Analyticity of the LFPP exponent for all positive parameters

Abstract

We prove that the Liouville first-passage percolation scaling exponent is real analytic at every positive parameter. The argument uses the canonical metric at a fixed real parameter and changes finitely many covariance scales through Gaussian likelihoods. Local geodesic replacement gives a contracting analytic normalization, while a separate marked expansion controls endpoint and macroscopic terms. Small positive moments of deterministic endpoint distances suffice. A real comparison on both sides of the base parameter identifies the resulting holomorphic limit with the exponent. The geometric arguments use short surrounding loops and lower semicontinuity, and apply when the metric has singular points.

Contents

1 Introduction . . . 2
2 The metric at a fixed parameter . . . 4
3 Compact sources and coherent mean variables . . . 8
4 Original coordinates and conditional trial estimates . . . 13
5 Paths near deterministic circles . . . 17
6 Exact traces in an unchanged region . . . 18
7 The pure occupation reference . . . 21
8 Gaussian comparison on strict buffers . . . 22
9 Fixed-parameter geometry and source forecasts . . . 24
10 The annular common submeasure . . . 31
11 Initial capture and uniform restart . . . 36
12 Finite-chaos moments at the actual root . . . 38
13 Boundary cells, Cauchy convergence, and scalar descent . . . 38
14 The precise interface with a structural activity norm . . . 42
15 Positive refinement of complete owners . . . 43
16 A source-free atlas of positive serial arms . . . 46
17 Scalar capture and complete currents . . . 52
18 Finite port bodies . . . 59
19 A fresh physical realization . . . 64
20 The paired local replacement . . . 67
21 Complete activities and their estimates . . . 78
22 Uniform analytic normalization . . . 93
23 Gaussian identities for the torus observation . . . 98
24 A small Mellin terminal and its logarithmic derivative . . . 102
25 The real pressure on both sides of the base parameter . . . 110
26 Passage to the analytic pressure . . . 115
A Finite target exploration and typical geodesics . . . 115
B Geometric estimates at every positive parameter . . . 121
References . . . 151

1 Introduction

Let h be a whole-plane Gaussian free field, with covariance singularity $\log(1/|x-y|)$ and a fixed additive normalization. Write h_ϵ for its heat regularization at scale ϵ. For $\xi > 0$, the Liouville first-passage percolation distance is

$$D_\epsilon^\xi(x,y) = \inf_{P:x\to y} \int_0^1 e^{\xi h_\epsilon(P(t))} |P'(t)|\, dt, \tag{1.1}$$

where the infimum is over piecewise continuously differentiable paths. Let $a_\epsilon(\xi)$ be the lower median of the left-right crossing distance of the unit square, with paths restricted to the square. The exponent λ is defined by

$$a_\epsilon(\xi) = \epsilon^{\lambda(\xi)+o(1)}, \qquad \lambda(\xi) = 1 - \xi Q(\xi). \tag{1.2}$$

It exists for every $\xi > 0$, and $Q(\xi) > 0$ throughout this interval [12, Proposition 1.1].

Theorem 1.1. *The function λ is real analytic on $(0,\infty)$. For every compact interval $J \Subset (0,\infty)$ there are constants $C_J, R_J > 0$ such that*

$$\sup_{\xi\in J} |\lambda^{(m)}(\xi)| \le C_J R_J^{-m} m!, \qquad m = 0,1,2,\dots. \tag{1.3}$$

Consequently $Q(\xi) = (1-\lambda(\xi))/\xi$ is real analytic on the same interval.

The proof works at one fixed real parameter $a > 0$. Put $q(a) = aQ(a)$. The canonical parameter-a metric exists for every $a > 0$ [4]. It is an extended length metric: when $Q(a) < 2$, some points are at infinite distance from every other point. We use finite geodesics, crossing and surrounding-circuit estimates, and small positive moments of deterministic endpoint distances. No upper bound for the diameter of a Euclidean neighborhood is needed.

1.1 Outline of the proof

We change the covariance of finitely many scales of the field and study the expected logarithm of a fixed torus distance. For real z near zero, the changed scales have amplitude $t = \sqrt{a^2+z}$, while the finer scales retain amplitude a. The metric is always evaluated at a real field. Complex values of z enter through Gaussian likelihoods.

Let $r_j = r_0 b^{-j}$, where $b > 1$ is fixed during the construction. We obtain holomorphic functions ℓ_N on a disk about zero, independent of N, whose real values are these expected logarithmic distances. The two estimates needed at the end are

$$\sup_{|z|<\rho} |\ell_N(z)| \le C(N+1) \tag{1.4}$$

and, for real z in that disk,

$$\lim_{N\to\infty} \frac{\ell_N(z) - \ell_N(0)}{N\log b} = q\big(\sqrt{a^2+z}\big) - q(a). \tag{1.5}$$

The bound gives a normal family. The values on the real interval determine every subsequential holomorphic limit, which proves analyticity near a. Since a is arbitrary, this proves the theorem. In particular, no uniform construction as the base parameter varies is required.

The main work is to obtain (1.4) on one disk for arbitrarily many scales. We explain its four parts.

Ordinary geodesic geometry. First we establish the geometric statements needed for local replacement at the fixed parameter. An exploration aimed at a finite target has finitely many extremal prefixes between two positive levels. This gives protected subsegments of geodesics between deterministic endpoints. A local forcing argument produces such subsegments along other geodesics. Short surrounding loops replace the upper metric continuity estimates which are unavailable in the supercritical regime. They also supply convergence of endpoint distances whenever we take a geodesic limit. The resulting confluence and interior uniqueness imply that an unchanged annulus has only finitely many complete geodesic traces. A separate separated-annulus argument gives moments of every order for their number.

A reference sampled from metric length. The scalar effect of a compact Gaussian source is measured by sampling its root according to geodesic length. An annular comparison matches the entire normalized inner field and the rooted occupation measure with positive probability. Its acceptance rule uses only revealed data. Repeating the comparison gives a linear scalar functional Π on complete physical source expressions. We define this functional through bounded distance tests and then pass to a limit on their first currents. This avoids assuming a finite first moment for a point distance.

To estimate the centered expression $P - \Pi(P)D_1$, we retain the full vector of local port costs. Three independent ordinary serial portions smooth the captured cost. A finite polytope decomposition then preserves both the scalar action and its whole first current. We realize the resulting rows by actual curves in a fresh local field. Compact lowerings approach the conditional essential infima of the ordinary row costs, and surrounding circuits give the required comparisons. Division by the fresh density is performed on the entire convolution support and separately for each retained exterior state.

A contracting analytic map. Compact covariance sources are estimated jointly, with all original Gaussian coordinates and previously read observations retained. Every estimating term uses one current factor and value representations for its other factors. Thus each estimate refers to one ordinary occupation measure. Positive spatial assignments preserve the scalar drift exactly. Large supports gain from the spatial weight; bounded supports use the local replacement at the last accepted scale. Its error tends to zero with the age of the source. Convolution with the exponentially decreasing unsuccessful suffix gives a strict contraction for one sufficiently long finite block. Connected-component estimates yield holomorphic scalar increments b_j satisfying $|b_j(z)| \le C|z|$ on a common disk.

Endpoint terms and real identification. Sources which reach an endpoint or cease to fit in a planar chart are kept as marked terms. For endpoint terms, collapsing a small endpoint neighborhood annihilates the protected source derivative. The difference is controlled using a sufficiently small positive endpoint moment and negative moments of an annular crossing away from the other endpoint. Macroscopic terms use an unused spatial weight. A joint estimate before taking any quotient, followed by a marked connected-component sum, controls the remaining attachments. The torus calculation retains the Gaussian constant and every auxiliary mean read by the expression.

We use a small Mellin parameter only to justify differentiation at the logarithmic observation. No zero-free assertion for a complex distance expectation is needed. Finally a real comparison of coarse LFPP paths with fine surrounding circuits proves (1.5) for both $t < a$ and $t > a$. This comparison is independent of the analytic normalization.

1.2 Organization

Section 2 records the all-parameter metric inputs. Sections 3–4 give the Gaussian estimates. The geometric conclusions and their use in the occupation reference are developed in Sections 5–7; their additional all-parameter proofs are given in the geometric appendices. The local source constructions, finite port representation, and physical realization lead to the paired replacement in Section 20. Sections 21 and 22 construct the analytic map. The torus and terminal estimates are in Sections 23 and 24. Section 25 proves the real comparison, and Section 26 completes the proof.

2 The metric at a fixed parameter

Fix $a > 0$ and put $Q = Q(a)$ and $q = aQ$. Throughout the proof the metric parameter is a. The response parameter will change the covariance of a Gaussian field, and may be complex. Every metric appearing in the argument is evaluated at a real field.

We use the canonical parameter-a LQG metric, whose existence and uniqueness for every $a > 0$ follow from [4, Theorems 1.3 and 1.8]. For the physical field $F = ah$, write this metric as M_F^a. It is an extended length metric, lower semicontinuous in the Euclidean topology, positive off the diagonal, and measurable from F. Its internal metric in an open set is measurable from the restriction of F to that set. A point is called nonsingular if it has finite distance from another point; any two nonsingular points have finite distance and are joined by a geodesic. Each fixed deterministic point is almost surely nonsingular. These statements, including the convention for infinite distances, are as in [4, Definition 1.7, Lemma 1.16, and Section 2.4].

A physical mean u is added to F, so its Weyl factor is e^u. The Weyl and dilation rules are

$$M_{F+u}^a = e^u \cdot M_F^a, \qquad M_F^a(rx, ry) = r^q M_{F(r\,\cdot)}^a(x, y). \tag{2.1}$$

The first equality means integration of e^u against the M_F^a-length of each path, followed by minimization. Both rules include the additive normalization of the field. The exponent satisfies $Q(a) > 0$ for every finite $a > 0$ by [12, Proposition 1.1]; hence $q > 0$.

For a smooth compact physical direction u and a cylindrical test Ψ, write

$$D_u\Psi(F) = \frac{d}{dt}\Psi(F + tu)\Big|_{t=0}.$$

For a symmetric finite-rank covariance direction K, the notation D_K^2 denotes contraction of the Hessian with K. Thus $D_K^2 = \sum_{i,j} K_{ij}\partial_i\partial_j$ in finite Gaussian coordinates. For metric tests, higher coefficients are defined by differentiating their Gaussian expectations through a likelihood. We use only first variations of individual distances.

Crossing and endpoint estimates. Let $h_r(z)$ be the circle average of h and set

$$s_r(z) = r^q e^{ah_r(z)}.$$

For any fixed bounded annulus with nontrivial boundary components, its normalized crossing distance and normalized shortest surrounding circuit have moments of every positive and negative order. The corresponding statement holds for internal distances between two disjoint deterministic connected compact sets which are not singletons, in a fixed open neighborhood of those sets. These bounds are uniform under translation and dilation; see [4, Lemma 2.6 and the paragraph following

it]. We always take a surrounding rectifiable circuit with length at most twice the infimum plus $s_r(z)$. Its normalized length consequently has every positive moment.

Fixed endpoint distances require a smaller moment range. For every $p \in (0, 2Q/a)$, a deterministic point in a disk can be joined to a prescribed deterministic point on its boundary with normalized internal distance having finite pth moment, uniformly under translation and dilation [4, Lemma 2.9]. A finite chain of such disks gives the corresponding bound for any fixed pair of deterministic endpoints in a connected open domain. We use this estimate only with an explicitly chosen $p < 2Q/a$. We do not assume that a point distance has a finite first moment, or that a Euclidean compact set has finite metric diameter.

