is open.

Put $q_* = \min_I q_d > 0$ and

$$\kappa = \min\{q_*/4, 1/4\}, \qquad \alpha(\psi) = Q_d(\psi) - \kappa/\psi.$$

Then $\alpha(\psi) \geq 3q_*/(4b)$, $\psi(Q_d - \alpha) = \kappa$, and $1 + \psi\alpha - q_d = 1 - \kappa \geq 3/4$. Choose one small $p > 0$, for example with $2pb^2 < q_*/4$. After choosing one common small slack, both point-chain moment exponents of orders $p, 2p$ are uniformly positive. These are the strict margins used in the exponent, core, and normalizer arguments.

For completeness the normalizer comparison must be made uniformly, not deduced from separate fixed-parameter bounds. Proposition 5.11 of [1] is uniform on I: its proof uses the fixed finite-range lower-quantile labels, the same Gaussian coarse oscillation bound, and the now uniform annular exponent convergence. Hence its large-gap threshold is common for every fixed slack. The envelope ratios and small positive moments in Lemma 8.1 are consequently uniform as well. If the final point/annular normalizer ratios were unbounded, choose varying (ψ_j, n_j) witnessing this. Bounded cutoffs have bounded ratios uniformly on I, so $n_j \to \infty$. Apply the finite-record selection separately to each prefix at ψ_j. Uniform exponent convergence gives the common $o(n_j)$ record growth. Fix, in the order of Sections 8–11 of [1], the common positive margins above, the envelope and attachment slacks, a fixed macroscopic endpoint radius, and finally the penalty coefficient. Every displayed estimate in the record contradiction then has a common constant. The contradiction is a uniform positive lower bound for a normalized point p-moment, while that point variable tends to zero in probability and has a common $2p$-moment. Uniform integrability is contradictory. This is exactly the varying-fractional-cutoff argument of Section 12.2, with ψ_j varying as well. It proves the common scalar comparison.

The remaining estimates now follow from the same local proofs with these fixed margins and quantiles. Integrating (4), split at $Cs\log(s+2)$; above that point $ct^2/\log(t+2)$ dominates $2st$. This gives (5). At depth k, a union over $C(1+k)^C 2^{dk}$ grid shells gives

$$\mathbb{P}[\text{some normalized crossing} < 2^{-\varepsilon k}] \le C(1+k)^C \exp\{C_d k - c\varepsilon^2 k^2/\log(\varepsilon k+2)\}. \tag{6}$$

The corresponding coarse oscillation estimate has the same form. Thus a small shell slack introduces only a power of ε^{-1} and logarithmic factors into the starting depth, not an unspecified quantile. The core accuracy dependence follows from its logarithmic starting depth and the exponential moment growth at that depth. The first displayed consequence is the all-grid lower estimate with a fixed slack, summed over scales. The second is the block Hölder estimate and the spatial-grid union bound; choose the outer exponent α small to pay its entropy. The third follows from exact scaling and Cauchy–Schwarz, using the common $2p$ point moment and decreasing p so that the Gaussian scale factor leaves a positive power of r. The same calculation with a small negative moment, at large radii, proves the fourth since $q_* > 0$. □

Here is the K dependence used repeatedly later. A thin-shell core with accuracy ε gives, for consecutive local spans,

$$T_j \le K\varepsilon(T_{j+1} - T_j) + KH_{j,\varepsilon}, \qquad \mathbb{E}[(H_{j,\varepsilon}/s)^p] \le C\varepsilon^{-A}e^{Bj}.$$

Choose $\varepsilon = c_0/K$ so that $q = c_0/(1+c_0)$ satisfies $q^p e^B < 1$. Iteration gives

$$T_0/s \le \sum_{j\ge 0} q^j K H_{j,\varepsilon}/s, \qquad \mathbb{E}[(T_0/s)^p] \le CK^{p+A}.$$

The terminal term vanishes for a finite chronological measure. Relative connectors without an additive mass similarly give $T_{\text{in}} \le KC(T_{\text{out}} - T_{\text{in}})$ on successful annuli. A positive density of these annuli yields an occupation exponent at least c/K^C. These arguments also show directly why the references are uniform over all allowed weights, rather than over one preselected path.

We use an elementary convention for quantitative bookkeeping: a quantity is *controlled* if it and its required reciprocal are bounded by $\mathsf{T}_j(v^{-C})$ for fixed j, C. A fixed finite composition of sums, products, powers, logarithms and exponentials of controlled quantities is controlled. Every construction below has a fixed number of such stages. Operations on many slots or many comparison coefficients are performed in parallel or with the same preassigned constants; they do not increase the composition depth.

3 Simultaneous endpoint chains and controlled contacts

The distinction between a finite curve and arbitrary points of a Euclidean neighborhood matters here. We prove access at the actual contacts needed later. No finite diameter is assigned to an open ball.

Lemma 3.1 (Continuum endpoint chains)**.** *In a fixed buffered compact region, write $h = H_k + g_k$, where H_k contains the coarse scales and g_k the scales below 2^{-k}. Uniformly for deterministic $\psi \in I$, there are events G_M with*

$$\mathbb{P}(G_M^c) \le \sum_{k \ge M-1} C 2^{dk} e^{-ck^2} \tag{7}$$

such that, simultaneously for every endpoint x, its scale-k chain cost is bounded by

$$U_k^\psi(x) = C(1+k)^{2d} 2^{-kq_d(\psi)} \exp\left(\psi \sup_{B_{C(1+k)^2 2^{-k}}(x)} H_k \right). \tag{8}$$

If the two corresponding series converge and $|x-y| \le c2^{-M}$, then

$$D_h^\psi(x, y; B_{C(1+M)^2 2^{-M}}(x)) \le C \sum_{k \ge M} (U_k^\psi(x) + U_k^\psi(y)). \tag{9}$$

Every point with $\limsup_k H_k(x)/(k \log 2) < Q_d(\psi)$ belongs to the finite component of the rational points. A common strict thickness bound gives a uniform internal upper Hölder bound on the set of points satisfying that bound.

Proof. This is the continuum version of [1, Lemma 6.2]; using exact continuum scaling avoids a cutoff-ratio constant depending on a small exponent slack. On $2^{-k}\mathbb{Z}^d$, call a vertex good when all the finitely many signed coordinate-edge distances for $D_{g_k}^\psi$ in their fixed buffers are at most $K_0 2^{-kq_d(\psi)}$. The uniform positive moment gives arbitrarily small bad probability for a fixed large K_0. Strong locality and the compact-kernel representation give a fixed finite dependence range in lattice units. Call a vertex nice if it is good and all the needed level-$(k+1)$ vertices in its fixed neighboring region are good.

A bad self-avoiding $*$-path of j sites contains at least $c_d j$ sites with disjoint determining regions. Summing its probability over at most $(3^d)^j$ choices gives exponential bad-component tails. Except with probability $C2^{dk}e^{-ck^2}$, all relevant bad and non-nice components have diameter $C(1+k)^2$ in lattice units. The deterministic lattice routing lemma connects nearby infinite-good component vertices with at most $C(1+k)^{2d}$ edges in a $C(1+k)^2$ neighborhood. Nice vertices connect consecutive levels; a coarse coordinate edge subdivides into two good fine edges. Thus every x, chosen after the field if desired, has good vertices $v_k \to x$, with $|v_k - x| \le C(1+k)^2 2^{-k}$, connected in successive shrinking tubes. Weyl restoration of H_k gives (8). The same-level lattice bridge joins the two chains of nearby endpoints. A union over levels proves (7).

If $\sum_k U_k^\psi(x) < \infty$, lower semicontinuity gives

$$D_h^\psi(v_j, x) \le \liminf_{l \to \infty} D_h^\psi(v_j, v_l) \le C \sum_{k \ge j} U_k^\psi(x).$$

Consequently the actual successive paths can be concatenated and appended at x, continuously in the metric, with finite total length. They stay in the prescribed Euclidean tube. This proves (9). Their lattice vertices are rational and belong to the canonical finite component.

For every $e > 0$, outside summable probabilities bounded by

$$C \exp\{Ck - ce^2 k^2 / \log(k+2)\},$$

the oscillation of H_k over all the indicated tubes is at most ek, simultaneously in their centers. If $H_k(x) \le \alpha k \log 2 + B$ for all $k \ge M$, with $\alpha < Q_d(\psi)$, choose e smaller than a fixed multiple of $Q_d(\psi) - \alpha$. The series in (9) is geometric after a polynomial factor. It is bounded by $Ce^{\psi B}2^{-bM}$ for any fixed $0 < b < \psi(Q_d(\psi) - \alpha)/2$. This proves the assertions. The initial depth for a prescribed failure probability has a power bound in the reciprocal gap, by the displayed Gaussian estimate. All endpoint quantifiers are on the one good lattice event. □

A short reference germ from such an x is obtained by adjoining to its chain a same-level good-lattice bridge to a good vertex in a fixed offset box. To keep it inside a ball of radius r, use lattice mesh of order $r/\log^2(r^{-1})$. It then reaches outside a fixed smaller multiple of that ball and has cost at most Cr^b, after decreasing b. This supplies reference germs, not cheap access to every preassigned nearby rational point.

Lemma 3.2 (Transfer of controlled geodesic contacts)**.** *Fix $K \ge 1$. On an arbitrarily high-probability local event, the interior contacts of every locally K-quasigeodesic crossing a fixed buffered region have a common thickness bound*

$$H_k(x) \le \left(Q_d(\psi) - \frac{\beta_K}{2\psi}\right) k \log 2 + B_K, \qquad k \ge k_K, \tag{10}$$

where $\beta_K \ge cK^{-C}$. The constants and the onset have controlled bounds in K and in the reciprocal failure probability. If $|\eta - \psi| \le c_I \beta_K$, the contacts have uniform short internal reference germs and uniform internal two-point upper Hölder bounds for both D_h^{ψ} and D_h^{η}.

The same conclusion holds for a whole controlled path when both endpoints have uniform short reference germs. In particular it applies to the finite carriers constructed below. For output geodesics under arbitrary weights in $[\theta, 1]$, use $K = \theta^{-1}$.

Proof. The local span envelope and occupation contraction following Proposition 2.1 give, simultaneously for all such curves, $\mathrm{len}_D(P \text{ between visits to } B_{2r}(x)) \le A_K r^{\beta_K}$ at an interior contact; the prescribed buffers keep its endpoints outside the local annuli. The all-grid lower-shell estimate with slack $e \le \beta_K/2$ gives

$$cr^{q_d(\psi)+e} e^{\psi H_{\log_2(1/r)}(x)} \le D_h^{\psi}(\text{across } A_{r,2r}(x)) \le A_K r^{\beta_K}.$$

Taking logarithms proves (10), after adjusting B_K and decreasing the exponent. The effective onset follows from (6). Local Lipschitz continuity of Q_d away from zero implies $Q_d(\eta) \ge Q_d(\psi) - \beta_K/(4\psi)$ when the parameter spacing is sufficiently small. Lemma 3.1 gives the claimed access, with a common exponent $\zeta_K > 0$ and constant H_K, in both metrics.

We spell out the endpoint issue because pointwise interior bounds do not give a uniform estimate on an entire arc. Suppose an endpoint x has reference germs of cost Cr^{ζ}. For a K-quasigeodesic issuing from it, an asymmetric same-metric core connects the arm after the last inner visit to such a reference germ. If $T_{\mathrm{in}}, T_{\mathrm{out}}$ are prefix lengths through the last visits,

$$T_{\mathrm{in}} \le K\varepsilon(T_{\mathrm{out}} - T_{\mathrm{in}}) + CKK_\varepsilon r^{\zeta}.$$

Choose $\varepsilon = 1/(2K)$. The core coefficient K_ε has a power bound in K. Iterating over the positive-density successful annuli gives a uniform endpoint last-visit bound $C_K r^{\zeta'}$ for some effective $\zeta' > 0$. Reverse at the other endpoint. For a ball of radius r at any point of the arc, either an endpoint is within $\sqrt{r}$, when this bound applies, or both are farther away, when the interior contraction applies between r and $\sqrt{r}$. This gives one occupation bound on the whole arc and hence one strict thickness bound. There is no assertion of rectifiability of the original curve in the other parameter metric. □

Lemma 3.3 (Repairing displaced endpoints)**.** *Let F be one of the two metrics. Suppose x, y have uniform short F-reference germs and internal upper access, with cost HR^{ζ} at scale R. Suppose between r and R there are J disjoint good same-F annuli at both endpoints. A finite F-path P between their r-balls can be repaired, inside the corresponding guard, at cost*

$$F(x, y) \le (1 + 2C/J)\,\mathrm{len}_F(P) + CHR^{\zeta}. \tag{11}$$

The statement is simultaneous over all these contacts and paths. For an all-point assertion with failure at most p_0 on a guard of diameter L, it suffices to choose

$$\log(R/r) \ge C\big(J + \log(1/R) + \log(L + 2) + \log(p_0^{-1})\big).$$

The additional terms pay the spatial-grid entropy.

Proof. If $|x - y| \le 4R$, use the direct upper access. Otherwise the recovery path and the reference germ cross each of the J annuli at an endpoint. Their arm lengths sum to at most their respective total lengths. At one annulus a relative core joins them at cost at most C/J times that sum. Truncate the recovery path there and append the appropriate reference prefix. Repeat at the other endpoint. The total is (11); images stay in the guard by the buffers. Selected contacts, rather than arbitrary points of a ball, are all that is used. □

Lemma 3.4 (A robust favorable whole arc)**.** *For $0 < v < v_0$, in the two mean-log normalizations there is a local event of probability at least cv^2 carrying a whole E-geodesic Γ between two fixed rational points, and numbers $a, \rho > 0$, such that*

$$\mathrm{len}_E(\Gamma) \le (1 + Cv)a, \qquad D(\overline{B}_\rho(u), \overline{B}_\rho(v)) > a. \tag{12}$$

The source lower test is certified in a fixed enlarged buffer against all exterior extensions. The confinement radius, moduli, positive distance bounds and ρ^{-1} have controlled bounds; before local rescaling these can be taken as powers of v^{-1}. In particular no uncontrolled random subarc endpoints are used to obtain the initial favorable ratio.

Proof. Set $Y = \log E(0, e_1) - \log D(0, e_1)$. Its mean is zero and its variance is uniformly bounded, by the point moments and crossing inverse tails. Cantelli's inequality gives

$$\mathbb{P}[Y \le \log(1 + v)] \ge cv^2.$$

Restrict point-distance upper and lower quantiles, rational endpoint access, and geodesic moduli with total failure at most $cv^2/4$. All bounds are powers of v^{-1} by fixed positive moments. At deterministic endpoints one has a sharper version of the repair lemma. Attach a crossing arm to a finite same-D core with accuracy ε, then to a rational anchor, and add a reference path from the deterministic endpoint u to that anchor. The excess cost $Z(u, r, \varepsilon)$ satisfies, for one sufficiently small fixed $p > 0$,

$$\mathbb{E}[Z(u, r, \varepsilon)^p] \le C\varepsilon^{-A} r^{\kappa}, \qquad \kappa > 0.$$

Indeed its normalized excess has a $2p$-moment bounded by a power of ε^{-1}: the core starts at depth $O(\log \varepsilon^{-1})$, and the rational reference paths have a small positive moment. Exact scaling and Cauchy–Schwarz bound the physical moment by $C\varepsilon^{-A} r^{p\psi Q_d(\psi) - p^2\psi^2}$, which has a uniform positive exponent for small p. No independence of the coarse factor and excess is needed. Applying the two anchored attachments to disjoint initial and final arms of a recovery path yields

$$D(u, v) \le (1 + \varepsilon) D_{r/4}(u, v) + Z(u, r, \varepsilon) + Z(v, r, \varepsilon).$$

Choose ε a sufficiently small multiple of v, and r a sufficiently high fixed power of v. Markov's inequality makes the additive error $o(v)$ times the point-distance lower cap, with failure $o(v^2)$. This proves the claimed power precision and (12), after decreasing the lower threshold slightly.

Finally cap both $D(0,e_1)$ and $E(0,e_1)$ by a power of v^{-1}. At a polynomially large radius require both respective outer crossing charges to exceed these caps. Uniform small inverse moments and $q_*>0$ make the additional failure $o(v^2)$. The E geodesic stays in that buffer; every competing D path leaving it pays the source exit charge. Hence the global relaxed source lower bound is certified by intrinsic data and the local exit barrier, and survives arbitrary changes outside. Translate and rescale the fixed template as needed. Thresholds may be rounded down on a $(1+cv)$-mesh, retaining the displayed slack. □

Locality convention. In annular applications retain the resulting inequalities for actual intrinsic metrics and locally admissible curves. They depend only on the field restriction in the fixed buffered annulus. Their probability is proved using the scale decomposition; the raw decomposition is not itself asserted to be measurable in the summed local field. Curve admissibility is encoded by rational-time source-length and local quasigeodesic tests. Failure of a universal contact inequality is the projection of a Borel set of local data, curves and contact times. Thus these events are measurable in the completed local Gaussian sigma field. The finite catalogs below use the same convention.

4 A rough comparison at finite depth

We first obtain a fixed comparison constant. All paths produced here have actual endpoints, displaced by a controlled Euclidean distance from the requested endpoints. The source path is never assumed to have finite length in the other metric.

Proposition 4.1 (Finite-depth prefix router)**.** *There are constants $C_*, c, a>0$, uniform on I, with the following property. For two buffered compact regions $K \Subset W$, for $n \ge n_0(K,W)$ and $|\psi-\eta|n \le c$, outside probability Ce^{-cn} every finite $D=D_h^{\psi}$ path $\Gamma\subset K$ of length at most e^{an} has an $E=D_h^{\eta}$ recovery path in W such that*

$$E^W(B_r(\Gamma(0)),B_r(\Gamma(1))) \le C_*\,\mathrm{len}_D(\Gamma), \qquad r=Ce^{-cn}. \tag{13}$$

The reciprocal statement holds on the same event. The result is simultaneous over all paths. A prescribed fixed exponential failure rate can be obtained by increasing C_* and the onset. For a bounded smooth shift f, the same construction applies when, in addition,

$$n \ge C(\|f\|_\infty + \log(2+\mathrm{Lip}(f))), \qquad |\psi-\eta|\|f\|_\infty \le c.$$

Here E^W denotes the constrained metric on a closed buffered guard; the construction stays a positive distance from its outer boundary. For the ball guards $K=\overline{B}_L$, $W=\overline{B}_{L+1}$, $L\ge 1$, one may take $n_0 \le C(1+\log(L+1))$ and exceptional probability $C(L+1)^C e^{-cn}$, with constants independent of L.

Proof. Choose a closed guard W' with $K \Subset \mathrm{int}\, W'$ and $W' \Subset W$, and run all constrained target prefixes in W'. The final recovery is therefore a positive distance from ∂W. We give both the local input and the deterministic routing argument. The asymmetric all-arm core in (P2) extends to the pair (D,E) at nearby parameters. At relative fine depth ℓ, the new upper/lower unit ratio, compared with the equal-parameter proof, is bounded on the simultaneous coarse cap by

$$\exp\{C|\psi-\eta|\ell\}. \tag{14}$$

More explicitly the target-chain/source-shell ratio is bounded by

$$C(1+\ell)^C 2^{-\ell(q_d(\eta)-q_d(\psi))}\exp\{(\eta-\psi)H_\ell+\eta\,\mathrm{osc}\,H_\ell+\epsilon_0\ell\}.$$

Here the oscillation and ϵ_0 are the fixed small slacks in the section estimate of [1, Lemma 9.1]. The exponent contribution is $2^{-\ell(q_d(\eta)-q_d(\psi))}$, and the coarse contribution is $\exp((\eta-\psi)H_\ell)$, with $|H_\ell|\le C\ell$. The equal-parameter section construction has a factor $2^{-\ell}$ before its fixed small exponent slacks. Fixing those slacks first, a sufficiently small fixed parameter difference leaves a summable geometric gain after (14). Thus the same finite core, fine-chain attachment and lsc passage in [1, Sections 17–18 and 23] yield, on an arbitrarily likely annular event, actual selected contacts for every pair of crossing arms A_D, A_E, with

$$E(\text{connector})\le \varepsilon\,\mathrm{len}_E(A_E)+K_0\,\mathrm{len}_D(A_D). \tag{15}$$

The target accuracy $\varepsilon>0$ is fixed first by increasing the starting relative depth; K_0 is then fixed. At physical scale s, the additional unit ratio is

$$s^{q_d(\eta)-q_d(\psi)}\exp((\eta-\psi)H_{\log_2(1/s)}).$$

For scales through depth n it is bounded above and below by fixed constants when $|\psi-\eta|n\le c$. All core paths and attachments concern intrinsic annular metrics.

We also need a seed core with a separated anchor. For a test crossing from radius $3s$ to $4s$, use the working annulus from s to $6s$, and fix a rational anchor q at radius $3s/2$ from its grid center. Truncate the first successful core depth at a deterministic level, losing an arbitrarily small probability, and retain the finite menu of all possible rational core vertices through that level. Augment the selected finite E core by internal rational-point paths to q. The finite geometries in that menu have uniform positive moments; high quantiles and a source crossing lower quantile therefore give, for every D arm, an actual contact joined to q at E-cost at most $K_0\,\mathrm{len}_D(A_D)$. This enlargement preserves annular locality.

Take good seed annuli at radii between $e^{-n/4}$ and $e^{-n/8}$, and good routing annuli between e^{-n} and $e^{-n/2}$. The all-point annular cover, using ancestor grid centers within a small fixed fraction of the corresponding radius, supplies at least κn pairwise scale-separated good annuli in each block. Increase the separation of successive levels so that each inner routing ball contains the next outer routing ball, even with these ancestor centers. The determining annuli about these centers lie in enlarged fixed annuli about the tested point. Increasing the separation preserves the block Hölder inequality. Choose the marginal failure probability sufficiently small to absorb the spatial-grid entropy. Together with the inverse spatial reference this has exceptional probability at most Ce^{-cn}. On this event, in W,

$$D(x,y),E(x,y)\ge e^{-an}(|x-y|\wedge 1)^b \tag{16}$$

for a fixed b, with a harmless change of a.

Parametrize $\Gamma:[0,L]\to K$ by D arclength, where $L\le e^{an}$. If its endpoints are within $Ce^{-n/8}$, a diagonal recovery pair suffices. Otherwise choose a seed annulus about $x=\Gamma(0)$ whose outer ball does not contain $y=\Gamma(L)$. Let t_- be the last visit to its closed radius-$3s$ ball and use the subsequent arm until the first radius-$4s$ exit. The seed produces a contact $\Gamma(t_0)$ and the anchor q with

$$E^{W'}(q,\Gamma(t_0))\le K_0L.$$

Every later source point is outside the radius-$3s$ ball, whereas q is at radius $3s/2$. Thus every target prefix from q to a later source point exits all the much smaller routing annuli about that point.

Here is the finite-depth balance used to extend a prefix. Given two finite germs from a point, one for D and one for E, which both exit the routing block, let a_i^F, b_i^F be their first

outer and inner exit lengths on the ith good annulus, ordered from outside to inside. Nesting gives $a_{i+1}^F \le b_i^F$. The last inner exit has length at least $ce^{-(a+b)n}$ by (16), and each whole germ has length at most $C_1 e^{an}$. Consequently

$$\sum_i \log \frac{a_i^D}{b_i^D} + \sum_i \log \frac{a_i^E}{b_i^E} \le C_2 n + 2\log C_1 + C_3.$$

Choose fixed $\theta_D, \theta_E > 0$ so small that fewer than $\kappa n/3$ good annuli can violate either corresponding bound. For large enough n, some good annulus satisfies

$$b^D \ge \theta_D a^D, \qquad b^E \ge \theta_E a^E. \tag{17}$$

Fix $\varepsilon < \theta_E/2$ in (15) and put $K_1 = K_0/\theta_D$.

The set

$$\mathcal{A} = \{t \in [t_0, L] : E^{W'}(q, \Gamma(t)) \le K_1 t + K_0 L\}$$

is nonempty and closed, by Euclidean continuity of Γ and lsc of the constrained metric. Let $t = \max \mathcal{A}$. Suppose $|\Gamma(t) - y| > Ce^{-n/2}$. The forward source germ and the reversal of a constrained E geodesic from q to $\Gamma(t)$ both exit every routing annulus. Their lengths obey the bounds used above. At an annulus satisfying (17), select the two arms after their last inner contacts and before first outer exits. The selected target contact removes a prefix tail of length $s_E \ge b^E$. The selected source contact is $\Gamma(t')$ with $t' - t \ge b^D$. Truncate the target prefix and append the mixed connector. Its length is at most

$$\begin{aligned} \mathrm{len}_E(P) - s_E + \varepsilon a^E + K_0 a^D &\le \mathrm{len}_E(P) - (1 - \varepsilon/\theta_E) s_E + (K_0/\theta_D)(t' - t) \\ &\le K_1 t' + K_0 L. \end{aligned}$$

If a constrained infimum needs approximation, its error is absorbed by the strictly positive removed-tail margin. Thus $t' \in \mathcal{A}$ with $t' > t$, a contradiction. The maximal prefix ends within $Ce^{-n/2}$ of y, starts within $Ce^{-n/8}$ of x, and has length at most $(K_1 + K_0)L$. This proves (13). The argument used only arclength of the source path, so it applies to every rectifiable source path, not just a geodesic. For the stated ball guards, the spatial-grid and coarse-field unions cost only a polynomial factor in $L + 1$. Their logarithmic entropy is $O(n + \log(L + 1))$. Taking $n \ge C(1 + \log(L + 1))$ gives precisely the uniform onset and failure form asserted above.

For smooth shifts, their oscillation at the routing scales is small as soon as $n \ge C \log(2 + \mathrm{Lip}(f))$. Their common constant component is included in both local units, whose ratio changes by at most $\exp(|\psi - \eta| \|f\|_\infty)$. The remaining onset absorbs $\|f\|_\infty$ in the height references. This proves the last assertion. □

5 Quantitative protected carriers in dimension at least three

Throughout this section $d \ge 3$, and the parameters belong to a fixed compact interval $J \Subset \{\psi > 0 : Q_d(\psi) > 0\}$. We write $D = D_h^\psi$ and $E = D_h^\eta$, with the mean-log normalization used in the comparison argument. All constants below are uniform on J. A fixed-height-tower bound means a bound $\mathsf{T}_H(Cv^{-C})$, where H, C are fixed, $\mathsf{T}_0(t) = t$, and $\mathsf{T}_{j+1}(t) = \exp(\mathsf{T}_j(t))$. The height never depends on a number of slots, trial boxes, or paths. Reference failures can be made any prescribed fixed small number before the accuracy v is chosen.

5.1 Minimum trees and local quasigeodesic control

Lemma 5.1 (Minimum finite network)**.** *Let F be a separating lower semicontinuous extended length metric with Euclidean compact finite-radius balls and Euclidean-continuous finite-pair*

geodesics. Given $N \geq 2$ distinct mutually finite terminals, there is an embedded finite geodesic tree T spanning them, with at most $N-2$ free vertices and $2N-3$ edges, such that

$$F(x,y) \leq d_T(x,y) \leq (2N-1)F(x,y), \qquad x,y \in T. \tag{18}$$

Here d_T is intrinsic tree arclength. The tree minimizes total edge length among abstract terminal-spanning trees whose free vertices have degree at least three; geometric intersections are not excluded from this variational class.

Proof. The degree identity bounds the number of free vertices by $N-2$, so there are finitely many combinatorial types. Fix a finite competitor of cost L. In every competitor of cost at most $L+1$, each free vertex is within F-distance $L+1$ of the first terminal. All its coordinates therefore lie in one Euclidean compact set. For each type the sum of edge distances is lower semicontinuous on a finite product of this set and attains its infimum. Contract zero-length edges, delete free leaves, and suppress free degree-two vertices. These operations cannot increase cost. A zero-edge contraction cannot identify distinct terminals, since F separates points. Repeat until every remaining edge has positive length, then realize it by a geodesic.