Finite surrounding circuits give two elementary comparisons used repeatedly below. If P is a geodesic whose endpoints lie outside a circuit Γ, then its segment between the first and last contacts with Γ has length at most $\mathrm{len}(\Gamma; M_F^a)$. This follows by replacing that segment by an arc of Γ. Also, every curve joining points on opposite sides of an annulus pays its crossing distance. These statements apply to internal geodesics whenever the replacement circuit lies in the domain of minimization.

Geodesic compactness. For each fixed pair of disjoint deterministic Euclidean compact sets, the geodesic between them is almost surely unique [4, Lemma 2.7]. We use this on a countable collection of deterministic endpoint pairs. Every finite geodesic is Euclidean continuous, and the identity from the nonsingular metric space to the Euclidean plane is locally Hölder continuous [3, Proposition 2.11]. Every metric-bounded set is Euclidean bounded [3, Lemma 2.7]. Lower semicontinuity and these estimates give compactness of bounded-length geodesics in a fixed bounded Euclidean region, as formulated in [3, Lemma 5.5]. In particular, if their endpoints converge in the metric and their lengths converge to the distance of the limiting endpoints, every subsequential Euclidean limit is a geodesic between those endpoints.

Every nonsingular point admits surrounding rectifiable loops whose Euclidean diameters, lengths, and distances from the point tend to zero [3, Lemma 3.2]. This is the form of local access used when an endpoint is random. The geometric appendices prove the simultaneous interior uniqueness and confluence statements required below from these inputs.

Throughout the plane we use

$$\mu(dx) = \frac{dx}{2\pi}, \quad \widehat{f}(q) = \int e^{-iq\cdot x} f(x)\, \mu(dx), \quad \langle f, g\rangle_\mu = \int fg\, d\mu\,.$$

The Fourier transform is unitary with this convention. The covariance multiplier of F on tests of total mass zero is $a^2/|q|^2$. Its Cameron–Martin space is the Dirichlet space modulo constants with norm

$$\|u\|_{\mathcal{H}_a}^2 = a^{-2}\langle \nabla u, \nabla u\rangle_\mu\,.$$

Every integration by parts on a degree-one metric observation is conditional on an exterior that includes its scalar pin. The source direction is compact inside that exterior, and therefore preserves the pin.

Lemma 2.1 (Local Gaussian scores)**.** *Let u, v be smooth directions compactly supported in a Dirichlet domain whose exterior contains the pin, and let Φ be a cylindrical test with integrable derivatives. Put*

$$I_v(F) = a^{-2}\langle F, -\Delta v\rangle_\mu\,.$$

Under the pinned pure law,

$$\mathbb{E}[I_v\Phi] = \mathbb{E}[D_v\Phi], \qquad \mathbb{E}[(I_uI_v - (u,v)_{\mathcal{H}_a})\Phi] = \mathbb{E}[D_uD_v\Phi]\,.$$

The score is unchanged by adding a function harmonic on a neighborhood of its support. Conditional on the exterior of a domain containing that support, it has the same centered Gaussian law as the score of the zero-boundary GFF in that domain.

Proof. Finite-dimensional Gaussian integration by parts has score $C^{-1}v$. At the pure covariance it is $a^{-2}(-\Delta)v$. Differentiating the first identity in direction u contributes $D_u I_v = (u,v)_{\mathcal{H}_a}$, giving the second. Approximation by Gaussian coordinates proves the stated integrable versions. If m is harmonic near $\operatorname{supp} v$, then $\langle m, -\Delta v\rangle_\mu = 0$. Apply this to the harmonic term in the domain Markov decomposition to obtain the conditional assertion. □

The pin condition matters. For the stationary field of covariance $a^2(1-e^{-T^2|q|^2})/|q|^2$, the same compact score instead gives $\mathbb{E}[I_v\Phi] = \mathbb{E}[D_{(I-e^{T^2\Delta})v}\Phi]$. That cutoff law is not used for the reference, the local Gaussian adjoints, or the local replacement. The compact radial bands below are covariance directions at the pure law. They are not treated as a separate field with the pure law's precision.

Fix one common lattice phase, independent of the field. All source grids, positive assignments, reference comparisons, and copied inner grids use that phase. A private random phase introduced for a finite geometric selection is a separate auxiliary variable. Physical source expressions are always integrated over any such private variables before they are declared complete.

Lemma 2.2 (Local construction and gluing on the torus)**.** *Let H^0 be the mean-zero GFF on the flat unit torus, let $\zeta \sim N(0, v_0)$ be independent, with $v_0 > 0$, and put $F = a(H^0 + \zeta)$. The canonical planar metric determines a measurable extended length metric M_F^a on the torus. Its internal metrics agree with the canonical local metric on every strict coordinate chart. It is positive off the diagonal and lower semicontinuous for the Euclidean topology. Every fixed deterministic pair has finite distance almost surely, and every pair at finite distance has a minimizing path. The metric satisfies simultaneous Weyl scaling by continuous functions.*

Proof. Here is a local absolute-continuity argument which includes the constant mode. The Cameron–Martin quadratic form of F is

$$a^{-2}\int_{\mathbb{T}^2} |\nabla v|^2\, d\mu + (a^2 v_0)^{-1}\bar{v}^{\,2}, \qquad \bar{v} = \int_{\mathbb{T}^2} v(x)\, dx.$$

Here $d\mu = (2\pi)^{-1}dx$, as in the planar convention. For a smooth disk U in a coordinate chart, Gaussian conditioning on the field outside U restricts this form to $H_0^1(U)$. The centered conditional innovation is therefore equivalent to $a\,\mathrm{GFF}_0(U)$: its density with respect to the latter is a positive constant times

$$\exp\{-\bar{\eta}^{\,2}/(2a^2 v_0)\}.$$

The linear functional $\bar{\eta} = \int_U \eta$ is a well-defined Gaussian variable. This formula follows first by finite Gaussian projection and then by convergence of that one linear functional. The conditional mean is smooth in the interior of U: its distributional Laplacian there is constant, by orthogonality to $H_0^1(U)$ in the displayed quadratic form. On $V \Subset U$, a cutoff of this mean is a Cameron–Martin function. Consequently the conditional, and hence the unconditional, law of $F|_V$ is mutually absolutely continuous with the restriction of $a\,\mathrm{GFF}_0(U)$.

The same restriction is mutually absolutely continuous with that of a whole-plane GFF pinned outside U. Indeed, its domain Markov decomposition on U consists of an independent zero-boundary field and a harmonic function; a cutoff of the latter on V is again a Cameron–Martin

function. Conditional equivalence in both constructions implies equivalence of their mixtures. We may therefore apply the measurable canonical planar internal-metric map on V.

Choose a countable basis of strict coordinate disks. For each disk, realize its local field as the restriction of a planar field whose law is absolutely continuous with respect to the pinned whole-plane law. For example, keep the latter's regular conditional exterior law given the prescribed restriction. Its global canonical metric is lower semicontinuous. Locality and translation covariance imply that these metrics induce the same path lengths on every common smaller disk, outside one null set. An admissible path is locally continuous for these metric structures. Define its length using these local lengths, and define M_F^a by the infimum over admissible paths. A compact path has a finite subdivision subordinate to the disks, so its length is independent of the subdivision and of the planar extensions. The internal metric on a strict chart is then the canonical one there. The local length functional and the condition of local metric continuity are jointly measurable on the space of Euclidean-continuous paths. This follows by using rational partitions and rational moduli of continuity; lower semicontinuity of the planar metric extends their inequalities to all parameter values. The infimum is universally measurable, and hence measurable on the completed Gaussian probability space.

We spell out the required compactness, since no finite Euclidean metric diameter is being asserted. For every fixed $\delta > 0$, a finite collection of chart annuli detects every Euclidean displacement of at least δ. Their crossing distances have a positive minimum almost surely. Taking $\delta = 2^{-m}$ gives a common modulus under which small metric length implies small Euclidean displacement. Bounded-length paths parametrized by length are therefore equicontinuous in the Euclidean topology. For a limiting path piece compactly inside a strict chart, use the global lower semicontinuous canonical metric of one planar extension. Its finite-partition lengths are lower semicontinuous under Euclidean uniform convergence. Taking their supremum and then a finite subdivision into chart pieces gives lower semicontinuity of the glued length. For paths parametrized by length, the same comparison on each chart piece shows that the limit is locally metric-Lipschitz. Thus the limit is admissible; Euclidean continuity alone would not suffice here. Arzelà–Ascoli now gives a minimizer at every finite distance and lower semicontinuity of the distance as the endpoints vary. The same annular lower bound proves positivity off the diagonal. A fixed chain of chart disks with deterministic intermediate endpoints, and the deterministic-endpoint moment estimate, gives finiteness for each fixed deterministic pair.

The planar Weyl axiom is simultaneous for all continuous functions [4, Definition 1.7, Axiom III]. Applying it on the countable charts and then to path lengths proves the same assertion for the glued metric. □

Lemma 2.3 (Torus first variations at deterministic endpoints)**.** *For fixed distinct deterministic p, q, and each prescribed smooth physical direction χ, the function $t \mapsto M_{F+t\chi}^a(p,q)$ is differentiable at $t = 0$ almost surely. Every M_F^a-minimizer P from p to q then has the same first variation, equal to $\int_P \chi \, d\ell_F$. One may choose a single null set for a countable family of directions. In particular the global geodesic for a fixed pair is almost surely unique, if uniqueness is needed.*

Proof. Weyl comparison and finiteness make the displayed function locally Lipschitz. Rademacher's theorem and Fubini give differentiability at almost every t, for almost every field. Translation by $t\chi$ is a Cameron–Martin equivalence for F, including its constant mode; hence differentiability holds at $t = 0$ almost surely. For any minimizing path P, Weyl scaling gives

$$M_{F+t\chi}^a(p,q) \le \int_P e^{t\chi} \, d\ell_F, \qquad D'_+(0) \le \int_P \chi \, d\ell_F \le D'_-(0).$$

Differentiability proves the claim. If one applies it to a countable uniformly dense family of smooth functions on the torus, all minimizing paths have the same finite occupation measure. Its support is the path trace. A minimizing path between distinct endpoints is simple, since the metric is positive off the diagonal; two simple arcs with that trace and those endpoints coincide after orientation and length parametrization. □

For the compact fine field used in Section 25, the global Fourier equivalence proved there puts its law in the same torus metric class. Thus its metric is a measurable function of that rough field itself. If $F = U + S$ is the actual decomposition there, with S smooth, simultaneous Weyl scaling gives

$$M_U^a = e^{-S} \cdot M_{U+S}^a.$$

The equality permits random S; the primary axiom already holds for every continuous shift on one event of probability one. This is the only inverse-Weyl construction required below. In particular an independent smooth complement remains independent of M_U^a.

3 Compact sources and coherent mean variables

All Gaussian adjoints in this section are taken at the pure parameter-a field. In particular, the Cameron–Martin space is

$$\mathcal{H}_a = \dot{H}^1(\mathbb{R}^2), \qquad \|u\|_{\mathcal{H}_a}^2 = (2\pi a^2)^{-1} \int |\nabla u|^2 .$$

Only compactly supported directions are used. Gaussian adjunction is performed conditionally in a Dirichlet domain containing the complete carrier; any scalar pin is supported in the exterior of that carrier. Thus the pin annihilates each physical source direction. The energy estimates below also make sense modulo constants, but the scalar mode of a pinned metric observation is retained. The complete constant drift is assembled as an identity of physical operators and currents. Its torus contribution is treated in Section 23.

Fix a dyadic integer $b_0 > 1$, put $r_j = b_0^{-j}$, and use the nonnegative spline partition $\{\psi_{j,z}\}_z$ constructed in Section 15. Its finite smoothness order is chosen above every derivative order used here. Its supports have radius Cr_j, bounded overlap, and uniformly bounded rescaled derivatives of those orders. Identities first written for smooth directions pass to this fixed partition by approximation in the indicated Cameron–Martin and spatial norms. For ϕ as in Section 25, define

$$\mathcal{G}_j = L^2(\mathbb{R}^2 \times [r_j/b_0, r_j], dy\, dr/r), \qquad T_j f(x) = \int \phi_r(x-y) f(y,r)\, dy\, \frac{dr}{r}, \qquad C_j = T_j T_j^* .$$

For $s = (j, z)$ choose a compact input cutoff η_s equal to one on every input that influences $\operatorname{supp} \psi_s$. Set

$$K_s = \tfrac{1}{2}(M_{\psi_s} C_j + C_j M_{\psi_s}), \qquad U_s^{\mathrm{ph}} = M_{\psi_s} T_j M_{\eta_s}, \qquad V_s^{\mathrm{ph}} = T_j M_{\eta_s} . \tag{3.1}$$

Here K_s is a symmetric signed covariance direction. Both of its physical variables have compact support in a fixed enlargement of the cell.

Proposition 3.1 (One joint Gaussian estimate)**.** *After whitening in $\mathcal{H}_a$, put*

$$R_s = 2^{-1/2}[U_s, V_s], \qquad J(f,g) = (g,f), \qquad A_s = R_s J R_s^*, \qquad S_s = R_s R_s^* .$$

Then

$$\sum_z K_{j,z} = C_j, \qquad \sum_s S_s \le CI, \qquad \operatorname{Tr} S_s \le C\,. \tag{3.2}$$

The synthesis of the deterministic mean directions ψ_s has the same bounded operator norm. For a finite set $\mathcal{S}$ of m physical slots, let

$$A(q) = \sum_{s\in\mathcal{S}} q_s A_s, \qquad h(t) = \sum_{s\in\mathcal{S}} t_s \psi_s\,.$$

For every fixed $1 < p < \infty$ there are $\rho_q, \rho_t > 0$, independent of m, such that the likelihood $L_{A(q),h(t)}$ is L^p-valued holomorphic on a neighborhood of the closed product polydisc with these radii and

$$\sup \|L_{A(q),h(t)}\|_{L^p} \le C_p^m\,. \tag{3.3}$$

Consequently its coefficient of $q^\alpha t^\beta$ has norm at most

$$C_p^m \rho_q^{-|\alpha|} \rho_t^{-|\beta|}\,.$$

The two coordinates at a physical slot are distinct axes; repeated occurrences of one axis are combined before extraction.

Proof. The choice of η_s gives $M_{\psi_s} T_j (I - M_{\eta_s^2}) = 0$, and hence

$$K_s = \tfrac{1}{2}\{U_s^{\mathrm{ph}} (V_s^{\mathrm{ph}})^* + V_s^{\mathrm{ph}} (U_s^{\mathrm{ph}})^*\}\,.$$

Summing the partition proves the first identity in (3.2).

For either physical feature map W_s, direct differentiation gives

$$\|\partial_x^\alpha W_s(x,\cdot)\|_{\mathcal{G}_j} \le C_\alpha r_j^{-|\alpha|} \mathbf{1}_{\{|x-z|\le Cr_j\}}\,. \tag{3.4}$$

The same bound holds for ψ_s with a one-dimensional feature space. If $r \le R$, integration by parts onto the coarse feature and (3.4) give

$$\|W_{R,z}^* (-\Delta) W_{r,w}\|_{\mathrm{HS}} \le C(r/R)^2 \mathbf{1}_{\{|z-w|\le CR\}}\,.$$

A coarse cell meets $O((R/r)^2)$ fine cells and a fine cell meets $O(1)$ coarse cells. The block Schur bound is therefore Cr/R. Summing the bounded scale matrix $b_0^{-|i-j|}$ proves the synthesis estimate. Integrating the squared first-derivative row bound over its support proves the individual trace bound.

Write $R = (R_s)_{s\in\mathcal{S}}$ and let $D(q)$ act by $q_s J$ on the corresponding feature block. Then

$$A(q) = RD(q)R^*, \qquad \|R\|^2 \le C, \qquad \|R\|_{\mathrm{HS}}^2 \le Cm\,.$$

Thus $\|A(q)\| \le C\rho_q$, $\|A(q)\|_1 \le C\rho_q m$, and $\|h(t)\|_{\mathcal{H}_a}^2 \le C\rho_t^2 m$. No operator inequality $|A_s| \le S_s$ is used.

In a finite-dimensional projection of the whitened space the likelihood equals

$$\det(I+A)^{-1/2} \exp\Big\{ \tfrac{1}{2} X^{\mathsf{T}} A (I+A)^{-1} X + ((I+A)^{-1}h)^{\mathsf{T}} X - \tfrac{1}{2} h^{\mathsf{T}} (I+A)^{-1} h \Big\}\,.$$

All quadratic forms are complex bilinear, and the determinant branch equals one at zero. The resolvent identity bounds the absolute quadratic exponent by $C\rho_q \|R^* X\|^2$. The determinant factor

is at most $\exp(C\rho_q m)$, and the linear Gaussian exponential has the required fixed moment by the mean norm bound. For a positive trace-class covariance S,

$$\mathbb{E}e^{u\|G_S\|^2} = \det(I-2uS)^{-1/2} \le \exp\left(\frac{u\,\mathrm{Tr}\,S}{1-2u\|S\|}\right) \quad (2u\|S\|<1)\,.$$

Apply this with $S = R^*R$, after one fixed Hölder inequality. Choosing the two radii small enough proves a uniform L^{2p} estimate on a slightly larger polydisc.

The finite-dimensional likelihoods form a complex martingale under Gaussian coordinate projection. Gaussian marginalization proves this for real parameters; the finite-dimensional identity theorem proves it for complex parameters. The uniform L^{2p} estimate gives convergence in L^p. Cauchy integration on smaller polydiscs passes through the limit, proving holomorphy and the coefficient estimate. □

3.1 Mean queries on nested domains

Fix a countable deterministic catalogue $\mathcal{D}$ of smooth domains, formed from finitely many shapes by translations and dilations at the b_0 scales. At every radius their overlap is bounded. All domains that will be used are included before sampling the field. Let Q_H be Dirichlet-energy projection onto $H_0^1(H)$. Fix an integer $s_0 \ge 4$ and $0 < \varepsilon_0 < 2$. If S_H rescales a function to the unit reference domain, define

$$\|v\|_{E(H_0)}^2 = \sum_{\substack{H\in\mathcal{D}\\ H\subset H_0}} (r_{H_0}/r_H)^{\varepsilon_0}\|S_H Q_H v\|_{H^{s_0}}^2\,. \tag{3.5}$$

Take the completion of the finite-norm smooth innovations, quotienting by functions harmonic on H_0. The term $H = H_0$ controls the smooth representative on every strict interior region. A harmonic physical shift is retained in the exterior likelihood; this quotient does not remove its physical effect.

Proposition 3.2 (Joint coherent queries)**.** *If $H_i \subset H_0$ are disjoint catalogue domains, then*

$$\sum_i (r_{H_0}/r_{H_i})^{\varepsilon_0}\|Q_{H_i}v\|_{E(H_i)}^2 \le \|v\|_{E(H_0)}^2\,. \tag{3.6}$$

Let T_R collect finitely many physical feature maps and mean directions at scale R, with m_R slots and $\sum_R m_R = m$. The map

$$(f_R)_R \longmapsto \left(Q_H \sum_{R\ge r_H} T_R f_R\right)_{H\in\mathcal{D}} \quad \textit{into} \bigoplus_{H\in\mathcal{D}} E(H)$$

has operator norm at most C and squared Hilbert–Schmidt norm at most Cm. The same assertion holds when the direct Cameron–Martin source target is included in the target space.

Proof. Nested Dirichlet projections satisfy $Q_K Q_H = Q_K$ for $K \subset H$. Multiplication of the K term in $E(H_i)$ by $(r_{H_0}/r_{H_i})^{\varepsilon_0}$ gives exactly its weight in $E(H_0)$. Disjoint children use disjoint subcollections of the catalogue. This proves (3.6), including its extension to the completions.

For a smooth function v, the function $Q_H v$ solves $\Delta Q_H v = \Delta v$ with zero boundary values. Rescaled elliptic estimates and bounded overlap imply, at radius h,

$$\sum_{r_H=h} \|S_H Q_H v\|_{H^{s_0}}^2 \le Ch^2 \sum_{k=0}^{s_0-2} h^{2k}\|D^k\Delta v\|_{L^2}^2\,.$$

The feature row bounds give

$$\|D^k \Delta T_R\| \le C_k R^{-1-k}, \qquad \|D^k \Delta T_R\|_{\mathrm{HS}}^2 \le C_k R^{-2-2k} m_R \,.$$

Thus queries at $h = b_0^{-j} R$ have operator norm $C b_0^{-j}$ and squared Hilbert–Schmidt norm $C b_0^{-2j} m_R$. For fixed j, different source radii have different input spaces and different output radii, so these bounds combine by orthogonal direct sum. Summation over j proves the unweighted assertion.

For a descendant at relative depth j_2 inside a queried carrier at relative depth j_1 below the source, the weight in (3.5) changes the bound to

$$C b_0^{-j_1} b_0^{-(1-\varepsilon_0/2) j_2} \,.$$

A descendant lies in only a bounded number of catalogue carriers of a specified radius. Both geometric sums converge. The operator estimate follows by the triangle inequality; the Hilbert–Schmidt estimate follows by the weighted Cauchy–Schwarz inequality with these geometric factors. Finally include the direct source target using Proposition 3.1. □

In particular, restriction to children of radius at most θr_{H_0} has joint norm at most $\theta^{\varepsilon_0/2}$. This is the strict contraction used at an actual separated conversion. A fine source is not inserted into the smooth mean space of a larger opaque child: it is either disjoint from that child or is consumed before the child is formed. All remaining intersecting sources have radius at least a fixed multiple of the child's radius. The finitely many comparable scales only change the constants above.

3.2 A common positive envelope for the first current

The likelihood estimate also controls a first metric variation, uniformly over its spatial test function. We obtain this bound by taking the supremum of the compact physical direction before integrating, rather than by taking a supremum of separate Gaussian adjoints.

Lemma 3.3 (Raw current envelope)**.** *Fix a finite ordinary moment order p. For a finite coefficient using m original physical slots, choose an occupied coordinate whose physical range annihilates every retained multiplier. We call this the protected coordinate. Its first-current representation has one nonnegative envelope Y_α satisfying*

$$\|Y_\alpha\|_{L^p} \le C_p^m \rho_q^{-|\alpha_q|} \rho_t^{-|\alpha_t|} \,.$$

The same envelope bounds every deterministic original output label of supremum norm at most one. It is measurable from the compact source observations and the original reserve observations retained in the expression.