For two abstract positions x, y, subdivide their edges, adding at most two vertices. Their tree path has at most $2N-1$ segments. Add an edge of cost $F(x,y)$ between them and remove a longest segment, of length a, on the old path. The result is a tree. Pruning and suppression give an admissible competitor of cost at most $\operatorname{len}(T;F) - a + F(x,y)$; thus $a \leq F(x,y)$. Summing the segment lengths proves the upper bound in (18); the lower bound is the triangle inequality. If two abstract positions had the same image, the upper bound would force their tree distance to vanish. The realization is therefore injective. It is a topological embedding because the abstract tree is compact and its realization is Euclidean continuous. This also excludes coincidences between free vertices and unrelated terminal or edge points. □

The all-arm estimates of [1] give the following quantitative extension of the usual geodesic occupation estimate. We record the recurrence because its dependence on the quasigeodesic constant is needed later.

Lemma 5.2 (Local envelope)**.** *For $K \geq 1$, all finite K-quasigeodesics for E whose endpoints lie outside a fixed outer buffer have, simultaneously, a bounded first-to-last visit span in a fixed inner buffer. In local metric units s, a failure-u bound for this span is at most $CK^C u^{-C}$. On further local reference events, their interior traces have a uniform inverse Euclidean modulus with exponent at least c/K and constants polynomial in K, u^{-1}. The forward modulus comes from the usual lower spatial metric estimate.*

Proof. Take disjoint thin shells increasing to the outer buffer. The adjustable all-arm core construction supplies connectors of cost $\varepsilon(\operatorname{len} A^- + \operatorname{len} A^+) + H_{j,\varepsilon}$ in shell j, where, for one fixed $p \in (0,1]$,

$$\mathbb{E}[(H_{j,\varepsilon}/s)^p] \leq C\varepsilon^{-A} e^{Bj}. \tag{19}$$

Indeed the core depth is $O(\log \varepsilon^{-1})$, its fixed moments grow exponentially in that depth, and the thin-shell rational grid and Gaussian scale increments contribute only e^{Bj}. The exponent B is independent of ε; these are exactly the core and point-moment bounds of [1]. If T_j is the full visit span, the two selected contacts bracket the inner span and their arm lengths sum to at most $T_{j+1} - T_j$. Hence

$$T_j \leq K\varepsilon(T_{j+1} - T_j) + KH_{j,\varepsilon}.$$

Choose $\varepsilon = c_0/K$ so that $q = c_0/(1+c_0)$ satisfies $q^p e^B < 1$. Iteration, using $q^n T_n \leq q^n \operatorname{len}(P;E) \to 0$, gives

$$T_0 \leq \frac{K}{1+c_0} \sum_{j \geq 0} q^j H_{j,c_0/K}, \qquad \mathbb{E}[(T_0/s)^p] \leq CK^{p+A}.$$

Markov's inequality proves the local cap, without any bound on the parent's global length. At a successful annulus with relative connector constant C_0, the sharper recurrence is $T_{\mathrm{in}} \le KC_0(T_{\mathrm{out}} - T_{\mathrm{in}})$. Thus its contraction factor is $\rho_K = KC_0/(1 + KC_0)$. The all-point positive-density annulus estimate of [1], applied in scale blocks $[m, 2m]$, yields an inverse exponent smaller than $-c \log \rho_K \asymp K^{-1}$. Its onset has exponential tails. Combining it with the local cap proves the asserted normalized bounds. All events use internal metrics in fixed buffered determining regions. □

For fixed deterministic rational terminals, short endpoint chains complete this estimate up to the endpoints. Every pair of points of a finite tree without free leaves lies on a terminal-to-terminal path. Consequently, for a fixed K and a fixed mass bound, one may impose with arbitrarily high fixed probability

$$d_T(x, y) \le H_0|x - y|^{\beta_0}, \qquad \beta_0 > 0, \tag{20}$$

on the whole initial tree. The constants here are fixed before v.

5.2 Quantitative holes and two leads

Lemma 5.3 (Local routing)**.** *Let a local trace have a forward Hölder exponent $a > 0$ and an inverse-time Hölder exponent $\beta > 0$, with constant M, on fixed nested buffers. There is a finite deterministic rational catalog of boxes, of mesh bounded below by a power of $(1 + M)^{-1}$ with exponent $O((a\beta)^{-1})$, one of which misses the trace with that clearance. If a second controlled arc is contained in a missed box, two disjoint polygonal lead tubes can join specified separated endpoint windows of the two traces, while avoiding their complementary portions. Their mesh, edge count, and mutual clearance have explicit power bounds in the moduli, buffer widths, and window clearances. The statement holds for traces with arbitrary excursions outside the local working domain.*

Proof. Collapse the complement of a working domain Ω to one point, with quotient distance $d_\Omega(x, y) = \min\{|x-y|, \operatorname{dist}(x, \Omega^c) + \operatorname{dist}(y, \Omega^c)\}$. Decrease the Hölder exponents to $a, \beta \in (0, 1]$ if necessary, changing the constants on these bounded domains. Write the trace as $\alpha : [0, A] \to \Omega/\Omega^c$, and assume

$$d_\Omega(\alpha(s), \alpha(t)) \le M|s - t|^a, \qquad |s - t| \le M|\alpha(s) - \alpha(t)|^\beta$$

when both images lie in the relevant inner compact set. On any compact portion C of this trace its inverse-time map has an M-Lipschitz extension for the snowflake distance $|x - y|^\beta$. Clip the extension F to $[0, A]$. For x within ε of C, its nearest trace point z satisfies

$$d_\Omega(\alpha(F(x)), z) \le M^{1+a}\varepsilon^{a\beta}.$$

In a fixed interior buffer this is Euclidean, so $R(x) = \alpha(F(x))$ obeys

$$|R(x) - x| \le (1 + M^{1+a})\varepsilon^{a\beta}. \tag{21}$$

If a ball $\overline{B}_r$ were entirely within ε of the trace, with the right-hand side below $r/4$, degree would give $B_{3r/4} \subset R(\overline{B}_r)$. But $F(\overline{B}_r)$ is an interval on which α is injective, so this image is an arc or a point and has empty interior. A rational grid at the resulting explicit scale gives the hole catalog. The same argument applies after deleting open time windows: the extension may land on the full trace, which suffices.

For completeness, here is the quantitative nonseparation step. Choose connected buffers $O_0 \Subset O_1 \Subset O_2 \Subset O_3 \Subset \Omega$ and endpoints in O_0, at distance at least u from the forbidden closed trace portions C_i. The complete old and new traces are separated by $\sigma > 0$ in this

region. Take ε so small that (21) is less than $\min(u, b)/10$, where b is the smallest buffer width, and $\varepsilon < \sigma/10$. In S^d put $B = O_2^c \cup \bigcup_i C_i^{+2\varepsilon}$. On each disjoint tube use the homotopy $z \mapsto z + t\zeta(z)(R_i(z) - z)$, with $\zeta = 1$ near $\overline{O_1}$ and $\zeta = 0$ near O_2^c, and use the identity elsewhere on B. It avoids the endpoints and at time one maps B into $Z = O_0^c \cup C_*$, where C_* is a compact subset of the full traces with small endpoint balls removed. The finite closed sum theorem gives $\dim C_* \le 1$. The nonseparation theorem for sets of covering dimension at most $d - 2$ implies that $O_0 \setminus C_*$ is connected [4]. Naturality of Alexander duality, with Čech cohomology for these compact sets, transports this nonseparation through the homotopy and shows that the endpoints lie in the same component of $S^d \setminus B$ [5]. This open complement is path connected.

Round a resulting path to cubes of mesh $c\varepsilon$, erase repeated cubes, and connect their centers inside O_3. It has at most $C\varepsilon^{-d}$ edges and retains clearance $c\varepsilon$. Do this for both leads, with disjoint small endpoint balls. A small rational perturbation makes the two polygonal paths simple and mutually disjoint when $d \ge 3$. First insert a free midpoint on any edge whose two endpoints are fixed, and exclude collinear triples as well as unintended intersections. Each forbidden nonincident intersection is contained in a proper algebraic condition of bounded degree on their free vertices. A rational grid of a fixed polynomially finer mesh avoids the finite collection of such conditions. The determinant bounds then give a polynomial lower bound for segment separation, as in the elementary routing argument of [3]. This proves the quantitative claim and includes excursions collapsed outside Ω. □

Lemma 5.4 (Actual controlled leads)**.** *The tubes in Lemma 5.3 admit actual E-paths with effective finite cost, simultaneously for all admitted actual stubs. They may be chosen quasi-geodesic for E, with effective constants, while retaining their tube confinement and their permitted attachment windows.*

Proof. At a stub $x = \gamma(t)$ choose a rational grid point q within $R/10$. An actual stub prefix exiting $B_{3R}(q)$ has length at most $LM(4R)^\beta$ when γ is L-Lipschitz in its metric parametrization. A fixed rational path from q to $q + 3Re_1$ inside $B_{5R}(q)$ crosses the same band. The common core attaches both actual arms; thus

$$E(x, q; B_{5R}(q)) \le CLMR^\beta + s_R(q)(Y_{q,R} + CZ_{q,R}). \tag{22}$$

The normalized variables have a uniform small positive moment by [1]. Scale covariance and Gaussian moments give a crude physical bound $\mathbb{E}[(s_R(Y + CZ))^p] \le CR^{-C}$. A union bound over the finite grid therefore controls all actual stubs. The rational tube skeleton is handled by internal point paths in consecutive buffered balls. For n distinct links of radius at least r, the inequality $(\sum a_i)^p \le \sum a_i^p$ bounds their total cost by $(Cnu^{-1}r^{-C})^{1/p}$ outside probability u. Only distinct links, not complete route labels, are counted.

Let P be one resulting candidate, of cost at most L, in an inner tube. Give its enclosing collar a lower E crossing charge $\lambda > 0$, using the lower spatial estimate. Choose a smooth weight $w \in [1, B]$, equal to one on the inner candidate region and to B on that collar, where $B\lambda > L$. A wE-geodesic between the endpoints stays in the outer tube, has original length at most L, and every subarc satisfies

$$\mathrm{len}_E(\gamma|_{[s,t]}) \le (wE)(\gamma(s), \gamma(t)) \le BE(\gamma(s), \gamma(t)).$$

Use disjoint outer tubes for the two leads. Trim each after its last old trace contact and before its first mark contact. Its cost and quasigeodesicity persist, and the tube exclusions confine these contacts to the specified windows. No field resampling is used in this construction. □

5.3 The protected carrier

Proposition 5.5 (Quantitative protected carrier)**.** *Fix a finite rational terminal menu $\mathcal{P}$ in an inner collector, a small $\varepsilon_0 > 0$, and a reference failure $p_* > 0$. These data precede v. For $0 < v < v_0$ and sufficiently close parameters, with a fixed-height-tower sufficient restriction on $|\psi - \eta|$, a local event of probability at least $1 - Cp_*$ supplies an embedded finite E-rectifiable tree T in the collector containing $\mathcal{P}$, of mass at most M independent of v. It has finitely many protected middles with the following properties.*

1. *A favorable full arc Γ_j has rational trial endpoints u_j, w_j, a robust radius ρ_j, and a lower certificate $a_j > 0$ with $\operatorname{len}_E \Gamma_j \le (1 + Cv)a_j$ and $D(B_{\rho_j}(u_j), B_{\rho_j}(w_j)) \ge a_j$. These lower tests are local certificates valid for every exterior extension. The protected middle's actual endpoints c_j^-, c_j^+ lie inside the robust endpoint balls, with buffers from the lead attachment windows.*
2. *Every simple tree arc containing no complete protected middle has original E-length at most ε_0.*
3. *At either cut $c_j^{\pm}$, opposite components have quantitative separated remnants. In particular, if opposite-side points are sufficiently close, both lie on the same favorable mark and their tree distance through the cut is at most $H_v|x - y|^{\beta_*}$, for a fixed $\beta_* > 0$.*
4. *All points of T, including every lead endpoint, have one common strict field-thickness bound below $Q_d(\eta)$. Its gap is at least $\exp(-v^{-C})$, and its constants and onset have fixed-height-tower bounds. Lemma 3.2 therefore supplies uniform internal D/E access on the carrier for the stated nearby parameters.*

The number of marks is bounded independently of v; their positive geometric and metric scales are at least $\exp(-v^{-C})$. All finite geometric codes and subsequent precision choices have fixed-height-tower bounds. No global inverse modulus for the modified intrinsic tree metric is asserted.

Proof. Initial confinement. Wire $\mathcal{P}$ by finitely many rational internal E-paths in a smaller collector. Fixed point moments give a total cap L_0, and lower crossing tails give a positive exit charge b_0 for a fixed outer collar, outside probability p_*. Put an auxiliary smooth weight $w_0 \in [1, B_0]$, equal to one on the wiring region and B_0 on the collar, with $B_0 b_0 > L_0$. Apply Lemma 5.1 to $w_0 E$. Its minimizing tree cannot cross that collar. It has E-mass at most L_0 and is $K_0 = (2N-1)B_0$ quasigeodesic for E on every tree arc, since $d_T^E \le d_T^{w_0 E} \le (2N-1)w_0 E \le K_0 E$. Impose (20) on this tree T_0, with fixed H_0, β_0. This construction is local: the internal competitor and exit certificate prevent any minimizing network from using an exterior extension. All these constants are fixed before v.

Parallel safe slots. Set $\delta_0 = \varepsilon_0/[1000(N+1)]$. Subdivide original edges of length at least δ_0 into blocks of length in $[\delta_0/2, \delta_0]$ and take their midpoints as slots. There are at most $S \le C(N + L_0/\delta_0)$ slots. Intrinsic windows of radius $\delta_0/20$ are disjoint and avoid vertices. Equation (20) gives disjoint safe Euclidean balls of radius $R_* < c(\delta_0/H_0)^{1/\beta_0}$ meeting only their own windows. Every replacement below is contained in its own safe ball and is selected using only the original tree, never a previously modified modulus.