Proof. Start with finite feature rank. Write

$$C(q) = \Sigma + \sum_s q_s K_s, \qquad h(t) = \sum_s t_s \psi_s \,,$$

and let $L_{q,t}$ be the likelihood relative to $N(0,\Sigma)$. Gaussian integration by parts gives

$$\begin{aligned} \partial_{q_s} \mathbb{E}_{C(q),h(t)} f &= \mathbb{E}_\Sigma \big[L_{q,t}\, Df\big(\tfrac{1}{2} K_s C(q)^{-1}(F - h(t))\big)\big] \,, \\ \partial_{t_s} \mathbb{E}_{C(q),h(t)} f &= \mathbb{E}_\Sigma [L_{q,t} D_{\psi_s} f] \,. \end{aligned}$$

These are initially identities for smooth cylindrical tests and real positive covariances. The fixed-contour likelihood estimates extend them holomorphically. The metric first-variation formula then gives the current, by its dominated approximation. The metric is always evaluated at the real base field.

If Y is a retained selector, the primary range annihilates Y by protection. Thus $D(Yf) = YDf$ in this range; the remaining source contacts stay in $L_{q,t}$. In whitened coordinates write $A_s = R_s J_s R_s^*$ and $A = \sum_s q_s A_s$. The random physical vector is

$$\tfrac{1}{2} R_s^{\mathrm{phys}} J_s R_s^* (I + A)^{-1} X \,.$$

The compact row bound $\|R_s^{\mathrm{phys}}\|_{\mathrm{feature}\to C^0} \le C$, the trace bound $\|R_s\|_{\mathrm{HS}}^2 \le C$, and $\|(I+A)^{-1}\| \le 2$ give all fixed moments of its supremum norm. The deterministic mean contribution is bounded by $C\rho\sqrt{m}$ and is absorbed into C_p^m.

For $\alpha_s > 0$,

$$[z^\alpha] H = \frac{1}{\alpha_s} [z^{\alpha - e_s}] \partial_s H \,.$$

Apply Cauchy's formula on the original contours and take the contour integral of $|L_{q,t}|$ times the supremum norm of the displayed physical vector. This defines Y_α. Hölder and Minkowski give the asserted bound; the factor $1/\alpha_s$ is favorable. A label multiplies this physical current at its original output point, so the same positive envelope applies to all unit labels.

To check locality, collect the original covariance and mean features in R. The resolvent identity

$$(I + RJR^*)^{-1} = I - RJ(I + R^* J R)^{-1} R^*$$

expresses all likelihoods and current scores through $R^* X$. At the pure Dirichlet base a compact feature v has local score $\langle F, -\Delta v \rangle / (2\pi a^2)$. For an original residual covariance $C = I - GG^* \ge \frac{1}{2} I$ as in Proposition 4.1, use

$$C^{-1} = I + G(I - G^* G)^{-1} G^* .$$

Thus residual scores are measurable from the same compact source scores together with the retained original reserve scores. Whitening changes their covariance bounds by at most a fixed factor, independently of the number of reserve coordinates. The local sigma field is preserved under the feature limits in probability. The fixed-moment bounds justify those limits and the metric-test approximation. For completeness, Hilbert–Schmidt approximation gives convergence of $R_s^*(I + A)^{-1} X$ in every fixed L^p in the feature Hilbert space. The row estimate (3.4) makes the physical feature map bounded into C^0, so the physical current directions converge in $L^p(C^0)$. Together with the joint likelihood bound, this justifies passage to infinite feature rank. Approximation of bounded metric tests uses the first-variation identity and the same integrable envelope. □

3.3 Shared auxiliary means and reserved coordinates

For an original residual covariance $C \ge \frac{1}{2} I$, a complex Hilbert–Schmidt map T, and a real standard auxiliary Gaussian family Y, Gaussian integration gives

$$\mathbb{E}_Y L^C_{m+TY}(X) = L^C_{TT^{\mathsf{T}},m}(X) \,.$$

The transpose is complex bilinear. This is an identity of likelihoods; it does not interpret a complex covariance as a probability measure. The real and imaginary parts of the shared mean TY, jointly in the direct source target and all eligible coherent queries, have positive covariance with operator

norm $O(\rho)$ and trace $O(\rho m)$. The determinant estimate above controls their joint quadratic exponential. The same Y must occur in every child and in the exterior.

For disjoint deterministic holes H_i, the original Dirichlet projections are orthogonal. If the reserved directions lie wholly on one side of every retained boundary, their residual covariance commutes with these projections. Conditional on the exterior and the exterior reserve coordinates, the inner residual blocks are independent; adding their own original coordinates reconstructs the ordinary independent Dirichlet fields. In a laminar collection one uses the orthogonal differences $Q_H - \sum_{H' \text{ child of } H} Q_{H'}$, rather than repeating a deepest projection at all its ancestors. The mean likelihood factors according to these orthogonal blocks.

Unused original coordinates may be deleted only under the marginal identity (23.11), after the entire integrand, including its selection record, no longer reads them. Fresh coordinates for a converted representation are introduced after its original numerator has been integrated. They receive no new likelihood for a mean that this numerator has already consumed.

4 Original coordinates and conditional trial estimates

The coordinates used to select a local configuration belong to the original Gaussian numerator. Coordinates used to realize a fresh body are introduced after that numerator is integrated. The distinction is essential for both conditioning and subsequent source insertion.

Proposition 4.1 (Original coordinates on a finite genealogy)**.** *Fix a finite deterministic genealogy of buffered domains $H_{0,v} \Subset H_{c,v}$. Each inner domain contains the complete domains, observations and coordinate supports of its descendants; sibling complete domains are disjoint. At v let G_v synthesize a finite collection of smooth directions supported in $H_{c,v} \setminus \overline{H}_{0,v}$, with $\|G_v\|^2 \le 1/2$. Put $G = (G_v)_v$ and $C = I - GG^*$ in the pure Cameron–Martin coordinates. Then*

$$C \ge \tfrac{1}{2} I, \qquad F = X + GA, \qquad X \sim N(0, C), \quad A \sim N(0, I), \quad X \perp A$$

has the original pure Gaussian marginal. The covariance C commutes with every protected Dirichlet projection in the genealogy. Moreover,

$$\mathbb{E}[A \mid F] = G^* F, \qquad \mathrm{Cov}(A \mid F) = I - G^* G\,.$$

The conditional coordinates from different shells are independent, and each shell's conditional mean is measurable from the field in that shell.

Place every source covariance K and mean v in the residual law $N(v, C + K)$, retaining the original law of A. If A_J is an unread collection of entire local coordinate families and the entire retained integrand factors through $(X + G_J A_J, A_I)$, then deletion replaces the residual by

$$X' = X + G_J A_J \sim N(v, I - G_I G_I^* + K), \qquad X' \perp A_I\,. \tag{4.1}$$

This identity holds on the common source germ and for all its coefficients.

Proof. Distinct shells have disjoint supports, so their Dirichlet ranges are orthogonal. Hence

$$\|GG^*\| = \sup_v \|G_v G_v^*\| \le \tfrac{1}{2}\,.$$

Relative to any protected boundary a reserved direction either belongs to its Dirichlet subspace or vanishes in a neighborhood of the domain. Its rank-one covariance commutes with that projection. Summation proves commutation for C.

The marginal covariance is $C + GG^* = I$. Gaussian regression gives the conditional formulas. Orthogonality of the shell ranges makes G^*G block diagonal. A component of G^*F is the pure score $\langle F, -\Delta g\rangle/(2\pi a^2)$, which is measurable in the support of the corresponding compact direction.

The perturbed marginal has covariance $I + K$ and mean v. Adding the independent unread coordinates to the residual gives exactly (4.1). Equivalently, conditional expectation of the residual likelihood onto $(X + G_J A_J, A_I)$ is its Gaussian likelihood relative to the reduced residual covariance. The assertion for complex sources follows from the finite-dimensional identity theorem and the fixed-moment likelihood convergence of Proposition 3.1. □

Lemma 4.2 (Reconstruction with retained readers)**.** *Let e denote retained exterior data, let (X, A) have its fixed real base augmented Gaussian law, and put $F = X + GA$. If $H = H(e, X, A)$ is an integrable retained multiplier, define*

$$\overline{H}(e, F) = E[H(e, X, A) \mid e, F].$$

For every bounded measurable $q(e, F)$,

$$E[H(e, X, A)q(e, F) \mid e] = E[\overline{H}(e, F)q(e, F) \mid e].$$

Thus reconstruction preserves the pure physical marginal but retains the posterior multiplier $\overline{H}$. It does not delete a previously read coordinate, selector or history. For $p \geq 1$, conditional expectation contracts the corresponding L^p norm; a bound $|H| \leq 1$ is also preserved.

If every coordinate and physical-field observation used by H is supported outside a protected domain D, and the original conditional coordinate kernel given F depends only on $F|_{D^c}$ for these readers, then $\overline{H}$ is exterior to D. Otherwise no exterior measurability is asserted. Unread families may be deleted by (4.1) only when the entire retained integrand has the factorization required there.

For a mean or covariance perturbation, include its single original residual likelihood in H and take this same real-base conditional expectation. Do not replace it by a centered posterior formula for a different Gaussian law. On a holomorphic parameter family with the stated common integrable envelope, reconstruction is consequently a bounded linear operation on the numerator and preserves holomorphy. It introduces no likelihood on a fresh companion.

Proof. The first identity is the tower property, and the norm bounds are conditional Jensen inequalities. In the stated exterior case, Gaussian regression for the original augmentation gives a conditional kernel for all the read coordinates depending only on the exterior field. Integrating H against this kernel proves exterior measurability. The compact score formula in Proposition 4.1 supplies precisely this kernel when the retained directions and readers lie outside D. □

All such augmentations are made only after the deterministic source positions and positive assignment labels have fixed a genealogy. There is no assertion that one full grid of arbitrary bumps commutes with every possible domain. If a large owner does not fit the next fixed-shape trial domain, its trial is postponed. Success and failure retain the same declared envelope.

For one trial path, the exact partition

$$1 = \prod_{j=1}^{n}(1 - \Gamma_j) + \sum_{k=1}^{n}\Gamma_k \prod_{j=k+1}^{n}(1 - \Gamma_j) \tag{4.2}$$

is inserted in the common augmented probability space before any deletion. The kth summand reads no earlier new trial outcome. Because its selected and later observations are outside the

complete earlier modules, the corresponding unread whole coordinate families can then be deleted. Previously read coordinates of an old input remain. A selected configuration may read the residual after all its candidates have been subtracted; its nonselected candidates cannot be deleted while that selection remains.

For a protected pair $H_0 \Subset H_c$, commutation gives independent original residual blocks in

$$\operatorname{Ran} Q_{H_0}, \qquad \operatorname{Ran}(Q_{H_c} - Q_{H_0}), \qquad \ker Q_{H_c}\,.$$

Their mean likelihoods multiply, and their squared score norms add. Reconstructing the inner original coordinates gives an ordinary Dirichlet field independent of the retained exterior and the still-unexposed selected outer coordinates. A covariance perturbation need not commute with the projections; its contacts remain in the one joint likelihood or one shared auxiliary mean.

4.1 A Gaussian product inequality

Lemma 4.3 (Gaussian Hölder inequality)**.** *Let $(Y_1, \dots, Y_m)$ be centered jointly Gaussian, with block covariances Σ_i. Suppose $\operatorname{Cov}(Y) \le K \operatorname{diag}(\Sigma_i)$. For nonnegative measurable f_i,*

$$\mathbb{E} \prod_{i=1}^{m} f_i(Y_i) \le \prod_{i=1}^{m} \{\mathbb{E} f_i(Y_i)^K\}^{1/K}\,. \tag{4.3}$$

In particular, if $\mathbb{P}(E_i) \le \epsilon$, then $\mathbb{P}(\bigcap_i E_i) \le \epsilon^{m/K}$.

Proof. For finite Gaussian blocks this is [11, Theorem 1(i)], with every exponent equal to K, restricting to the linear span of each block if necessary. For field observations, approximate a bounded nonnegative function of each observation by its conditional expectations on an increasing sequence of finite Gaussian coordinates. The covariance inequality survives compression. These approximations converge almost surely and remain uniformly bounded, so dominated convergence gives the asserted inequality for bounded functions. Truncation and monotone convergence give the general nonnegative case. Taking $f_i = \mathbf{1}_{E_i}$ proves the final assertion. □

Proposition 4.4 (Counts in a new conditional query)**.** *Let A_v be a finite family of fixed-shape annular observation regions. Each has a connected fixed-ratio buffer B_v, and the buffers have overlap at most K_0. Observations are restrictions of the pure field modulo a scalar on the corresponding annulus. The original residual coordinates of Proposition 4.1 may be included in their own observation blocks. There is a constant K_1, independent of the family size and depths, such that the joint conditional centered observation covariance inside any deterministic query domain H containing every B_v is bounded above by K_1 times its block-diagonal covariance.*

Suppose the chosen local numerical and geometric event has conditional failure probability at most ϵ in every such block with centered mean in a fixed bounded class. For a conditional query whose centered block means all belong to that class,

$$\mathbb{P}_H\left(\bigcap_{v=1}^{m} \Gamma_v^c\right) \le \epsilon^{m/K_1}\,. \tag{4.4}$$

Fix a prescribed rate $\lambda_f > 0$. Starting from such a query, if a smooth incoming innovation $x \in E(H)$ is added, the stored bound $e^{-\lambda_f m}$ may be retained at a further cost $\exp(C\lambda_f \|x\|^2/M^2)$, provided the ordinary cutoffs give a sufficiently large failure exponent on the enlarged mean class with tolerance M.

Proof. Let P_v be the Cameron–Martin projection onto the annular restriction modulo constants. Choose $\chi_v = 1$ on A_v, supported in B_v, with rescaled derivatives bounded. If $\bar{h}_{B_v}$ is the average of h on B_v, then

$$\|P_v h\|_{\mathcal{H}_a}^2 \le \|\chi_v(h - \bar{h}_{B_v})\|_{\mathcal{H}_a}^2 \le C \int_{B_v} |\nabla h|^2 .$$

Poincaré's inequality proves the second bound. Sum over v to get $\sum_v P_v \le CK_0 I$, equivalently a fixed covariance angle bound for the unconditioned normalized observations.

The same cutoff also proves a lower covariance comparison. Any whole-plane observation can be represented with at most a fixed energy multiple by a zero-boundary function in B_v. By the variational formula for Gaussian covariance, the B_v-Dirichlet observation covariance dominates a fixed multiple of the whole-plane observation covariance. The conditional H covariance dominates this B_v covariance on the diagonal blocks, while its full joint covariance is no larger than the whole-plane covariance. These two statements, together, prove the asserted angle bound after conditioning. Monotonicity alone would not suffice.

The original residual covariance lies between one half and one times the pure covariance, and it commutes with the retained cuts. This changes the comparison constants by a fixed factor. The original coordinate families assigned to distinct shells are independent and are included only once in their local blocks. Their inclusion preserves a fixed angle constant. Apply Lemma 4.3 to the centered observations after subtracting the deterministic means. This proves (4.4).

For the final assertion, write $R = r_H$. After removal of its local scalar, the smooth incoming mean changes an observation of radius r_v by at most $C(r_v/R)\|x\|$. Bounded overlap gives $\sum_v r_v^2 \le CR^2$. At most $N_{\mathrm{spoil}} \le C\|x\|^2/M^2$ observations exceed the tolerance. On all other observations use a failure rate $\tau > \lambda_f$. Keeping the stored factor gives

$$-\tau(m - N_{\mathrm{spoil}}) + \lambda_f m \le \lambda_f N_{\mathrm{spoil}} .$$

This is the stated extra cost. The tolerance is chosen after the stored exponent and before the local ordinary cutoff. Consequently increasing the available ordinary rate does not enlarge the extra real quadratic coefficient. □

The application uses predicates measurable in these same local blocks: the selected finite configuration, its complete original amplitude window, the strict-interior harmonic prediction, local free-row tests, retained middle-score bound, and required length comparisons. The pure domain Markov property makes the interior harmonic prediction measurable from the exterior annular germ. Since each retained boundary is separated from all original coordinate supports on its two sides, the same assertion holds for the residual-selected predicates. Strict-buffer Gaussian comparison transfers an arbitrarily small ordinary failure probability to a uniformly small conditional probability on each fixed bounded centered-mean class.

If a local selection uses an independent private phase, condition on all such phases before applying the product inequality. A phase-dependent failure probability $p_v(\theta_v)$ contributes $p_v(\theta_v)^{1/K_1}$. Independence of the private phases and concavity give

$$\mathbb{E}\prod_v p_v(\theta_v)^{1/K_1} \le \prod_v (\mathbb{E}p_v(\theta_v))^{1/K_1} .$$

Thus a high-probability statement averaged over the private phases is sufficient. The common physical refinement phase remains fixed throughout.

The fixed-moment current transfer is applied once after this joint Gaussian estimate. First work under the ordinary conditional Gaussian law. If a common positive envelope V satisfies $\|V\|_{L^{2p_0}} \le$

$C_{p_0}^m$ and the ordinary failed-trial event has rate at least $2p_0\lambda_f$, then

$$W = C_*^{-m} e^{\lambda_f N_f} \mathbf{1}_{\{\text{all retained trials failed}\}} V \quad \text{satisfies} \quad \|W\|_{L^{p_0}} \le C$$

for C_* fixed sufficiently large, by one ordinary Hölder inequality. Apply Proposition 14.1 to this one envelope. Choosing the ordinary moment and the GOOD count rate with a strict margin gives a rooted exponent $p > 1$. A subsequent exterior event then costs a power of its probability without changing the already stored factor $e^{-\lambda_f N_f}$. Both steps concern the complete record, not its ancestors separately.

5 Paths near deterministic circles

The following fact prevents a bounded-length minimizing path from moving along the limiting boundary of a narrowing ordinary region. It will be applied to circles chosen from a fixed countable catalogue.

Lemma 5.1 (Deterministic circles)**.** *Almost surely, every nonconstant Euclidean-continuous curve contained in a fixed deterministic circle has infinite M_F^a-length. Consequently, if K is a fixed Euclidean compact set containing that circle and $\eta > 0$, the infimum of the lengths of curves in its δ-neighborhood in K, with Euclidean endpoint separation at least η, tends to infinity as $\delta \downarrow 0$.*

Proof. We first check the numerical condition needed for the metric KPZ bound. The exponent bounds in [2, Lemma 1.1 and (2.9)] give

$$q \le \frac{3}{2}, \qquad q \le \frac{3}{4} + \frac{a^2}{2} \quad \text{if } a \ge \frac{1}{\sqrt{6}}.$$

For $a < 1/\sqrt{6}$, monotonicity of Q and $Q(1/\sqrt{6}) = 5/\sqrt{6}$ imply that the parameter is subcritical; see [3, Section 1.1]. The relation $a = \gamma/d_\gamma$ and [8, Theorem 1.9] then give $q = (2 + \gamma^2/2)/d_\gamma \le 1$. In particular,

$$q \le \frac{3}{2}, \qquad q < 1 + \frac{a^2}{2} \tag{5.1}$$

for every $a > 0$.

For a deterministic nondegenerate circle arc I, the lower KPZ bound [1, Theorem 1.15] gives

$$\dim_{M_F^a} I \ge \begin{cases} a^{-1}(Q - \sqrt{Q^2 - 2}), & Q \ge \sqrt{2}, \\ \infty, & Q < \sqrt{2}. \end{cases}$$

In the first case, $a = q/Q \le 3/(2Q) < Q$ and

$$(Q - a)^2 - (Q^2 - 2) = a^2 - 2q + 2 > 0.$$

Hence $Q - \sqrt{Q^2 - 2} > a$, so the dimension is strictly greater than one in both cases. Apply this simultaneously to the countably many subarcs with rational angular endpoints. Every nonconstant continuous curve in the circle contains such a subarc, whereas the image of a rectifiable curve has metric Hausdorff dimension at most one.

Here Euclidean continuity and finite metric length suffice to obtain metric continuity. Indeed, for $t_n \uparrow t$, finite variation makes the curve values a metric Cauchy sequence. Lower semicontinuity,

applied with one endpoint fixed and then $n \to \infty$, identifies their metric limit with the Euclidean limit at t. The same argument applies from the right. Thus arclength parametrization is available.

If the last assertion failed, there would be curves in shrinking neighborhoods of the circle, each of length at most a common constant L, whose endpoints stay Euclidean distance at least η apart. Parametrize them on $[0,1]$ at speed at most L. Positivity and lower semicontinuity on $K \times K$ imply Euclidean equicontinuity: for each $\rho > 0$, the infimum of $M_F^a(x,y)$ over $x, y \in K$ with $|x-y| \geq \rho$ is positive. A uniformly convergent subsequence therefore has a nonconstant limit in the circle. Lower semicontinuity bounds every partition sum of the limit by L, contradicting the first assertion. □

The conclusion is uniform under every family of local laws whose densities with respect to the fixed pure law have a common L^p bound for some $p > 1$. Indeed, if E_δ is the event that the infimum in Lemma 5.1 is at most A, Hölder's inequality gives

$$\sup_W \mathbb{E}[W \mathbf{1}_{E_\delta}] \leq \sup_W \|W\|_{L^p} \mathbb{P}(E_\delta)^{1-1/p} \longrightarrow 0.$$

The same assertion holds after division by any fixed positive finite local length unit. Exact scaling gives the assertion at every physical radius, and a finite union treats the circles of one configuration. Thus the strict-buffer Gaussian comparisons below supply the required conditional uniformity without compactness of a family of continuous metrics.

6 Exact traces in an unchanged region

Theorem B.14 proves simultaneous strong confluence and strict interior uniqueness for the canonical metric at every positive parameter, including ordinary internal geodesics in a compact subdomain. We first record the deterministic consequence used to select complete traces. Its proof uses Euclidean continuity of finite geodesics; it does not require an upper bound for the metric on a Euclidean compact set.