Favorable local marks. At a deterministic unit pair, $Y = \log E(u, w) - \log D(u, w)$ is centered with uniformly bounded variance. Cantelli gives $\mathbb{P}[Y \le v] \ge cv^2$. Deterministic-point lifting follows directly from the core: an arm starting in $B_{r/4}(u)$ can be attached to u at cost $\varepsilon \operatorname{len}(A) + Z_u$, with $\mathbb{E}Z_u^p \le C\varepsilon^{-C} r^c$ for sufficiently small fixed $p > 0$. Indeed the excess has the form $r^{\psi Q(\psi)} e^{\psi h_r(u)} W$; Hölder's inequality and the core moment give exponent $p\psi Q(\psi) - Cp^2\psi^2 > 0$. Apply this at both deterministic endpoints to obtain

$$D(u, w) \le (1 + \varepsilon) D(B_{r/4}(u), B_{r/4}(w)) + Z_u + Z_w.$$

Choose polynomial point-distance caps, $\varepsilon \ll v$, and polynomially small r so all added failures are $o(v^2)$. This yields a whole favorable E-geodesic with robust coefficient $1 + Cv$ and

polynomial reference bounds. Localize both metrics: cap both $D(u,w)$ and $E(u,w)$ and require their respective outer exit charges to exceed twice those caps. Large-scale negative crossing moments are bounded by $CR^{-p\psi Q(\psi)+Cp^2\psi^2}$, and similarly for η. A polynomially large buffer makes both exit failures $o(v^2)$. The internal source relaxed distance, truncated by its collar charge, is then a local lower certificate; the target geodesic stays in its buffer.

At a local old-tree crossing, Lemma 5.2 gives a scale-independent class of controlled traces. Lemma 5.3 gives finitely many off-trace candidate holes. In each hole place trial buffers at geometrically decreasing, separated, offset annular scales about one deterministic parent center, with pinning functionals inside their respective annuli. The annular Gaussian block Hölder inequality of [1] bounds the probability that n trials all fail by $(1-cv^2)^{n/P_0}$. Taking $n = Cv^{-2}\log(J_{\mathrm{hole}}/p_*)$ puts a favorable mark in every hole outside probability p_*. The smallest trial scale is $\exp(-v^{-C})$. A joint Gaussian pin and oscillation cap of size Cv^{-C} controls unit conversion; the extra factor is $\exp(O(|\eta-\psi|v^{-C}))$, absorbed when it is $1+O(v)$. No independence of a dense spatial array is invoked.

Insertion and its physical cost. Choose a missed hole and its mark. The old lead windows lie in the first and last quarters of a controlled crossing subarc. Parametrize the mark on $[0,1]$ at constant E speed. Choose $t_v \geq \exp(-v^{-C})$ so that its larger endpoint tails $[0,4t_v]$ and $[1-4t_v,1]$ lie in the robust endpoint balls with buffer. Put lead windows in $[0,t_v]$ and $[1-t_v,1]$, and define the protected middle as $\Gamma[3t_v,1-3t_v]$, with actual endpoints c^-,c^+. Lemmas 5.3–5.4 give two disjoint actual leads with normalized total replacement cost $K(v)\leq\exp(v^{-C})$ and quasigeodesic constants $B_v\leq\exp(v^{-C})$. The finite-grid moment union counts at most $\exp(v^{-C})$ distinct links. The lead tubes avoid the complete old trace and mark outside their own windows. Last/first contact trimming therefore gives an embedded replacement, removes an old interval of length at least $m_0 s_r/2$, and retains the full protected middle.

The local crossing lower bound is $m_0 s_r$, whereas the old tree span is at most $CH_0 r^{\beta_0}$. Thus $s_r \leq CH_0 r^{\beta_0}/m_0$. Choose a depth $M_v \leq v^{-C}$ so that $2^{-M_v} < R_*$ and $K(v)CH_0 2^{-\beta_0 M_v}/m_0 < \delta_0/100$. All tests just used have arbitrarily high fixed probability and one fixed annular determining buffer. The all-point annular density estimate of [1], with its finest-grid union bound, supplies a successful scale in $[M_v,2M_v]$ at every adaptive slot, outside Ce^{-cM_v}. The physical coarse caps in this block retain the same effective unit bounds. Perform all replacements in parallel. Their supports are disjoint, their costs are at most $\delta_0/100$, and $\mathrm{len}_E T \leq L_0 + S\delta_0/100 =: M$. The scale caps give $a_j,\rho_j \geq \exp(-v^{-C})$. If needed, round a_j down on a $(1+cv)$ geometric mesh; the favorable coefficient absorbs this factor.

Small arcs and cut separation. Replacing disjoint old edge intervals by embedded arcs meeting the retained tree only at their endpoints preserves the tree and all terminals. A simple arc with no complete protected middle can contain incomplete replacements only at its two ends. In between, on each original edge it runs through a gap of length at most $2\delta_0$; there are at most $2N-3$ original edges. Its length is at most $2(2N-3)\delta_0 + 2\delta_0/100 < \varepsilon_0$. At either cut $c^\pm$ keep a mark window of parameter radius $t_v/2$. Opposite retained old branches have an intrinsic gap at least $m_0 s_r/2$, so (20) separates them quantitatively. Lead tubes avoid the opposite old component, opposite lead, and opposite mark portion. The missed hole separates the mark from the entire old trace. Other inserted arcs lie in different original safe balls. These cases exhaust all pairs outside the kept mark window and give separation at least $\exp(-v^{-C})$. Thus closer opposite-side points both lie on the mark, where its own inverse modulus gives the asserted cut-specific estimate. For a gate $T\cap B_g(c^\pm)$, the forward mark modulus gives positive parameter separation of its remote remnants, and the inverse modulus converts this to Euclidean separation. Gates contain the actual protected endpoints with buffer; deleting an unvisited gate therefore excludes a complete protected middle, including in an isolated chamber component.

Uniform control including lead endpoints. Original tree and mark points already have common strict thickness bounds by their occupation estimates and Lemma 3.2. At an endpoint

x of a B-quasigeodesic lead, the endpoint chains provide a reference E arm across a radius-r annulus of cost at most Cr^{ζ}. Use chain mesh $r/[1+\log(1/r)]^2$ and a good rational vertex in an offset box outside $B_{2r}(x)$; the same-level lattice bridge gives this arm without assuming that an arbitrary nearby rational point is cheap. On an asymmetric core event its connection to the lead's outgoing arm costs at most ε times the lead arm plus K_ε times the reference arm. The last-visit prefix times therefore satisfy

$$T_{\mathrm{in}} \le B\varepsilon(T_{\mathrm{out}} - T_{\mathrm{in}}) + CBK_\varepsilon r^{\zeta}.$$

Take $\varepsilon = 1/(2B)$ and iterate along the positive-density good annuli. This gives a uniform endpoint prefix bound $C_B r^{\zeta'}$ with effective $\zeta' > 0$. For a radius-r ball centered anywhere on the lead, either an endpoint is within $\sqrt{r}$, so this bound applies, or both endpoints are farther away and the interior contraction applies between r and $\sqrt{r}$. Hence the whole lead, including its endpoints, has span $A_v r^{\beta_v}$ with $\beta_v \ge \exp(-v^{-C})$. The uniform lower-shell bound with slack below $\beta_v/2$ gives $H_k(x) \le [Q(\eta) - \beta_v/(2\eta)]k\log 2 + C_v$ on every lead point beyond one common onset. The explicit inverse-crossing tails of [1] bound that onset effectively. Taking the minimum gap over the finitely many pieces proves whole-carrier thinness and the claimed cross-parameter internal access.

Finally encode covered-and-hit cells, slot and trial indices, tubes, and endpoint windows on a rational mesh μ below their separation reserves. There are $O(\mu^{-d})$ cells and finitely many bounded choices per cell; the logarithm of the catalog size has a fixed-power bound in μ^{-1} and the preceding index bounds. Closed-slack compact path certificates or universally measurable selection on the completed Gaussian space give measurable labels. All choices were simultaneous over slots. Trial amplification adds one exponential scale; routing and moments take powers, physical scale selection takes a logarithm, and endpoint control adds only a fixed number of exponentiations. This proves the stated dependence. □

6 Separate companions, local wells, and comparison

Throughout this section $I \Subset \{\psi > 0 : Q_d(\psi) > 0\}$ is fixed, $d \ge 3$, $D = D_h^{\psi}$, $E = D_h^{\eta}$, and $\tau = |\psi - \eta|$. We use mean-log normalization until the final normalization step. Constants may depend on I, d and on the fixed buffered domains. A bound called effective below is bounded by $\mathsf{T}_j(v^{-C})$, with fixed j, C; the same convention applies to reciprocals of positive constants. The height is increased only a fixed number of times in this section. The one-metric core and moment inputs are from [1]. The finite-count and confidence-reset organization follows [2, 3]; the companion, restoration, and sewing steps needed in the present regime are proved below.

6.1 A separate source companion

Lemma 6.1 (Finite-depth companion)**.** *Let T be a connected finite union of E-rectifiable arcs, with total mass at most M, contained in a fixed compact subset of an open buffer U. Suppose all its points satisfy the common thin-point bounds of Lemma 3.2. For $w, a > 0$, sufficiently small, there is a finite source-rectifiable connected network $K_D \subset T^{+w}$ with mass at most CM, containing a finite set $F \subset T$ such that*

$$\sup_{t \in T} D(t, F; U) \le a. \tag{6.1.1}$$

It may also contain any prescribed finite subset of T. The mass bound is independent of w, a. The assertion holds on the simultaneous finite-depth rough-comparison and core events. The required depth is effective in w^{-1}, a^{-1} and the displayed control constants, and is used only when compatible with the parameter restriction in Proposition 4.1.

Proof. Write the controlled ambient bounds on T as

$$E(x,y) \le H|x-y|^{\beta}, \qquad D(x,y;U) \le H|x-y|^{\zeta}, \qquad E(x,y), D(x,y) \ge h(|x-y| \wedge 1)^{b}.$$

Only the appropriate separate lower exponents are needed; taking their maximum simplifies notation. Choose a maximal ambient-E q-separated net. An Euler traversal of the given arc system has length at most $2M$. Consequently the net has at most $1 + 2M/q$ vertices, and its ambient-E minimum spanning tree has total edge weight at most $2M$ and largest edge at most $2q$. Distinct net vertices have Euclidean separation at least $(q/H)^{1/\beta}$. The covering property and the lower E bound give

$$D(t,F;U) \le H(q/h)^{\zeta/b}.$$

Thus q may be chosen to ensure the required a-access, and then decreased so that $3C_R q < h(w/10)^b$, where C_R is the rough-comparison constant. If extra terminals are prescribed, decrease q further so that their source access to the net has total cost at most $M/10$ and stays in $T^{+w/4}$; the common internal thin-point bound makes this possible without requiring separation between those extra terminals. The exit barriers identify the short ambient distances with buffered intrinsic distances whenever this is required.

Apply Proposition 4.1 to each spanning-tree edge, with endpoint displacement r and additive allowance u. The resulting source paths have total length at most $2C_R M + Nu$, where $N \le 1 + 2M/q$. Choose $u \ll M/N$. Each path stays in $T^{+w/4}$: a first exit from the corresponding $w/4$ ball would contradict the source lower spatial bound. At a net vertex, all incident paths start in its r-ball and leave a ball of radius

$$R \ll \min\{w, (q/H)^{1/\beta}\}.$$

Their arms cross J disjoint successful source annuli between r and R. For any two incident paths, one such annulus has combined arm length at most the sum of their entire lengths divided by J. Join to a fixed incident path using the source all-arm connector. Choosing $J \ge CN$ makes the total added length at most a fixed small fraction of the original total. These connectors remain in T^{+w}.

At each exact terminal include a short source reference germ, supplied uniformly by its thin-point bound, and join it in the same manner. Decrease R so that the sum of the N reference lengths is at most $M/10$. This proves the asserted mass bound and exact containment of the terminals. Append the already chosen accesses from the extra prescribed terminals. The resulting object is a finite union of rectifiable paths; its prescribed junctions give a finite weighted graph, so no assertion about the number of incidental intersections is needed. For the all-point event, choose the available logarithmic depth comparable to

$$J + \log R^{-1} + \log(1/p) + \log(1 + \operatorname{diam} U),$$

where p is its allowed failure probability. The $\log R^{-1}$ term pays the spatial-grid entropy; the marginal confidence is chosen first to absorb the finest-grid part of that entropy. All choices therefore have the asserted effective bounds. In particular this argument does not send w to zero at a fixed unequal parameter pair: the finite-depth restriction is retained. □

6.2 Repairing endpoints at controlled contacts

Lemma 6.2 (Exact contact repair)**.** *Let F be either metric. Suppose x, y are controlled interior contacts of a bounded weighted source geodesic, and the thin-point estimates give short F-reference germs from each contact, of cost at most HR^{ζ} through radius R. Suppose also that nearby controlled contacts have the corresponding internal upper bound. If Q is an F-path between $B_r(x)$ and $B_r(y)$, then, on the simultaneous source-independent F-core event,*

$$F(x,y) \le (1 + 2C_{\mathrm{arm}}/J) \operatorname{len}_F Q + CHR^{\zeta}, \tag{6.2.1}$$

provided $r \le cRe^{-CJ}$ and the relevant buffers are available. The conclusion is internal when all its ingredients are internal. There is also an asymmetric form: for $\epsilon \in (0,1)$,

$$F(x,y) \le (1+2\epsilon)\operatorname{len}_F Q + CK_\epsilon HR^\zeta, \qquad K_\epsilon \le C\epsilon^{-C}, \tag{23}$$

on a fixed arbitrarily high-confidence local reference event. Uniformly over a compact guard, the latter can be imposed with failure p using r for which

$$\log(R/r) \le C\{\log R^{-1} + \log \epsilon^{-1} + \log p^{-1} + \log(1+\operatorname{diam} U)\}. \tag{*}$$

The constants in this polynomial precision bound may depend on the fixed thin-point controls and on the prescribed annular confidence.

Proof. If $|x-y| \le 4R$, use the controlled-contact upper bound directly. Otherwise Q crosses the J annuli about both endpoints. At either end, pair its arms with the short reference germ and pigeonhole over the disjoint annuli. The connector costs at most $C_{\mathrm{arm}}J^{-1}$ times the sum of the two path lengths. Truncating Q and adding the two reference pieces proves the bound. For the asymmetric form choose the same-metric core depth $L = C\log\epsilon^{-1}$, so each arm attaches to the common core at cost at most ϵ times its length. Its normalized mass has a fixed positive moment polynomial in ϵ^{-1}, and the normalized crossing has a negative moment. Hölder's inequality therefore bounds the corresponding mass-to-crossing ratio by $K_\epsilon \le C\epsilon^{-C}$ on a fixed high-confidence event. Charge the core mass to the reference arm's crossing cost. The resulting connector costs at most $\epsilon \operatorname{len} Q_{\mathrm{arm}} + K_\epsilon \operatorname{len} R_{\mathrm{ref}}$. One successful annulus at each endpoint proves the asymmetric bound. The all-point annular cover, with entropy term $\log R^{-1}$ and the chosen failure allowance, gives (*). For source repair the reference germs may be taken from the original source path and bounded by its visit spans. For target repair, Lemma 3.2 supplies the internal reference paths. Only interior contacts are repaired; unrestricted final endpoints are left displaced. No target length of a source path is used. □

6.3 The pure tube well

Proposition 6.3 (Local active–inactive certificate)**.** *Fix a small $v > 0$. There is a finite family of deterministic nonnegative smooth profiles f_ℓ, supported in a fixed annular collector $V \Subset U \Subset W$, away from the site pin, with the following properties. Its cardinality J_v, the squared Cameron–Martin norms, the reciprocal of a threshold $b_v > 0$, and all reference constants are effective. Its source validity events have union of probability at least $1-p$, for any prescribed fixed small $p > 0$. On a valid source, a selected profile has:*

1. *a gain $g > 0$ on every eligible original source crossing;*
2. *for every output source geodesic, under any other discounts in disjoint working regions, its entire first-to-last support interval contained in W;*
3. *either an interval of output source length at least b_v with an actual target shortcut in W of ratio at most $1 + Cv$, or an exact original-source restoration of all support visits at cost at most a_{inact}, where $a_{\mathrm{inact}} < g/16$;*
4. *an original-source restoration of cost at most A_0 in all cases.*

The constants g, A_0 can be fixed before v. The parameter restriction and the required finite-depth references have fixed-height-tower bounds.

Proof. Use the same-source all-arm core and the local unweighted span cap to capture the two arms of each eligible crossing by a finite rational menu, at cost at most η_{cap} per arm. The captured interval has source cost at least $a_0 > 0$. Fix $\eta_{\mathrm{cap}} < a_0/100$. The menu is the entire deterministic menu allowed by a high-probability core-depth truncation. Apply Proposition 5.5 with this menu and with a small unmarked-arc allowance ε_0. Choose ε_0 so that $C_R\varepsilon_0 \ll a_0$.

We obtain an embedded target tree T, of mass at most M independently of v, with robust protected middles and controlled gates. Let M_D be the bound in Lemma 6.1.

The gates are chosen before the discount. Each is a carrier stub about an actual protected endpoint, contained with buffer in its robust source endpoint neighborhood. Its target tail cost is at most $cva_{\min}$. The compact remnants obtained by deleting smaller gate interiors have only the intended finite tree incidences and a positive separation δ_{sep} between nonincident pieces. This minimum includes opposite remote branches, not merely remote pieces versus the protected stubs. Choose $\theta > 0$ so small that

$$2\eta_{\mathrm{cap}} + \theta M_D < a_0/4, \qquad 3\theta M_D < \min\{a_{\mathrm{exit}}, h(\delta_{\mathrm{sep}}/4)^b\}. \tag{24}$$

These are effective choices by the carrier bounds.

Next impose the local references for all θ^{-1}-quasigeodesics. Lemma 3.2 gives $H_\theta, \zeta_\theta > 0$ such that an actual output support contact within Euclidean distance w of a carrier point has internal access to it in both source and target metrics of cost at most $H_\theta w^{\zeta_\theta}$. Output endpoints lie outside W, so all these contacts are interior. Choose $\eta_{\mathrm{net}} \ll \theta v a_{\min}$, and refine the ambient-E carrier net until its source access is at most η_{net}. Include every net vertex as an exact companion terminal. Its source mass bound remains M_D. Only now choose w with

$$H_\theta w^{\zeta_\theta} \ll \theta v a_{\min}, \qquad w \ll \delta_{\mathrm{sep}}, \tag{6.3.2}$$

and construct the source companion in $T^{+w/10}$. Choose a smooth pure discount with source weight θ on both T and its companion, between θ and one everywhere, and support $S \subset T^{+w} \Subset V$.

The original crossing can be replaced by its two capture paths and the discounted companion, at cost at most $2\eta_{\mathrm{cap}} + \theta M_D < a_0/4$. This proves a fixed insertion gain g. For any two actual output support contacts, source access to the carrier, the fine carrier net, and the companion give a modified source competitor of cost

$$B_\theta = 2H_\theta w^{\zeta_\theta} + 2\eta_{\mathrm{net}} + \theta M_D \le 2\theta M_D. \tag{6.3.3}$$

The output geodesic is a θ^{-1}-quasigeodesic for its source metric, so the preceding local references apply uniformly in the exterior field. The exit barrier and (24) confine its entire support span to W. Every excursion outside S between support contacts has unchanged source cost at most B_θ.

A gate is visited only at a point of the closed support intersected with the w-thickening of its carrier stub. For an unvisited gate about a protected cut c, assign a support contact the left or right component of $T \setminus \{c\}$ containing its nearby carrier point. The opposite branch remnants outside that gate have separation at least δ_{sep}, so the label is unambiguous. If both labels occurred, the last contact of one label before the first contact of the other would give an outside-support excursion of source cost at least $h(\delta_{\mathrm{sep}}/2)^b > B_\theta$. Thus every support contact has the same label at every unvisited gate. This argument does not assume that the intersection of the carrier with a Euclidean gate ball is connected. If both gates of a protected middle are visited, choose their ordered actual support contacts. The robust lower test makes the intervening weighted source length at least θa_e. Target access to the carrier is paid at full density; the two small carrier tails and the protected middle receive the discount. Since the target weight differs from the source weight by at most $e^{\pm\tau\|f\|_\infty}$, their total cost is at most

$$2H_\theta w^{\zeta_\theta} + e^{\tau\|f\|_\infty}\theta(1 + Cv)a_e \le (1 + C'v)\theta a_e.$$

Require $\tau\|f\|_\infty \ll v$ and set $b_v = \theta a_{\min}/2$. This gives the active alternative with a strict reserve.

If no opposite gate pair is visited, choose one unvisited endpoint gate for every protected middle. The carrier projections of the first and last support contacts are on the same side of each chosen cut, so their unique tree arc crosses none of them. It therefore contains no full protected middle. The gates contain the actual protected endpoints with a positive buffer, so no endpoint exception occurs. The carrier's small-image property and the two accesses give an internal original-target connection between the first and last support contacts of cost at most $2\varepsilon_0$. The reciprocal rough comparison and Lemma 6.2, in the source metric, make this an exact source restoration costing at most $4C_R\varepsilon_0 < g/16$. The undiscounted source companion gives general restoration at cost $A_0 = M_D + 1$. All these images stay in W.

For finite coding, record the carrier by covered-and-hit cells on a rational mesh much finer than w. Its finite incidence and strict inequalities persist under this coding. A mollified indicator of the resulting plateau gives the required deterministic profiles. There are at most $\exp(Cw^{-C})$ additional codes, and

$$\|f\|_{\mathrm{CM}}^2 \le C\psi^{-2}(\log\theta^{-1})^2 w^{-C}.$$

The homogeneous $\dot{H}^{d/2}$ norm is invariant under site dilation. Validity is expressed by the finite geometric certificates, the rough comparison event, and the simultaneous thin-point and core references. Their analytic or coanalytic formulations are measurable in the completed local Gaussian sigma-field. Selecting the first valid finite code loses no factor equal to the number of codes in the probability of validity. The carrier construction, the order θ, then references, net, then w, and the finite-depth companion give a fixed-height tower bound. The local validity tests use each metric divided by its *own* scale unit, so are invariant under adding a constant to the local field. At a physical site the separate coarse-pin event gives $|\log(s_E/s_D)| \le cv$ when $\tau n \le cv$. This event is paid globally, not inserted into the annular mixing event. Its factor is absorbed in the reserved $O(v)$ shortcut slack and in the fixed restoration constants. All local and global pins are outside the corresponding profile supports. □

6.4 The finite Gaussian forcing count

Lemma 6.4 (Gaussian overlap)**.** *For deterministic Cameron–Martin shifts f_i, $i \in \mathcal{A}$, with $\#\mathcal{A} = N$ and $\|f_i\|_{\mathrm{CM}} \le M$, suppose $\sum_i 1_{A_i}(h) \ge L_- 1_B(h)$ and $\sum_i 1_{A_i}(h+f_i) \le L_+$. Then*

$$\mathbb{P}(B) \le \frac{L_+}{L_-}\min\{N, \exp(M\sqrt{2\log N})\}.$$

Proof. Change variables in each Gaussian integral. The resulting densities have pth moments at most $\exp((p-1)M^2/2)$ after taking the pth root. Bound their maximum by the ℓ^p norm and choose $p = 1 + \sqrt{2\log N}/M$. Summing their expectations gives the alternative bound N. Zero-norm and singleton cases follow directly. □

Proposition 6.5 (Uniform coefficient improvement)**.** *Fix $C_b > C_R$, and $c \in [1 + C_0 v, C_b]$, with C_0 sufficiently large. There are $g_v, e_v, t_v > 0$, with effective reciprocals, such that a high-probability global comparison of coefficient c and error e_v on a fixed guarded ball implies a local annular comparison seed of coefficient $c - g_v$ with any prescribed sufficiently high fixed confidence. The constants are uniform in c in this interval, provided $\tau \le t_v$. Here is the precise grouped seed, in unit coordinates and source units. Put*

$$K = \{1/2 \le |z|_\infty \le 3\}, \qquad U = \{1/4 < |z|_\infty < 4\}, \qquad S_t = \partial B_t^\infty.$$

The reference event includes a lower charge $m_* > 0$ for crossing S_1 to S_2, lower source and target exit charges from $\{1 \le |z|_\infty \le 2\}$ to K^c and U^c, respectively, and the local controlled-contact estimates. For every intrinsic D^K-geodesic parent from $S_{1/2}$ to S_3 of length at most

M_0, take its first S_2 hit and last preceding S_1 hit. Divide that middle crossing into N_0 equal source-arclength pieces. The fixed integer N_0 is large enough that each piece is shorter than the source exit charge and than the target exit charge divided by $2C_b$. Each piece has length at least m_*/N_0 and its endpoints have fixed buffers from the parent endpoints and the working boundary. The seed asserts that each such piece $[x, y]$ has

$$E(B_e(x), B_e(y); U) \le (c - g_v)\,\mathrm{len}_D[x, y]. \tag{S}$$

It suffices for amplification to retain one piece, so one may weaken "each" to "at least one." The numbers m_*, M_0, N_0, and the finite offset geometries are fixed before the improvement accuracy; the precision e is relative to the site radius. The normalized seed includes its pin sphere in its fixed determining annulus. The forcing calculation below preserves the original global pin and each counting-site pin.

Proof. We give the counting calculation and the localization, since these are the steps where a boundary-diameter argument would be invalid. To preserve the original field normalization, perform this calculation after one fixed homothety of the displayed seed geometry. Choose its factor $r_0 = r_0(d) > 0$ so that the complete closed source working guard W lies in $B_{1/2}$, while the field retains its original unit-sphere pin. Every counting profile is supported in the interior of W and therefore vanishes on that pin sphere. Source distances in this count are the constrained distances D^W. The localized source pieces are shorter than the exit charge from a smaller compact guard inside W, so their constrained minimizers remain in that smaller guard. The canonical discounted minimizers are taken in W; the local references only use their quasigeodesic property in the buffered site regions. Include an untouched outer collar whose source crossing charge is larger than U_0, and take every counting region in its inner compact complement. The fixed subdivision number N_0 is increased first so that its piece upper bound U_0 is below this charge. Original and discounted constrained minimizers then stay a positive distance from ∂W, since they have cost at most U_0 and crossing the collar receives no discount. The old global comparison still applies to every source piece because $D(x, y) \le \mathrm{len}_{D^W}(P[x, y])$. Let $K_i \Subset L_i$, $i = 1, 2$, be two separated compact endpoint sets. At a source H suppose

$$E_H(L_1, L_2) \ge (c - g_\Delta) D_H^W(K_1, K_2), \qquad D_H^W(K_1, K_2) \le U_0. \tag{6.5.1}$$

Require the previous comparison, a guard for all discounted geodesics, and a common-source reference for all weighted source geodesics with weights in $[\theta, 1]$ and original source length at most $\theta^{-1}U_0$. This last reference supplies target contact access and target core sewing at all needed junctions. Its exceptional probability is paid once.

For the all-point cover use nested ancestor grids: at radius r, centers lie on $cr\mathbb{Z}^d$, with a fixed small $c > 0$. For each deterministic finest-grid point z, its radius-r ancestor z_r obeys $|z - z_r| \le cr$. The complete determining footprint and pin sphere of its local event lie in a fixed annulus $A(z; (a - c)r, (b + c)r)$, with $a - c > 0$. Enlarging to these fixed-center annuli and separating levels by a fixed larger factor permits the same Gaussian block inequality; no stability under a translation of the event is asserted. At each radius a fixed number of lattice colors makes the full working regions disjoint. Fix one radius, color, and deterministic unit bin *before* the count. Its entire physical-site universe, including sites appearing only in shifted outputs, has disjoint working regions. It has at most $C\Delta^{-D_0}$ sites and source units in $[Q, 2Q]$, where $Q \ge \Delta^{K_0}$. Its original source geodesic visits at least $N_{\mathrm{site}} \ge \Delta^{-\rho}$ eligible central regions, for some fixed $\rho > 0$, once Δ is sufficiently small. To see the latter bound, the two endpoint sets have fixed separation, whereas every ball in the cover has diameter at most $C\Delta$. The fixed colors lose only fixed factors; the radius and unit bins lose powers of $\log \Delta^{-1}$. The common pin cap gives the stated lower bound for Q. Every catalogue inspection is required to avoid a fixed small Euclidean neighborhood of the larger endpoint sets L_i, with width a

fixed fraction of their uniform positive separation. A continuous crossing still travels a fixed Euclidean distance between these neighborhoods, so the preceding supply bound is unchanged. This gives a fixed buffer from every possible original endpoint, independent of the eventual endpoint-grid precision. The common-source contact-reference constants therefore are fixed before Δ. The local event in Proposition 6.3 has sufficiently high marginal probability for this cover, uniformly over the source paths.