Lemma 6.1 (Finite complete traces)**.** *Let d be an extended planar length metric whose finite geodesics are simple and Euclidean continuous. Fix concentric closed disks $C = \overline{B}_r \subset B_R$ and $D = \overline{B}_R$ in its domain. Consider a family of finite geodesics with full endpoints strictly outside D, each meeting $\operatorname{int} C$. Suppose every strict interior subsegment of every path is the unique geodesic between its endpoints. If at most N pairwise closed-disjoint proper crosscuts of $D \setminus \operatorname{int} C$ occur as subsegments of these paths, then there are at most $1 + 2N^2$ distinct complete traces on C. Equal traces have equal restricted arclength measures.*

Proof. Orient each path P. Let t_-, t_+ be its first and last contacts with C, and let s_- be its last outer-circle contact before t_- and s_+ its first outer-circle contact after t_+. The incoming arm $A_P = P[s_-, t_-]$ and outgoing arm $B_P = P[t_+, s_+]$ are disjoint proper closed crosscuts. Every point of either arm is strictly interior to the full path.

If A_P meets A_Q at x and B_P meets B_Q at y, then $x \neq y$. Both x-to-y subpaths contain all core visits of their respective full paths. Interior uniqueness makes the subpaths identical. Thus distinct complete traces satisfy

$$A_P \cap A_Q \neq \varnothing \implies B_P \cap B_Q = \varnothing\,, \tag{6.1}$$

and the same implication with A, B interchanged. This includes returns to the core and isolated boundary contacts.

Fix one proper crosscut A_0. Cut the annulus along it. Equivalently, in the universal-cover strip use a lift of A_0 and its adjacent translate to obtain a Jordan fundamental region. A topological rectangle parametrizes that region. Every crosscut closed-disjoint from A_0 lies in it and joins its two horizontal sides without touching the vertical sides.

Order such crosscuts by $A < B$ when they are disjoint and A lies to the left of B. Separation by a crosscut proves transitivity. Two labels are comparable exactly when their closed images are disjoint; shared endpoints, tangencies and coincident images are intersections. A finite subfamily of distinct traces whose incoming arms avoid A_0 has chains of size at most N. In an antichain the incoming arms pairwise intersect, so (6.1) makes the outgoing arms pairwise disjoint. Its antichains also have size at most N. Rank by the longest chain ending at each element. There are at most N ranks and each is an antichain, so the subfamily has at most N^2 members.

For the other traces, incoming arms meet A_0, hence outgoing arms avoid B_0. Slit along B_0 and apply the same argument, giving another N^2 members. Add the reference path. Any larger infinite family would contain a forbidden finite subfamily.

A geodesic is simple, and its arclength is the metric one-dimensional Hausdorff measure on its image. Equal trace images therefore have equal restricted occupation measures. □

The strict endpoint clearance is used to invoke interior uniqueness at both arm intersections. An assertion with endpoints on the outer circle is not substituted for this lemma.

Proposition 6.2 (Moments of the local trace count)**.** *For the pure metric at every fixed $a > 0$, let $N_{z,r}$ be the maximal size of a closed-disjoint family of proper crossings of $A_{r,2r}(z)$, each minimizing among curves in $\overline{A_{r/2,4r}(z)}$. Then $\mathbb{E}N_{z,r}^p < \infty$ for every $p < \infty$, with bounds independent of z, r. The same conclusion holds for complete traces in a smaller closed disk of geodesics in one common open comparison domain, provided their full endpoints clear the larger observation disk.*

Proof. This is Proposition B.15. Its proof obtains qualitative finiteness from two ordinary surrounding circuits and internal confluence, then amplifies the one-annulus tail using the Gaussian inequality of Section 4. Lemma 6.1 converts the crossing count to the complete trace count. □

6.1 Measurability and arbitrary fillings

Let W be an unchanged source-free open region. Code its internal metric as a jointly measurable lower-semicontinuous extended kernel. The path relations and completed-measurable selection are justified in the measurability argument following Proposition B.15. Select a representative, then repeat after excluding its complete trace. Outside one exterior-only null set, Lemma 6.1 makes this procedure terminate. All selected traces and their restricted occupations are measurable from the ordinary field in W.

Fix such a good exterior. Under any compatible filling outside W, a compact segment of a new minimizing path is a geodesic for the same unchanged W-internal metric whenever it stays in W. If its endpoints clear the outer buffer and it meets the open core, its trace belongs to the selected list. This is valid for every filling, including a filling with ties. A global path that leaves W and returns may use several listed traces; its entire trace is not asserted to be one list member.

There is a different statement for a variable central core: its complete trace count uses interior uniqueness through that core and holds under an ordinary conditional Gaussian filling almost surely. It is used only for moment envelopes. It is not promoted to all possible tied fillings.

The reference comparison uses one forced deterministic-endpoint witness for each accepted canonical middle. It does not require every member of the universal internal frame to have a deterministic-endpoint representative. The witness is retained in its descendants. Its uniqueness

then proves uniqueness of the common middle by concatenation with its actual outside legs, as shown in Lemma 10.2. No uniqueness assertion for newly selected random anchor endpoints is used.

For the unchanged region the union occupation counts common pieces once. If T_i are the selected traces with occupations ν_i, it equals

$$\nu_{\text{union}}|_O = \sum_i \mathbf{1}_{O\cap(T_i\setminus\bigcup_{j<i} T_j)}\,\nu_i$$

on a smaller observation region O. The sum of all representative occupations is an upper bound, not the definition of the union measure.

6.2 Finite source graphs and protected pieces

For finitely many disjoint perturbation regions, choose the surrounding ordinary separators and ports supplied by Theorem B.16. Their counts have moments of every order. Retain the true closed internal rows and the genuine exterior-restricted kernels between these ports. The theorem gives exact equality of the resulting finite graph infimum with the original query, for every compatible positive filling. It proves this equality by tightening only ordinary fragments in one fixed buffer; no minimizing path inside a changed region is frozen.

The separator is constructed from whole representative geodesics. Subdividing their finite union into edges of sufficiently small Euclidean diameter bounds the number of ports. No finiteness of intersections with a prescribed circle is required. A graph edge used later in an estimating body is an actual restricted path cost, and all unchanged avoiding routes remain available.

Protected pieces are supplied by Lemma B.4 and the two-attachment construction of Lemma B.18. The lowered metric equals a constant multiple of the ordinary metric on each complete protecting neighborhood. Ordinary local minimization and the two unique barriers then force the exact prescribed subsegment. The construction includes the complete source supports and their transitions. Finite exhaustion chooses a deterministic catalogue from the residual field before its original Gaussian amplitudes are exposed; a fixed enlarged amplitude box preserves its strict comparisons.

6.3 The first variation

For a finite family of smooth compact Weyl directions, the distance is locally Lipschitz in their real amplitudes after the usual distance truncation. Indeed every path cost changes between the factors $e^{-\|f\|_\infty}$ and $e^{\|f\|_\infty}$. At a differentiability point in a direction v, comparison with each minimizing path gives

$$D_v f(D) = f'(D)\int v\,d\nu_D\,.$$

All minimizing paths have the same integral at such a point. The Gaussian coordinate laws are absolutely continuous, so this identity holds almost surely for each of the finitely many directions in a finite calculation. Localizing by a surrounding loop gives all required moments when the terminals are outside the protected region. The joint likelihood bounds extend the identity to the source coefficients by dominated approximation.

Finite port kernels are therefore locally Gaussian-Sobolev in the retained real directions. If a slack vanishes, its weak gradient vanishes almost everywhere on that level set. This locality of weak gradients is the reason that the full first variation stays tangent to the active incidence constraints when finite cone formulas are used. It does not justify differentiating a selected combinatorial label and dropping its boundary contribution.

7 The pure occupation reference

Fix a Weyl parameter $a > 0$, let h be a whole-plane GFF modulo constants, and put $F = ah$. Write $Q = Q(a)$ and $M_F = D^a_{F/a}$. Our conventions are

$$\langle f, g\rangle_\mu = \frac{1}{2\pi}\int_{\mathbb{R}^2} fg\,\mathrm{d}x, \qquad (u,v)_{\mathcal{H}_a} = a^{-2}\langle \nabla u, \nabla v\rangle_\mu, \qquad I_v(F) = a^{-2}\langle F, -\Delta v\rangle_\mu\,.$$

The metric obeys

$$M_{F+c} = e^c M_F, \qquad M_F(Hx, Hy) = H^{aQ} M_{F(H\cdot)}(x,y)\,. \tag{7.1}$$

Fix deterministic endpoints p, q, a central capture region, and a larger coding region, with fixed positive clearances. The endpoints are outside every enlarged observation and normalizing annulus used at the initial level. Fix a smooth probability average ℓ outside the coding region. All later observation domains and source supports are inside that coding region. Let

$$\ell_H(x) = H^{-2}\ell(x/H), \qquad F^{[H]} = F - \ell_H(F)\,.$$

Let $K_H = M_{F^{[H]}}(Hp, Hq)$, and let ν_H be the length measure of its almost surely unique geodesic. The original endpoint templates may be replaced by any member of a fixed finite collection satisfying the same clearances. The pin specifies the reference law. Physical mean derivatives act on the metric functional after this representative has been selected; the pin is not reapplied after a physical mean shift. Every compact direction used below preserves the pin, so this convention agrees with differentiating before pinning for those directions.

We use capture disks $C_{H,t} = HC_t$, where t ranges over a compact interval $[t_-, t_+] \subset (0,\infty)$, and all disks lie in the central region. Fix a nonnegative smooth density χ supported in (t_-, t_+) and of integral one. Set

$$w_H(x) = \int \chi(t)\mathbf{1}_{C_{H,t}}(x)\,\mathrm{d}t, \quad d_H = \mathbb{E}\int w_H\,\mathrm{d}\nu_H, \quad \mathsf{Q}_H(\,\mathrm{d}F,\,\mathrm{d}x) = d_H^{-1} w_H(x)\nu_H(\,\mathrm{d}x)\mathbb{P}(\,\mathrm{d}F)\,. \tag{7.2}$$

In particular $\|\nabla w_H\|_\infty \le C/H$. When the radius label is retained, the joint measure is

$$d_H^{-1}\chi(t)\mathbf{1}_{C_{H,t}}(x)\nu_H(\,\mathrm{d}x)\mathbb{P}(\,\mathrm{d}F)\,\mathrm{d}t\,.$$

Fix a lattice $\tau + r\mathbb{Z}^2$ and the fixed microscopic ratio b_0. Observation radii are powers of b_0 times r; all observations use one common phase. This ratio and the forecast delay below are fixed before the later analytic block length; the terminal block ratio $b = b_0^L$ is a different quantity. A finite local current is a finite sum, or a finite feature integral, of the form

$$q(F,x) = \sum_{c,\alpha} u_{c,\alpha}(x) : \prod_{j=1}^{d_\alpha} I_{v_{c,\alpha,j}}(F) : \;. \tag{7.3}$$

At each x, only finitely many anchors c contribute. Whenever $u_{c,\alpha}(x) \neq 0$, all the indicated directions are supported in $B_D(x)$, where $D \ge r$ is fixed. Directions and deterministic output coefficients are smooth, compactly supported, and translated with the fixed lattice. Feature integrals are required to converge in the Gaussian Sobolev norms of the finitely many output derivatives used below. This is exactly the class obtained by resolving one derivative of a finite compact mean/covariance source expression and taking the pure Gaussian adjoint of its remaining derivatives.