Select any k supplied sites. Their original confined crossing insertions save at least gkQ in the source metric. Delete an inactive selected site, recompute the canonical output geodesic, and repeat. An unvisited site costs zero; an inactive deletion costs at most $2a_{\mathrm{inact}}Q$ by exact source restoration. The final set Z therefore retains a saving of at least $gkQ/2$, and every survivor is active. General source restoration gives

$$|Z| \geq \kappa k, \qquad \kappa = g/(4A_0). \tag{6.5.2}$$

Fix one valid label at each supplied site before this pruning. A survivor of size m has at most N_{site}^{k-m} preimages. Thus the source event is covered with multiplicity at least

$$L_- \geq N_{\mathrm{site}}^{\lceil \kappa k \rceil}/k^k.$$

At an output h, every active index Z has source $H = h + f_Z$ and discounted field exactly h. It therefore uses the same selected constrained source geodesic P_h between K_1, K_2. The original global pin, every counting-site pin, and all site units are unchanged. In particular no repinning scalar enters the counting calculation; such a scalar could not be common to the two parameters. Let $\mathcal{V}$ be the physical sites occurring in active indices. Each supplies a block of length $L_i \geq b_v Q$ and a target shortcut of ratio at most $c_{\mathrm{act}} = 1 + Cv$. Its entire support span is in its own working region, so these intervals are pairwise time-disjoint.

Fix one active index Z_0 and put $H_0 = h + f_{Z_0}$. Retain the target shortcuts at the $m = |\mathcal{V} \setminus Z_0|$ other sites; they are unchanged in H_0, since their images avoid the Z_0 regions. Restore the Z_0 support spans by source-D_{H_0} connections, at cost at most $2A_0kQ$. Apply only the comparison at H_0 to the remaining source pieces and these restoration jumps. Repair every interior target recovery endpoint by Lemma 6.2. The single common-source reference applies: P_h is a weighted constrained $D_{H_0}^W$ geodesic and has original length at most $\theta^{-1}U_0$. The final endpoints stay displaced inside L_1, L_2. Writing $L_h = D_h^W(K_1, K_2)$, we obtain

$$\begin{aligned} E_{H_0}(L_1, L_2) \leq cL_h - (c - c_{\mathrm{act}}) \sum_{i=1}^{m} L_i + 2cA_0kQ \\ + C\epsilon_{\mathrm{sew}}(U_0 + A_0kQ) + C(m + 2k + 1)(u + u_*). \end{aligned} \tag{25}$$

Here u is the old additive error, u_* the absolute repair error, and ϵ_{sew} the relative repair error. No target length is assigned to a source piece. Choose

$$g_\Delta U_0, \quad \epsilon_{\mathrm{sew}} U_0, \quad u + u_* \leq c_*(c - c_{\mathrm{act}})b_v\Delta^{K_0}, \tag{26}$$

with c_* small. Source monotonicity gives $L_h \leq D_{H_0}^W(K_1, K_2)$. Comparing (25) with the near-bad inequality at H_0 and absorbing the term linear in m gives

$$|\mathcal{V}| \leq K_1 k, \qquad K_1 \leq C\left(1 + \frac{C_b A_0}{v b_v}\right).$$

Thus the output multiplicity is at most $L_+ \leq (CK_1J_v/\kappa)^k$. Only this one source comparison has entered the upper bound.

There are at most $(CJ_v\Delta^{-D_0})^k$ deterministic indices. In every dimension the safe norm bound is $\|f_Z\|_{\mathrm{CM}} \le kM_v$; disjoint supports need not be orthogonal in $\dot{H}^{d/2}$. Lemma 6.4, with $\ell = \log \Delta^{-1}$, gives

$$\begin{aligned}\log \mathbb{P}(B) \le & -\rho\kappa k\ell + k\log k + k\log(CK_1J_v/\kappa) \\ & + kM_v\sqrt{2k(D_0\ell + \log(CJ_v))}. \end{aligned} \tag{27}$$

Choose k to beat the endpoint entropy and then choose ℓ larger than a fixed polynomial in $k, M_v, \kappa^{-1}, \log J_v, \log K_1$. The count is then smaller than any prescribed fixed power of Δ. The source-reference failures are charged once, before endpoint union.

For completeness, the grouped local seed is obtained as follows. Use the parent class and the N_0 pieces just specified. Each piece is globally source minimizing, since a shorter global competitor leaving K would spend more than its length on the source exit charge. Its endpoints are separated by a fixed positive Euclidean amount on the controlled-contact reference event. A failure of (S), with a fixed strict reserve in g_Δ, gives a bad pair of endpoint balls. Round their centers to a fine rational grid, with $K_i \Subset L_i$ and L_i still inside the bad balls. Any target path leaving U costs more than C_b times the source-piece length, so the bad internal lower test is also a global lower test. Since $D^W(K_1, K_2)$ is at most that source-piece length, the near-bad inequality in (6.5.1) persists with a slightly larger reserve. If localization starts with exact endpoint distances, the asymmetric form of Lemma 6.2 first supplies these robust balls. All these predicates use intrinsic metrics in their fixed buffers and certified exit barriers, so are local under every exterior extension. The number of endpoint problems is at most $C\Delta^{-D_1}$, with effective D_1 fixed before k. This proves the local seed after the preceding union bound.

In (26) take $g_\Delta = cvb_v\Delta^{K_0}/U_0$. Use the *asymmetric* form of Lemma 6.2 for localization and for the upper count. Set ϵ to a small multiple of g_Δ/U_0 and choose R so that $K_\epsilon H_\theta R^{\zeta_\theta} \ll g_\Delta$. After absorbing fixed v-dependent prefactors by increasing ℓ, both R and the available recovery radius are powers of Δ, with exponents fixed before k and ℓ. This is why the endpoint entropy exponent D_1 above is fixed; using $J \asymp g_\Delta^{-1}$ for that localization would not suffice. More explicitly, fix an output seed precision Δ^{A_v} with $A_v > LK_0 + 1$, where L is the fixed seed-precision exponent in Lemma 6.7; absorb the fixed v-dependent prefactor by the eventual lower bound on ℓ. Choose the previous global precision Δ^{B_v} fine enough for the upper count and choose its additive error at most $cvb_v\Delta^{K_0}$. The exponents A_v, B_v, D_1 are chosen first, then k, then ℓ. Thus no endpoint entropy is hidden in the choice of the counting depth. The required depth is at most $C(v)(1+\ell)$. Require τ times this depth, and $\tau\|f\|_\infty$, to be sufficiently small. The resulting g_v, e_v, t_v have effective reciprocals. All choices use only $c - c_{\mathrm{act}} \ge cv$ and $c \le C_b$, hence are uniform in the current coefficient c. Finally undo the fixed homothety, explicitly retaining its metric units $s_D = r_0^{q_d(\psi)}e^{\psi h_{r_0}}$ and $s_E = r_0^{q_d(\eta)}e^{\eta h_{r_0}}$. A fixed high-probability coarse-pin cap gives

$$\left|\log \frac{s_E}{s_D}\right| \le C\tau(1 + |h_{r_0}|).$$

Require this bound to be less than $g_v/(10C_b)$ and reserve a fixed fraction of the improvement before naming g_v. The seed in its own source and target units then has exactly the asserted coefficient. The rescaled, repinned field has the original law, so this proves the probability bound for the displayed seed geometry. The coarse-pin cap was paid in this probability estimate and was not inserted into its annular determining sigma-field. Cameron–Martin changes of variables above always used the original pin, which all profiles preserve. The auxiliary homothety pin h_{r_0} need not be preserved by those shifts: it was used only in this source-probability transfer, never in the output count. The final seed itself uses only its intrinsic buffered metrics and its own pin sphere, and is invariant under adding a constant to the local field after division by the two respective units. □

6.5 Local contact sewing and annular amplification

We specify the seed interface used in the count. All its geometries and constants are fixed before the accuracy. In site coordinates a source parent is a confined source-geodesic excursion, of length at most M, which crosses a fixed annular band. Its prescribed middle pieces have length at least $m > 0$, lie in a fixed inner part of the working annulus, and have a fixed Euclidean margin from the parent's endpoints and the working boundary. A grouped seed of coefficient c_0 and precision e asserts that every such parent has one prescribed middle piece $[x, y]$ with an internal target recovery of cost at most $c_0 \operatorname{len}_D([x, y])$, whose endpoints are in $B_e(x), B_e(y)$. One can choose the prescribed pieces by finitely many thin radial shells and a fixed equal-arclength subdivision; the local cap, crossing lower bounds, and spatial lower estimate fix m, M and the margins. The exit references make every recovery used in this definition internal to one fixed annular determining buffer. Requiring the seed at finitely many fixed offset geometries gives this grouped form. This is the seed provided by Proposition 6.5.

Lemma 6.6 (Sewing in each site's own units)**.** *There are fixed $L, C > 0$ with the following property. For $0 < u < u_0$ put $e_u = cu^L$. On a local event of arbitrarily high fixed confidence, every eligible source parent has the following target attachments, uniformly over its actual controlled contacts. At a contact x associated with a site of radius r and source unit s, a target path whose recovery endpoint lies within $e_u r$ of x can be attached exactly to x, at extra cost*

$$\epsilon \operatorname{len}_E(Q) + \lambda s, \qquad \epsilon \le cu, \quad \lambda \le cu. \tag{28}$$

Here the other end is treated as follows: if its associated source contact is too close for the path to cross an attachment annulus, replace the path by the direct controlled-contact target connection, of cost at most λs. The numerical constants c can be chosen arbitrarily small in advance. All added attachment paths stay in the fixed local buffer. For a path with two distant ends, one may perform the two attachments in different local units. In this statement its other recovery error is also at most $e_u r$. In the applications this is automatic: a marked recovery has both errors in its own site units, and an old gap has both errors smaller than the minimum allowable radius over all selected sites.

Proof. The source local span bound and Lemma 3.2 give, in normalized site coordinates, a fixed exponent $\zeta > 0$ and constant H for target reference germs and nearby pairs of source contacts. This local family includes all points of the same original source geodesic in the contact's smaller spatial buffer, even if visited at another time: the parent span estimate controls the entire visit interval, and the original endpoints are outside the outer buffer. The reference event is imposed uniformly for this family, not for a single chosen geodesic.

Use the asymmetric same-target core of Lemma 6.2, with coefficient $\epsilon = cu$ and $K_\epsilon \le Cu^{-C}$. Choose the outer normalized attachment radius $R \le cu^{L_1}$ so that $CK_\epsilon HR^\zeta \le \lambda$. The all-point annular cover at relative scales between u^{L_2} and u^{L_1} supplies a successful annulus about every contact with the required fixed confidence. Its finest-grid entropy is proportional to $\log u^{-1}$; a sufficiently high fixed marginal confidence absorbs it. Choose fixed $L > L_2$ and decrease c in e_u so every recovery endpoint lies in that annulus's inner region. This uses only powers of u. A short reference target arm from x crosses its outer region. If the recovery path does too, the asymmetric connector, with the core mass charged to the reference arm, gives (28) after truncation. If the other associated source contact y lies too close, then y belongs to the same local controlled-contact family. Its direct connection to x has cost at most $H(CR)^\zeta \le \lambda s$ after reducing the chosen radius. Otherwise the recovery path necessarily crosses the annulus, since its other endpoint is close to y. At two ends the selected contacts need not occur in chronological order along Q: the union of Q and both attachments is connected and has mass bounded by their sum, hence contains a path of that cost. The physical source/target unit

ratio is retained; a bound $1 + O(\tau n)$ at the inspected scales changes the constants by that factor only. □

Lemma 6.7 (Amplification and confidence reset)**.** *There are fixed $L, C, c > 0$ such that the following holds. Suppose the oriented grouped seed of coefficient $c_0 \le C_b$ holds with sufficiently high fixed confidence at precision at most $e_u = cu^L$. For source geodesics in a guard B_W of length at most e^{an}, where $a > 0$ is the fixed rough-router height exponent, the conclusion is an internal target recovery in B_{W+1} of cost at most $(c_0 + u)\,\mathrm{len}_D(P)$ and endpoint displacement $e^{-cu^L n}$, simultaneously over those geodesics, outside probability*

$$C(W+1)^C u^{-C} \exp(-cu^L n), \tag{29}$$

provided $n \ge Cu^{-L}(1 + \log(W+1))$ and $\tau n \le cu$. The rough-router references on the same guard are included. A finer seed precision can be enlarged to e_u without affecting these bounds. Thus confidence $1 - q$ and final precision r_* can be reset using

$$n \le Cu^{-L}\{1 + \log(W+1) + \log q^{-1} + \log r_*^{-1}\}, \tag{30}$$

after enlarging C, L. Treating the reverse orientation gives the reciprocal conclusion. The comparisons throughout this proof are multiplicative, with no additive per-piece cost.

Proof. Write $\vartheta = m/(2M)$, decrease it if necessary, and put $q_0 = (1 + \vartheta)^{-1} < 1$. Combine the seed, the fixed span and crossing references, and Lemma 6.6 into a single annulus-local good event. Choose its attachment constants so that the total sewing error below is at most $\lambda_0 L$, where $\lambda_0 \le (1 - q_0)u/20$. This only changes the fixed power in e_u. Its confidence is fixed sufficiently high before constructing any spatial cover. The grouped seed uses finitely many offsets, whose failures are charged inside this one local event.

Choose

$$N = \lceil C \log((C_b + 1)/u) \rceil, \qquad a_j = a_1 \Lambda^{-(j-1)}, \quad 1 \le j \le N,$$

where $\Lambda > 8$ is fixed and $a_1 > 0$ is small compared with the rough-router recovery exponent. Then $a_N \ge cu^{L_0}$ for a fixed L_0. In band j use site radii $[e^{-4a_j n}, e^{-a_j n}]$. The annular block estimate and a finest-grid union bound supply a good central site around every point of B_W in that band. Their failure is at most $C(W+1)^C e^{-ca_j n}$. The determining regions have fixed annular geometry; moving centers are handled by their deterministic ancestor grids as in the count. A separate Gaussian cap at all inspected scales makes the physical source/target unit ratios $1 + O(\tau n)$. Enlarge the coefficient to $b = c_0 + u/10$ to absorb this drift. All these events precede the choice of the source path.

We prove the deterministic improvement in one band. Assume the old comparison gives, for every subgeodesic of the same source path, a target recovery of cost at most C_j times its length, with endpoint radius ρ_{old} and $C_j \le C_b$. Set $r_{\max} = e^{-a_j n}$ and $\rho = Ar_{\max}$, with a fixed large A exceeding all working-buffer constants. If the original endpoints are within 4ρ, their enlarged recovery balls intersect and a constant path suffices. Otherwise, stop first on the terminal closed ρ-ball and take the last preceding visit to the initial closed ρ-ball. Denote this retained source segment by $P : [0, L] \to B_W$. Its two endpoints are within ρ of the original endpoints, and its entire image avoids the two open endpoint balls. Consequently every site working buffer centered near this segment misses the original source endpoints. The local full-visit and contact references therefore apply, including to nearby source contacts from any other part of this same geodesic.

For each time of the retained segment choose a good site for which the path is in its central ball. Use the temporal component of the preimage of that site's fixed outer ball which contains this central visit. These are open intervals covering $[0, L]$; their images, including all excursions of each component, stay in the corresponding outer balls. The local envelope

bounds each full component by Ms_i. Take a finite minimal subcover after restriction to $[0, L]$. Its multiplicity is at most two: among three intervals with common interior, the middle one is redundant. At most one member is clipped at either endpoint. Let their union have length lL. If the two clipped members overlap, their outer balls intersect and both retained endpoints are within $Cr_{\max}$ of each other; enlarging the final recovery radius again gives the diagonal case. Otherwise their union is an initial and a terminal interval.

Every unclipped component contains an eligible parent and hence a marked middle block of length at least ms_i. Two-color the interval cover and retain the color with larger total marked length. Its components have disjoint interiors, and the retained marked time T satisfies

$$T \geq \vartheta(1-l)L, \qquad \sum_{i \text{ retained}} s_i \leq T/m \leq L/m. \tag{31}$$

No lower bound on gaps between consecutive marked blocks is needed. Apply the local seed to the marked blocks and the old comparison to the intervening source gaps. For each whole old gap, first decide whether either endpoint requires the direct short-gap replacement; if so perform that replacement once and do not attach that gap again. For all other gaps, at each marked endpoint attach both incident target recoveries using its own site in Lemma 6.6. If an old gap is too short to cross one attachment annulus, replace it by the direct target connection between its two source contacts in that site's controlled family. If this is an initial or terminal gap, its outer source contact is an allowed recovery point within the original endpoint ball. Otherwise keep its outer recovery endpoint free. All remaining internal junctions are exact. The relative attachment errors sum to at most $C\epsilon C_b L$; the absolute errors sum to at most $C\lambda \sum_i s_i \leq C\lambda L/m$ by (31). Thus the resulting cost is at most

$$\{C_j - \vartheta(C_j - b)(1-l) + \lambda_0\}L \tag{32}$$

when $C_j \geq b$. Keep an already better coefficient otherwise, and retain the old comparison whenever it is the cheaper competitor. This calculation assigns target length only to actual target paths. It never compares the target length of a source segment to its source length, nor requires a recovery ball at a large site to be smaller than a global separation between all marked endpoints.

There is a second competitor: discard the two clipped temporal components and apply the old comparison to the remaining source subgeodesic. Its cost is $C_j(1-l)L$. Each newly retained endpoint is on the boundary of the corresponding site's outer ball, and is within $Cr_{\max}$ of the former retained endpoint. Both competitors therefore have final recovery displacement at most $Cr_{\max} + \rho_{\text{old}}$ from the original endpoints. For $C \geq b$ and $l \in [0, 1]$,

$$\min\{C - \vartheta(C-b)(1-l), C(1-l)\} \leq b + \frac{C-b}{1+\vartheta}.$$

Combining the two routes proves the pure multiplicative recurrence

$$(C_{j+1} - b)_+ \leq q_0(C_j - b)_+ + \lambda_0. \tag{33}$$

Discarding the original two endpoint pieces only decreases source cost, so this holds relative to the untrimmed source length as well.

The old recovery radius must fit each local attachment, namely

$$\rho_{\text{old}} \leq ce_u e^{-4a_j n}. \tag{34}$$

For $j = 1$ this follows from the pure rough comparison and small a_1. For later bands the preceding radius is $Ce^{-a_{j-1}n}$; $\Lambda > 8$ and $n \geq Ca_N^{-1}(1 + \log u^{-1} + \log(N+1))$ ensure (34). There is no additive old error to divide by the smallest source charge. Initialize with Proposition 4.1

and iterate (33) for N bands. Increasing the fixed constant in N makes the final coefficient at most $c_0 + u$. The final radius is $Ce^{-a_N n}$. Summing the band cover failures and the fixed rough and Gaussian failures gives (29), after increasing L. The same increase absorbs $\log u^{-1}$ and the finite band count in the minimum depth. Solving the resulting radius and failure inequalities proves (30). The argument only used a seed at e_u; an arbitrarily finer input seed is first enlarged to that precision. Thus the reset constants are independent of any finer precision chosen in the forcing count. □

6.6 Completion of the reciprocal estimate

Proof of Theorem 1.2. Fix a requested accuracy δ, and take $v = c\delta$ so that the coefficient in Proposition 6.3 is below $1 + \delta/10$. We make the probability reset precise. Choose the forcing depth ℓ, and hence $\Delta = e^{-\ell}$, once as in Proposition 6.5. Fix its previous global radius $r_* = \Delta^{B_v}$, its additive allowance $u_* = cvb_v\Delta^{K_0}$, and a confidence $1 - q_*$ sufficiently high for the forcing implication. If desired decrease r_* so that the same number bounds both errors. Let $\mathcal{G}(c)$ denote the global comparison with coefficient c, these fixed errors, and this fixed guard. Its auxiliary common references have fixed failure allowances chosen before the endpoint union. Start at the rough coefficient C_R. The rational-net and large-barrier argument below, applied directly to Proposition 4.1, supplies $\mathbb{P}[\mathcal{G}(C_R)] \geq 1 - q_*$; this use does not involve a near-one comparison. Proposition 6.5 improves every current coefficient in $[1 + C_0 v, C_b]$ by the same $g_v > 0$ at local-seed confidence. Apply Lemma 6.7 with allowance $g_v/2$ and requested radius r_*. It gives a pure multiplicative comparison, hence the same predicate $\mathcal{G}(c - g_v/2)$, with probability at least $1 - q_*$. The required depth is at most

$$Cg_v^{-L}\{1 + B_v\ell + \log q_*^{-1}\}.$$

Its parameter restriction is included once in t_v. The local seed precision Δ^{A_v} was chosen below cg_v^L, so the amplifier is applicable at this exact depth; no precision needs to be reduced during the iteration. After at most $2C_b/g_v$ repetitions the coefficient is at most $1 + C_0 v$. All depths, profile families and precision allowances were chosen uniformly in the current coefficient. No tower is iterated in this repetition, and failure probabilities do not accumulate: each step is an implication between probability bounds at the same confidence. The final forcing step already supplies a grouped seed of coefficient at most $1 + C_0 v$. A final amplification of this seed, with allowance $\delta/2$, gives the required bounded-height estimate in both directions.

We explain the removal of the height and guard restrictions without changing any of the already fixed local events. Let A be a sufficiently large effective constant, chosen after the amplifier constants, and put $\sigma = n/A$. Take a rational net $\mathcal{N}$ of mesh $e^{-\sigma}$ in B_R. Uniform small positive point moments, including their polynomial dependence on the locations in B_R, and a union bound over net pairs give

$$\max_{p,q\in\mathcal{N}} E(p,q) \leq e^{B\sigma}(R+1)^B$$

outside $C(R+1)^C e^{-c\sigma}$, with fixed sufficiently large B. The same holds for D. A source pair above this height has a target recovery through the net at no greater cost. Infinite source distances are harmless. The remaining finite source geodesics have length at most this height. By the large-radius crossing estimate, they stay in

$$B_W, \qquad W = e^{B'\sigma}(R+1)^{B'},$$

outside another $C(R+1)^C e^{-c\sigma}$ exception, for fixed sufficiently large B'.

When $n \geq CA(1 + \log(R+1))$, their logarithmic height and $\log W$ are both at most $C'n/A$. Choose A large enough that this height is below the rough-router cutoff and that the

established amplifier's failure exponent dominates its spatial factor $(W+1)^C$. Its minimum-depth condition also then holds. The amplifier displacement is at most $e^{-n/A}$ after increasing A once more. This uses the already fixed confidence and geometry; it does not try to increase a marginal success probability after its geometric constants have been chosen. Consequently the bounded-height estimate applies simultaneously to all pairs in B_R, with errors and exceptional probabilities of the required form.

Finally enlarge the fixed tower height and its power to absorb all finite construction depths, reset depths, spatial factors, and minimum parameter restrictions. There is

$$\mathcal{H}(\delta) = \max\{2, \mathsf{T}_H(\delta^{-C})\}$$

such that, whenever $\tau \le \mathcal{H}(\delta)^{-1}$, $n \ge \mathcal{H}(\delta)(1+\log(R+1))$, and $\tau n \le \mathcal{H}(\delta)^{-1}$, the simultaneous reciprocal comparisons have displacement $e^{-n/\mathcal{H}(\delta)}$, additive error at most that displacement, and exception at most

$$\mathcal{H}(\delta)(R+1)^{\mathcal{H}(\delta)} e^{-n/\mathcal{H}(\delta)}.$$

This is precisely the assertion of Theorem 1.2. The planar input [2] gives the same statement in $d = 2$; the argument above supplies every $d \ge 3$. □

7 Completion of the parameter-continuity argument

We work on a compact interval $I = [a, b] \Subset \{\psi > 0 : Q_d(\psi) > 0\}$. Write D_t for the normalization used in Theorem 1.2. All estimates below are uniform for deterministic $t \in I$. We use Proposition 2.1 and the simultaneous same-metric core estimates of [1]. Every probability-one assertion involving all parameters will be obtained from countably many parameter-grid events.

7.1 A complete space for lower semicontinuous convergence

Let $\mathcal{F}$ consist of the symmetric lower semicontinuous functions $f : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty]$ that vanish on the diagonal. For $z = (x, y)$ and $w = (u, v)$ put $d_X(z, w) = |x - u| + |y - v|$, and define

$$T_m f(z) = \inf_w \{f(w) + m d_X(z, w)\}, \qquad K_R = \overline{B}_R \times \overline{B}_R, \tag{35}$$

$$\rho(f, g) = \sum_{m,R \ge 1} 2^{-m-R}(1 \wedge \|T_m f - T_m g\|_{K_R}). \tag{36}$$

Lemma 7.1. *The metric ρ is complete and separable. Its convergence is exactly lower semicontinuous, or epigraph, convergence: $f_n \to f$ if and only if*

$$f(z) \le \liminf_n f_n(z_n) \quad (z_n \to z), \qquad \forall z\ \exists z_n \to z :\ f_n(z_n) \to f(z).$$

Proof. Each $T_m f$ is finite and m-Lipschitz, satisfies $0 \le T_m f(x, y) \le m|x - y|$, and increases to f. For $z \in K_R$, its minimizing points lie in K_{3R}, by the diagonal competitor. Epigraph convergence therefore gives pointwise, hence locally uniform, convergence of every envelope. Conversely, envelope convergence gives the lower-limit inequality by fixing m and then sending $m \to \infty$; approximate minimizers with $m = m_n \to \infty$ sufficiently slowly give recovery sequences.

For completeness, let f_n be ρ-Cauchy and let g_m be the locally uniform limit of $T_m f_n$. The identity $T_\ell T_m = T_{\min(\ell,m)}$ gives $T_\ell g_m = g_{\min(\ell,m)}$. Put $f = \sup_m g_m$. For $m \ge \ell$, a minimizer w_m of $g_m(w) + \ell d_X(z, w)$ stays in a fixed compact set. Along a convergent subsequence, for every fixed j, $g_j(w) + \ell d_X(z, w) \le g_\ell(z)$. Taking the supremum over j proves $T_\ell f \le g_\ell$. The

reverse inequality follows from $f \geq g_\ell$ and $T_\ell g_\ell = g_\ell$. Thus $T_\ell f = g_\ell$, proving completeness. The envelope embedding into the countable product of the separable spaces $C(K_R)$ proves separability. □

Write $E \preceq_{\delta,r} D$ on B_R if every finite $D(x,y)$, with $x, y \in B_R$, admits x', y' within r of x, y such that $E(x', y') \leq (1+\delta)D(x,y) + r$. If both orientations hold on B_{3S}, the envelope competitors give

$$\rho(D, E) \leq C(2^{-M} + 2^{-S} + MS\delta + Mr). \tag{37}$$

Indeed an envelope minimizer for $m \leq M$, $R \leq S$ has cost at most $2mR$; moving both endpoints adds at most $(2m+1)r$. The discarded weights give the two geometric tails.

7.2 Explicit grids and the continuous modification

After increasing its fixed constants, Theorem 1.2 has the following unified form. For $H(\delta) = \max\{2, \mathsf{T}_H(\delta^{-C})\}$, where $\mathsf{T}_0(x) = x$ and $\mathsf{T}_{j+1}(x) = e^{\mathsf{T}_j(x)}$, set $r = e^{-n/H(\delta)}$. If

$$|s-t| \leq H(\delta)^{-1}, \qquad n \geq H(\delta)(1 + \log(R+1)), \qquad |s-t|n \leq H(\delta)^{-1}, \tag{38}$$

the reciprocal comparisons on B_R fail with probability at most $H(\delta)(R+1)^{H(\delta)} e^{-n/H(\delta)}$.

Choose

$$\delta_j = 2^{-j-10}(j+1)^{-2}, \qquad k_j = \left\lceil \mathsf{T}_D(2^{Bj}) \right\rceil, \qquad h_j = (b-a)2^{-k_j},$$

where $D \geq H + 2$ and B is a sufficiently large fixed constant. Let $Q_j = a + h_j\{0, \dots, 2^{k_j}\}$, so the grids are nested. Write $q_j(t)$ for the left projection onto Q_j, with $q_j(b) = b$. Put

$$H_j = H(\delta_j), \quad n_j = \left\lfloor (4H_j h_j)^{-1} \right\rfloor, \quad A_j = n_j/H_j, \quad r_j = e^{-A_j}, \quad s_j = \sum_{i \geq j} r_i. \tag{39}$$

Fix a sufficiently large C_0 and use spatial comparison windows B_{4R_j} with $R_j = e^{C_0 k_j}$. The elementary tower estimates give

$$\frac{A_j}{k_{j+1} + H_j k_j + \log H_j + j} \longrightarrow \infty, \qquad \frac{A_{j+1}}{A_j} \longrightarrow \infty, \qquad s_j \leq 2r_j \tag{40}$$

for all sufficiently large j. Here fixed tower height is essential: an unspecified cutoff depending on δ would not justify (40).