Theorem 7.1 (Pure reference limit)**.** *The reference measures in (7.2) are well defined, and*

$$0 < c \le H^{-aQ} d_H \le C < \infty\,.$$

Every positive moment of $H^{-aQ}\int w_H \,\mathrm{d}\nu_H$ is bounded uniformly in H. For every fixed current (7.3) and every $0 < p < \infty$,

$$\sup_{H \ge CD} \mathbb{E}_{\mathbb{Q}_H} |q(F,x)|^p < \infty\,. \tag{7.4}$$

Let $H_j = b_0^j D$, with a fixed rounding of D to a compatible lattice radius. Then there are $C_q < \infty$ and $\eta_q > 0$ such that

$$\left|\mathbb{E}_{\mathbb{Q}_{H_k}} q - \mathbb{E}_{\mathbb{Q}_{H_j}} q\right| \le C_q (D/H_j)^{\eta_q}, \qquad k \ge j\,. \tag{7.5}$$

Consequently the expectations have a limit. The same limit is obtained from any of the fixed endpoint templates and smooth capture densities, provided the source phase is copied before each reference is generated. For a fixed translated prototype current, the constants are uniform over its lattice phase in one fundamental cell, and the limit is independent of that phase. The assertion concerns the original finite current, after all its internal history partitions have been recombined.

8 Gaussian comparison on strict buffers

Let $V \Subset D$ be fixed smooth domains. After translating and dilating them, the same notation denotes their copies at any physical radius $R > 0$. Let X_D be the zero-boundary pure field in D. Its Cameron–Martin space is $H_0^1(D)$ with the norm inherited from $\mathcal{H}_a$. A pin supported outside D gives the ordinary decomposition

$$F = X_D + m, \qquad X_D \perp m, \qquad \Delta m = 0 \text{ in } D\,. \tag{8.1}$$

The harmonic function includes the representative's scalar.

Lemma 8.1 (Comparison at all radii)**.** *Fix an average A on a buffered annulus outside D, and a smooth cutoff equal to one on V and compactly supported in D. For the ordinary pure law, the cutoff extension Y of $m - A$ is a centered Gaussian random element of $H_0^1(D)$ satisfying*

$$\mathbb{E}\|Y\|_{\mathcal{H}_a}^2 \le K\,,$$

with K independent of R. Write Q_D for the law of X_D and P_D for the law of $X_D + Y$ on D. Then

$$\frac{\mathrm{d}P_D}{\mathrm{d}Q_D} \ge e^{-K/2} \quad Q_D\text{-almost surely}\,. \tag{8.2}$$

For every deterministic $u \in H_0^1(D)$ with $\|u\|_{\mathcal{H}_a} \le M$, every $p > 1$, and $Q_{D,u} = \mathrm{Law}(X_D + u)$,

$$\int \left(\frac{\mathrm{d}Q_{D,u}}{\mathrm{d}P_D}\right)^p \mathrm{d}P_D \le \exp\left(\frac{(p-1)K + p(p-1)M^2}{2}\right)\,. \tag{8.3}$$

There is also $h > 0$, depending only on K, such that

$$\int \left(\frac{\mathrm{d}P_D}{\mathrm{d}Q_D}\right)^{1+h} \mathrm{d}Q_D \le e^{hK}\,. \tag{8.4}$$

The same estimates hold after restriction to V. In particular they apply to the centered conditional field under any exterior whose strict-buffer mean has the displayed Cameron–Martin bound.

Proof. At unit scale the covariance of $m - A$ is smooth on a neighborhood of $\overline{V} \times \overline{V}$. Its derivatives of every prescribed finite order are bounded there. Cutoff energy therefore has finite expectation. The pure Dirichlet energy and the centered GFF law are invariant under dilation, proving the same bound for every $R > 0$.

Conditioning on Y, the Cameron–Martin formula gives

$$r(X_D) := \frac{\mathrm{d}P_D}{\mathrm{d}Q_D}(X_D) = \mathbb{E}_Y \exp\{I_Y(X_D) - \tfrac{1}{2}\|Y\|^2_{\mathcal{H}_a}\}\,.$$

This identity follows first for finite-dimensional projections and then by Gaussian L^1 convergence. For Q_D-almost every field, $\mathbb{E}_Y I_Y(X_D) = 0$; Fubini is justified by the finite second moment. Jensen gives (8.2). The density of $Q_{D,u}$ relative to Q_D is $r_u = \exp\{I_u - \|u\|^2/2\}$. Hence

$$\mathbb{E}_{Q_D}[r_u^p r^{1-p}] \le e^{(p-1)K/2}\mathbb{E}_{Q_D} r_u^p \le e^{((p-1)K + p(p-1)M^2)/2}\,,$$

which proves (8.3).

Let $\beta_i \ge 0$ be the eigenvalues of the covariance of Y in the Cameron–Martin space. Their sum is at most K. Diagonal Gaussian integration gives

$$\mathbb{E}_{Q_D} r^{1+h} = \prod_i (1+\beta_i)^{-h/2}(1-h\beta_i)^{-1/2}\,.$$

Take $0 < h < [2\max(K,1)]^{-1}$. Using $-\log(1-h\beta_i) \le 2h\beta_i$ bounds the product by e^{hK}. Restriction to V replaces each density by its conditional expectation and preserves the bounds by conditional Jensen. The lower bound also survives restriction. □

Lemma 8.2 (Harmonic prediction counts)**.** *Along any deterministic nested address of buffered disks D_j, of radii $R_j = b_0^{-j}R_0$, let m_j be the exterior harmonic prediction. Let A_j be its prescribed exterior average, and let B_jF denote a cutoff extension of $m_j - A_j$, measured in a fixed rescaled H^s norm, $s \ge 4$. The observation and average buffers are chosen with fixed clearances. There exist $t, C > 0$ such that, for every finite address of length n,*

$$\mathbb{E}\exp\left(t\sum_{j=1}^{n}\|B_jF\|^2_{H^s}\right) \le e^{Cn}\,. \tag{8.5}$$

The same bound under a bounded GOOD starting exterior has an additional fixed prefactor. Thus the proportion of prediction norms exceeding a sufficiently large deterministic cutoff has any prescribed exponential tail rate.

Proof. Let Q_j be the Dirichlet projection onto $H^1_0(D_j)$. The differences $E_i = Q_i - Q_{i+1}$ are orthogonal. The operator B_j annihilates E_i for $i \ge j$. For $i < j$, the function E_iu is harmonic on the intervening inner domain. Interior harmonic estimates, subtraction of A_j, and rescaling give

$$\|B_jE_iu\|_{H^s} \le Cb_0^{-(j-i)}\|E_iu\|_{\mathcal{H}_a}\,.$$

A fixed number of adjacent indices are absorbed by enlarging C. The initial outer harmonic component has the same geometrically decaying estimate. The convolution matrix $b_0^{-|i-j|}$ is bounded on ℓ^2, so the joint prediction map has bounded operator norm. At each scale its squared Hilbert–Schmidt norm is bounded by the smooth covariance calculation of Lemma 8.1. The

covariance of $(B_jF)_{j\le n}$ therefore has bounded operator norm and trace at most Cn. The Gaussian determinant formula proves (8.5) for small fixed t.

Conditioning on the starting exterior replaces the Gaussian covariance by a compression, preserving both bounds. A starting centered harmonic mean bounded in a strict-buffer rescaled H^s norm contributes a deterministic sequence with square sum at most CM^2: its gradient contributes R_j/R_0 and its higher derivatives contribute still smaller powers. Completing the square supplies the fixed prefactor. Markov's inequality applied to (8.5) proves the count assertion. □

Lemma 8.3 (Separated observations)**.** *Let $G, V \Subset D$ have disjoint buffered closures. Under the Dirichlet conditional law in D, their Gaussian observation spaces have canonical correlation at most $\rho < 1$, uniformly in physical radius and in the exterior mean. Consequently, for events $A \in \sigma(F|_G)$ and $B \in \sigma(F|_V)$,*

$$P_e(A \cap B) \ge [P_e(A)P_e(B)]^{\gamma}, \qquad \gamma = (1-\rho)^{-1}. \tag{8.6}$$

Proof. Choose a cutoff equal to one near G and zero near V. Multiplication by this cutoff has norm at most C on $H^1_0(D)$, by Poincaré's inequality. The variational formula for conditional Gaussian variance gives

$$\operatorname{Var}(\lambda(F)|F|_V, e) \ge C^{-2} \operatorname{Var}(\lambda(F)|e)$$

for every linear observation λ on G. Thus one may take $\rho \le \sqrt{1-C^{-2}} < 1$. For finite observations the block covariance is at least $(1-\rho)$ times its block diagonal. The reverse Gaussian Brascamp–Lieb inequality [11, Theorem 1(ii)], applied with both exponents $1-\rho$, proves (8.6). Bounded cylinder conditional expectations converge in L^2, so the same inequality holds for the field sigma-fields. Means are absorbed in the functions to which that inequality is applied. □

9 Fixed-parameter geometry and source forecasts

We use the fixed-parameter geometric estimates to construct the occupation-weighted conditional transitions.

The fixed-shape crossing and surrounding-circuit moments are recorded in Section 2. Theorem B.14 and Proposition B.15 give the simultaneous internal geodesic statements and the separate all-order trace estimate.

Lemma 9.1 (Closed exterior locality)**.** *Let D be a bounded smooth domain and let F be a pure GFF pinned by a fixed functional supported in the open exterior of $\overline{D}$. The length functional on curves contained in the closed exterior of D, and every intrinsic kernel obtained by taking an infimum over a fixed measurable class of such curves, are measurable from the field on the open exterior of D.*

Proof. Choose smooth $D_j \Subset D_{j+1} \Subset D$ exhausting D. The closed exterior of D lies in the open exterior V_j of $\overline{D}_j$. Internalizing a length metric in an open neighborhood does not change the length of a curve contained in that neighborhood. Locality of the LQG metric therefore makes the asserted lengths and kernels measurable from $F|_{V_j}$ for every j. Measurability of an infimum may be interpreted in the completed sigma-field: curve length is a supremum over rational partitions, and the relevant existence sets are analytic.

For the pure GFF, the orthogonal complement of the exterior Gaussian observation space is $H^1_0(D_j)$. The increasing union of these spaces is dense in $H^1_0(D)$, since it contains every compactly supported smooth function in D. Consequently the decreasing exterior Gaussian observation

spaces have intersection equal to the observation space of F outside D. The same is true of their completed Gaussian sigma-fields. One way to check the last assertion is to apply the converging orthogonal projections to each finite Gaussian chaos; Gaussian polynomials are dense in L^2, so the corresponding conditional-expectation operators converge to conditioning on the intersection observation space. The reverse martingale theorem now proves the claim. □

9.1 Local lower support and compact lowering

All metric statements in this subsection concern the ordinary metric at the fixed parameter $a > 0$. A restricted estimating edge is the length of an actual physical curve; it need not minimize the ordinary metric.

Lemma 9.2 (Simultaneous conditional lower support)**.** *Fix an exterior state e, a bounded open domain D, and finitely many ordinary local row costs $K_i(e, X)$, where X has the conditional Dirichlet Gaussian law in D. Let*

$$m_i(e) = \operatorname*{ess\,inf}_X K_i(e, X).$$

There is a deterministic increasing sequence $f_n \in C_c^\infty(D)$, $f_n \geq 0$, cofinal under domination of compactly supported smooth functions, such that, almost surely and simultaneously in i,

$$K_i(e, X - f_n) \downarrow m_i(e).$$

The conclusion is unchanged by a nondegenerate finite-rank covariance subtraction with covariance at least half the original covariance. If an old estimating row lifts to physical curves and retains every genuine route avoiding D, its cost satisfies

$$m_i(e) \leq K_i^{\mathrm{old}} \leq B_i(e)$$

under every reconstructed law absolutely continuous with respect to that conditional ordinary law. Here B_i is the genuine avoiding-D cost.