Apply the comparison at depth n_j to every parent edge $q \in Q_{j+1} \mapsto q_j(q)$ and every adjacent pair in Q_j. There are at most $C2^{k_{j+1}}$ such edges, each of parameter length at most h_j. Their total failure probability is bounded by

$$\exp\{Ck_{j+1} + CH_j k_j + \log H_j - A_j\},$$

which is summable. Thus all these comparisons hold eventually almost surely. Formula (37), with $M = S = j$, gives on each retained edge

$$\rho(D_u, D_v) \leq e_j := C(2^{-j} + j^2\delta_j + jr_j), \qquad \sum_j e_j < \infty. \tag{41}$$

Completeness now supplies limits $F_t = \lim_j D_{q_j(t)}$, uniformly in $t \in I$ for ρ. If $|s-t| < h_j$, their level-j projections coincide or are adjacent; (41) and its summable tail therefore prove continuity of $t \mapsto F_t$. The approximants are jointly measurable step functions, so the limit is jointly measurable. Define F_t to be the Euclidean metric off the common probability-one event.

For each fixed deterministic t, (38) implies $\rho(D_{q_j(t)}, D_t) \to 0$ in probability. To see this, first fix δ, M, S, take $n \asymp [H(\delta)|q_j(t) - t|]^{-1}$, and apply (37); then send $\delta \downarrow 0$ and $M, S \to \infty$ in that order. Equality of limits in probability gives

$$\mathbb{P}[F_t = D_t] = 1 \qquad \text{for every fixed deterministic } t. \tag{42}$$

This is a modification assertion, rather than simultaneous equality with arbitrary separately selected versions at every parameter.

7.3 Backward recovery, positivity, and triangle inequalities

Let $P_j = \prod_{i \geq j}(1 + \delta_i)$, so $P_j \downarrow 1$. Composing the reciprocal comparisons backwards gives, simultaneously for every t, x, y with $F_t(x, y) < \infty$, points x_j, y_j such that

$$|x_j - x|, |y_j - y| \leq s_j, \qquad D_{q_j(t)}(x_j, y_j) \leq P_j(F_t(x, y) + s_j). \tag{43}$$

Indeed start from an epigraph recovery pair at level N, compose to level j, and then let $N \to \infty$. The accumulated displacement is at most s_j, and the additive cost is at most $P_j s_j$. Compactness of the endpoint neighborhoods and lower semicontinuity of $D_{q_j(t)}$ give (43). For any fixed spatial compact set the comparison windows contain all these endpoints once j is large enough.

For each fixed spatial radius R, Proposition 2.1 gives an exponent χ_R and an inverse spatial constant with a polynomial tail. A union bound over Q_j and Borel–Cantelli therefore give, with c_R sufficiently large,

$$D_\theta(x, y) \geq e^{-c_R k_j}(|x - y| \wedge 1)^{\chi_R}, \quad \theta \in Q_j, \quad x, y \in B_R, \tag{44}$$

eventually. Intersect over integer R. If $x \neq y$ but $F_t(x, y) = 0$, equations (43)–(44) contradict $A_j/k_j \to \infty$. Thus every F_t is positive off the diagonal.

We next use the simultaneous all-arm conclusion of [1]. For u small, outside probability $C_R u^q$, every center in B_R has at least $c \log(1/u)$ pairwise disjoint good annuli between B_{2u} and B_{u^α}, for some fixed $\alpha, q > 0$. Each joins any two actual crossing arms at cost at most C times their combined metric lengths. Union over Q_j with $u = s_j$: the failures $C_R 2^{k_j} s_j^q$ are summable by (40).

Fix t and distinct x, y, z with $S = F_t(x, y) + F_t(y, z) < \infty$. Recover the two legs by (43) in $D_j = D_{q_j(t)}$, and take paths of lengths $a_j \to F_t(x, y)$ and $b_j \to F_t(y, z)$. Their two endpoints near y lie inside every good annulus, and their outer endpoints eventually lie outside all of them. In one of the $J_j \geq c \log(1/s_j)$ annuli, the combined lengths of the two crossing arms are at most $(a_j + b_j)/J_j$. Trim at the selected contacts and insert its connector. The result gives

$$D_j(x_j, z_j) \leq (a_j + b_j)(1 + C/J_j) + o(1).$$

The epigraph lower-limit inequality proves $F_t(x, z) \leq S$. Coincident points and infinite right sides are immediate. The events were simultaneous before the triple and parameter were chosen, so every F_t is an extended metric on this one event.

7.4 Continuity at deterministic points

Fix any deterministic countable set $Z \subset \mathbb{R}^d$, enlarged to contain $0, e_1$. Uniform point moments give, for $z \in Z$, some $p, \kappa > 0$ such that $\mathbb{E}[D_\theta(z, z + re_1)^p] \leq C_z r^\kappa$ for small r. Choose $0 < \beta < \alpha\kappa/p$. A reference path from z to $z + 2u^\alpha e_1$ has length at most u^β, outside probability $C_z u^{\alpha\kappa - \beta p}$. Combined with the preceding annuli, this gives, simultaneously for every finite path P meeting $B_u(z)$ and leaving $B_{u^\alpha}(z)$,

$$D_\theta(z, P) \leq C_z u^\beta + \frac{C_z}{\log(1/u)} \operatorname{len}(P; D_\theta), \tag{45}$$

outside probability $C_z u^{b_z}$, for some $b_z > 0$. Indeed choose a good annulus whose P-arm costs at most $\mathrm{len}(P)/J$ and attach it to the reference arm, then to z. This is a universal-path event; no length cutoff for P is required.

Impose (45) at scale r_j for every $\theta \in Q_{j+1}$, and at scale s_j for every $\theta \in Q_j$. Their failures are summable by $A_j/k_{j+1} \to \infty$. Intersect over Z. For a fixed pair $z \neq w$ in Z, attach the displaced endpoints in each edge comparison back to z, w. This gives reciprocal inequalities

$$D_u(z,w) \le (1+\epsilon_j) D_v(z,w) + b_j, \qquad D_v(z,w) \le (1+\epsilon_j) D_u(z,w) + b_j,$$

where $\sum_j \epsilon_j + \sum_j b_j < \infty$; one may take $\epsilon_j = (1+\delta_j)(1+C/A_j) - 1$ and $b_j = C(r_j + r_j^{\beta_z \wedge \beta_w})$. All initial finite-grid point distances are finite almost surely. The parent inequalities give a finite random uniform bound for later grid values by iterating $M_{j+1} \le (1+\epsilon_j) M_j + b_j$. The reciprocal inequalities consequently become summable absolute bounds. The parent and adjacent-edge argument gives a finite continuous limit $X_t(z,w)$, uniformly in t.

The epigraph lower-limit inequality gives $F_t(z,w) \le X_t(z,w)$. Conversely apply (43) and then (45) at scale s_j to its recovering path:

$$D_{q_j(t)}(z,w) \le (1+o(1)) P_j(F_t(z,w) + s_j) + o(1).$$

Thus $X_t(z,w) \le F_t(z,w)$. We have proved, simultaneously for all t and $z, w \in Z$, that $F_t(z,w)$ is finite and continuous in t.

7.5 Uniform Euclidean confinement and geodesics

The large-scale crossing estimates in Proposition 2.1 give, for some $p, q > 0$,

$$\mathbb{P}[X_\theta(R) < L] \le C L^q R^{-p} \quad (R \ge 2,\ L \ge 1), \tag{46}$$

where $X_\theta(R)$ is the crossing length of $A(R/4, R/2)$. Choose the earlier C_0 with $C_0 p > \log 2 + 2$. Union over Q_j in (46), with $R = R_j$ and $L = j$, is summable. Hence eventually $X_\theta(R_j) \ge j$ for every $\theta \in Q_j$.

Fix such a j. For $k \ge j$, compose comparisons from $q_k(t)$ back to $q_j(t)$ for pairs $x \in B_{R_j/8}$, $y \in \partial B_{3R_j/4}$. The displaced endpoints still bracket $A(R_j/4, R_j/2)$, so

$$j \le P_j(D_{q_k(t)}(x,y) + s_j).$$

Thus these distances are at least $j/2$, for large j. Each finer-grid metric is a length metric with Euclidean-continuous finite paths. First passage through $\partial B_{3R_j/4}$ extends the same lower bound to every $y \notin B_{R_j}$. Backward recovery at level k, followed by $k \to \infty$, gives

$$F_t(x,y) \ge j/2, \qquad x \in B_{R_j/9}, \quad y \notin B_{2R_j}, \quad t \in I. \tag{47}$$

In particular finite-radius metric balls are Euclidean compact, with uniform Euclidean confinement for parameters and centers in compacts. We do not infer compactness for the metric topology from this statement.

Continuity for ρ implies joint lower semicontinuity of $(t,x,y) \mapsto F_t(x,y)$. Positivity and compactness therefore give, for every $R, \varepsilon > 0$,

$$\inf\{F_t(x,y) : t \in I,\ x, y \in \overline{B}_R,\ |x-y| \ge \varepsilon\} > 0. \tag{48}$$

The countable grid metrics equal their modification simultaneously, so this is also their common inverse Euclidean modulus.

Fix $F_t(x,y) = L < \infty$. Recovery gives endpoints x_k, y_k with $D_{q_k(t)}(x_k, y_k) \to L$. Take almost-minimizing paths of lengths $L_k \to L$, parametrized on $[0,1]$ at constant metric speed.

Their images stay in one Euclidean compact set by the preceding grid barriers. Equation (48) gives Euclidean equicontinuity, so a subsequence converges uniformly to a curve γ. For $u, v \in [0,1]$, joint lower semicontinuity gives $F_t(\gamma(u), \gamma(v)) \le L|u-v|$. The triangle inequality and $F_t(x,y) = L$ force equality. Thus every finite pair, at every parameter, has a constant-speed Euclidean-continuous geodesic.

7.6 Median normalization and the modification quantifier

Lemma 7.2. *For each fixed parameter, the distance between* 0 *and* e_1 *has an atomless law with full support in* $(0,\infty)$.

Proof. Choose a smooth nonnegative pin-preserving Cameron–Martin function f, supported inside $B_{1/4}(0)$ and equal to one on a smaller fixed shell. Every path from 0 to e_1 crosses that shell. Its original crossing length X_h is strictly positive. For $s_2 > s_1$, Weyl scaling gives the strict bound

$$D_{h+s_2 f}(0,e_1) \ge D_{h+s_1 f}(0,e_1) + (e^{\psi(s_2-s_1)} - 1)e^{\psi s_1} X_h.$$

This follows before taking path infima, so strictness is preserved. Weyl bounds also give continuity in s. Gaussian disintegration along the Cameron–Martin direction f therefore makes the distance a strictly increasing function of a non-atomic Gaussian coordinate, proving atomlessness.

For full support, choose a deterministic thin connected open tube C containing $[0,e_1]$, and a smooth pin-preserving function $g = 1$ on C. A compensating bump on a disjoint cap of the pinning sphere makes the pin average zero. Local rational rectifiable connectivity gives $D_h(0,e_1;C) < \infty$ almost surely, and hence $D_{h-Mg}(0,e_1) \le e^{-\psi M} D_h(0,e_1;C) \to 0$. For each fixed M, adding sf increases this distance continuously to infinity. Given $(a,b) \Subset (0,\infty)$, countably many integer M and nonnegative rational s therefore give shifted events whose union has probability one. One has positive probability. Cameron–Martin equivalence transfers this positive probability to $\{D_h(0,e_1) \in (a,b)\}$. □

Let $m(t)$ be the unique median of $D_t(0,e_1)$. The deterministic-point continuity already proved implies weak convergence of these laws as $t_n \to t$. Lemma 7.2 makes their distribution functions continuous and strictly increasing on $(0,\infty)$. Two thresholds on opposite sides of $m(t)$ trap $m(t_n)$ by weak convergence, proving that m is positive and continuous. Multiplication by a convergent positive scalar preserves both defining conditions of epigraph convergence. Thus $F_t/m(t)$ is an almost-surely continuous family in the prescribed median normalization.

Completion of Theorem 1.1. Apply the construction to a countable cover of the positive-Q region by compact intervals with rational endpoints. On overlaps the modifications agree almost surely at every rational parameter by (42); continuity makes them agree everywhere. A countable intersection patches them into one jointly measurable family on the whole region. Its values are positive extended geodesic metrics, its parameter map is almost surely continuous for lower semicontinuous convergence, and it agrees almost surely with the canonical metric at each fixed deterministic parameter. For any prescribed deterministic parameter sequence and its limit, these identities hold simultaneously, so the original canonical metrics converge almost surely along that sequence. An assertion for all parameter sequences on one event is understood for the continuous modification. □

Sources and status of the inputs

The references labelled PQ, Planar and HD are the supplied project manuscripts. The present theorem is conditional on the fixed-parameter existence, strong characterization and local

estimates in PQ, in the same sense that HD takes its supplied fixed-parameter characterization as an input. The new assertions proved here are the compact-parameter uniformity, controlled-contact replacements for boundary diameter arguments, quantitative protected carriers and companions, and the resulting all-parameter comparison and completion.

References

[1] *Exponential metrics in higher dimensions: Tightness and uniqueness for positive Q*. Supplied project manuscript, `Higher_dimensional_LQG_tightness_and_uniqueness.pdf`, 188-page version inspected for this proof. In particular Sections 4, 6–12, 17–18, 21–25, 30–31, 46–47 and 50–54.

[2] *Almost sure continuity of the LQG metric*. Supplied project manuscript, 38-page version. Project manuscript (PDF).

[3] *Almost sure continuity of the higher-dimensional LQG metric*. Supplied project manuscript, 34-page version. Project manuscript (PDF).

[4] R. Engelking, *Dimension Theory*, North-Holland, 1978. Theorem 1.8.19 (Mazurkiewicz nonseparation theorem). Book (PDF).

[5] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002. Section 3.3, Alexander duality and its Čech-cohomology extension for arbitrary compacta. Author's edition (PDF).