Proof. Choose the sequence to diverge uniformly on compact subsets. It can be constant near any prescribed finite collection of compact reserve supports and surrounding circuits. Put $L_i(X) = \inf_n K_i(e, X - f_n)$. For a compact smooth v, cofinality and Weyl monotonicity give $L_i(X + v) = L_i(X)$. First impose these identities for a countable dense subgroup of the Cameron–Martin space. Translation continuity in probability extends them to its closure. A measurable random variable invariant under these translations is constant: after a bounded injective transformation, expand it in Gaussian chaoses; invariance under every finite-dimensional Gaussian translation makes all nonconstant coefficients zero. Write this constant as c_i. Each fixed translate is equivalent to the original law, so $L_i \geq m_i$. Since $L_i \leq K_i(e, X)$, constancy gives $c_i \leq m_i$. Thus $c_i = m_i$. The finite intersection handles all rows. The residual Gaussian law has the same Cameron–Martin space and null sets, so the argument also applies there. Covariance comparability by itself is not an equivalence theorem in infinite dimension: applications here retain a finite menu. An infinite-rank limit is allowed only after establishing the Hilbert–Schmidt equivalence criterion and preservation of the relevant null sets; no such limit is inferred from the lower covariance bound.

An estimating row has cost at least the ordinary row because each of its allowed paths is physical. Its retained avoiding routes give the upper bound. Absolute continuity preserves the ordinary lower-support null set. No claim is made for arbitrary positive numbers substituted as opaque row costs without such a physical lift. □

Lemma 9.3 (A local realization of the lower support)**.** *Fix the observation block O and local row domains before any numerical cutoff. Let $\overline{H} \subset \operatorname{int} H' \Subset D \Subset O$. Every row endpoint lies outside the retained circuit around H'. Only the interior visit to H' is freed; the incoming and outgoing ordinary legs and their original labels remain actual physical paths. Far exterior graph costs are left symbolic. Choose increasing cofinal $g_n \in C_c^\infty(D \setminus \overline{H'})$, and let $F_{i,n}$ be the ordinary local row for $F - g_n$ with its visit to H' free. Then*

$$F_{i,n} \longrightarrow m_i(e)$$

almost surely, simultaneously in the finite row set. All the quantities on the left are measurable in the same fixed block outside H'.

Proof. Retain a finite ordinary circuit Γ surrounding H' and compact in $D \setminus \overline{H'}$. Its length ℓ_n after lowering tends to zero by dominated convergence. First/last-contact replacement gives

$$F_{i,n} \le K_i(F - g_n) \le F_{i,n} + \ell_n.$$

The right-hand row is at least m_i by equivalence under the fixed shift. For the reverse inequality take an almost-minimizer for $K_i(F - f_N)$ from Lemma 9.2. Replace its entire visit inside Γ by a circuit arc. On the retained compact portions outside Γ, eventually $g_n \ge f_N$; outside D both are zero. The circuit cost tends to zero. Consequently

$$\limsup_n K_i(F - g_n) \le K_i(F - f_N).$$

Let the minimizing error tend to zero, then $N \to \infty$. Closed-exterior locality, applied with the fixed open observation buffer, gives the asserted measurability. All joins above are actual contacts of ordinary curves. There is no switch between labelled restricted atoms. □

Lemma 9.4 (Two attachments under arbitrary inner fillings)**.** *Fix the parameter $a > 0$, disks $H \Subset D$, and an unchanged surrounding rectifiable loop outside D. All full query endpoints lie outside that loop. There is a finite catalogue of compactly supported smooth shifts in $D \setminus \overline{H}$, on an ordinary event of positive probability, such that every minimizing query meeting the retained middle or either of two private source cores traverses two prescribed ordinary arms in opposite directions. The portion between their inward anchors stays in a fixed compact subset of D and avoids the full private source supports. All other attachments lie in this middle.*

The conclusion holds throughout a fixed enlarged box of additional source amplitudes and for every compatible positive filling of H. For fillings with unattained infima, the same restricted-row identities hold after ordinary tightening and passage to infima. An ordinary symmetric filling has one common scalar middle; a directed filling retains its two oriented middles and the rowwise formula (17.3)*.*

Proof. Lemma B.18 gives the two complete ordinary attachments, the matched-domain confinement, and uniformity over the amplitude box. Splitting at the anchors gives one inequality for the row identity. Concatenating arbitrarily accurate restricted representatives gives the other. Symmetry identifies the two oriented middles only for an ordinary filling. All ordinary replacements occur outside H. □

Closed-exterior kernels use the completed exterior measurability proved in Lemma 9.1.

9.2 A deterministic-address count

For the local annular and punctured-parent observations along one deterministic address, let P_j denote their Gaussian Hilbert-space projections, modulo constants. Cut off a Cameron–Martin function on each observation buffer and subtract its local average. Poincaré and bounded overlap give

$$\sum_j \|P_j u\|_{\mathcal{H}_a}^2 \le K_{\mathrm{ref}} \|u\|_{\mathcal{H}_a}^2 . \tag{9.1}$$

The overlap constant depends on the fixed delay L_{ref}, not on subsequent port meshes or numerical cutoffs. Gaussian Brascamp–Lieb therefore gives, for events B_j of probabilities at most δ,

$$\mathbb{P}\{\#\{j \le n : B_j\} \ge \eta n\} \le 2^n \delta^{\eta n / K_{\mathrm{ref}}} . \tag{9.2}$$

The harmonic prediction tests are handled separately by Lemma 8.2.

A state e consists of the physical exterior of a disk D_e, a central root cell C_e, endpoints and a retained domain, the common grid phase, and revealed auxiliary radii. Eligible endpoints lie outside a fixed enlarged annulus of D_e. Every canonical descendant retains the label of its original witness from a fixed countable deterministic-endpoint class. The successive outside-leg plus middle distance identities compose for that same witness. Let P_e be the ordinary Dirichlet conditional Gaussian law. Let L_e be the occupation in C_e, either for one path or for the universal union frame specified before sampling the core. Its rooted kernel is

$$\Pi_e(\,\mathrm{d}F,\,\mathrm{d}x) = Z_e^{-1} P_e(\,\mathrm{d}F)\mathbf{1}_{C_e}(x)\nu_e(\,\mathrm{d}x), \qquad Z_e = E_e L_e . \tag{9.3}$$

Fix the state geometry as follows; changing these numerical ratios only changes constants. Write $D_e = B(z_e, 4R_e)$, and put the normalizing annulus in $B(z_e, 6R_e)\setminus\overline{B}(z_e, 5R_e)$. Eligible full endpoints lie outside $\overline{B}(z_e, 8R_e)$. Choose a circuit in the normalizing annulus by a fixed measurable rule from a finite family of rectangle crossings, and let $U_e > 0$ be its length. These crossings and their complete buffers are in the physical exterior of D_e. The circuit separates D_e from the eligible endpoints, has all positive moments in the local length unit, and is bounded below by a fixed-shape crossing, which has all negative moments. It is chosen once from the exterior, before any core filling. For one path, $L_e \le U_e$ for every filling: replace the portion between its first and last circuit contacts by an arc of that circuit. For the union frame,

$$L_e/U_e \le C(1 + N_e^2) , \tag{9.4}$$

where N_e is the ordinary disjoint-crossing count in a buffered inner annulus. Here the frame is the union occupation, counting shared pieces once. It is bounded by the sum of one representative occupation per complete trace, so (9.4) follows from Lemma 6.1.

Fix a smooth nonnegative probability average ℓ_e in the normalizing annulus, obtained from one fixed template by translation and dilation, and put

$$A_e = \ell_e(F), \qquad s_e = R_e^{aQ}, \qquad \kappa_e = \log(U_e/s_e), \qquad c_e = \kappa_e - A_e. \tag{G0}$$

Let m_e be the exterior harmonic prediction on D_e. Choose once and for all $s_0 \ge 4$, a strict retained buffer $V_e \Subset D_e$, and a rescaled smooth cutoff ζ_e equal to one on V_e and supported in D_e. All retained seed and observation regions lie in V_e after a fixed bounded number of levels. Define the exterior-measurable prediction norm

$$b_e = \big\|\zeta_e(z_e + R_e\,\cdot)[m_e(z_e + R_e\,\cdot) - A_e]\big\|_{H^{s_0}(B_4)}. \tag{G1}$$

For numerical constants $0 < p_0 < 1$ and $M_0, M_1 < \infty$, the exact GOOD test is

$$e \text{ is eligible}, \qquad q_e \geq p_0, \qquad |c_e| \leq M_0, \qquad b_e \leq M_1 \,. \tag{9.5}$$

The forecast q_e is specified in the next lemma. In particular, on GOOD,

$$\big\|\zeta_e(m_e - \kappa_e)\big\|_{\mathcal{H}_a} \leq C_a(M_0 + M_1), \qquad e^{-M_0} \leq U_e/(e^{A_e} s_e) \leq e^{M_0}. \tag{G2}$$

The first norm is unchanged by rescaling. These, rather than an implicit bound on a pinned scalar, are the mean and normalization bounds used below.

9.3 Source events in a fixed observation block

The following construction replaces boundary-point access by finitely many actual contacts. Fix first the observation domains, delayed hidden alternatives, reference scale, port cutoff, and Gaussian overlap constant. The ordinary protected-reference and finite-trace conclusions of Lemma B.4 and Proposition B.15 are used inside these domains.

Lemma 9.5 (Local source tests and their count)**.** *For any prescribed raw failure probability $\varepsilon > 0$, the numerical source tests can be chosen inside one predetermined observation block so that their simultaneous failure has ordinary probability at most ε. The event applies to every eligible local port pattern and every prescribed hidden alternative. It implies the parent source event for every positive hidden filling. Its construction does not read the far costs used later to select a saving pattern.*

Proof. Let N bound the number of local edges of a simple port pattern, and fix a saving tolerance $\delta > 0$. Replace its local rows by their lower supports m_i, keeping the far edges exact and symbolic; denote the resulting graph value by F_m. Lemma 9.3 lets us choose one deterministic finite n_0 so that all local row errors and circuit errors are at most $\delta U/(4N)$ outside an arbitrarily small raw event. The minimum of finitely many path sums changes by at most the largest pathwise sum of row errors. Hence on the saving flag $T_{\mathrm{ext}} - F_m \geq \delta U$,

$$F_{n_0} \leq F_m + \delta U/2 \leq T_{\mathrm{ext}} - \delta U/2.$$

Every subsequent path and deleted subarc is selected in this fixed baseline. An all-exterior pattern supplies no saving, and an ordinary minimizing representative contains no positive loop. For each remaining eligible pattern the protected-reference construction supplies an ordinary subarc to delete, outside the delayed hidden enlargement. At the fixed n_0 these finitely many lengths are positive. A deterministic cutoff therefore makes their minimum at least $L_0 U > 0$, with arbitrarily small additional failure probability. We do not increase n_0 after this step.

All paths selected by this raw event lie in the observed block outside the delayed hidden enlargement H'. In particular, no loop entering H' is selected or tested. If a frozen comparison uses the free visit to H', trim its two actual exterior legs at their first and last contacts with the retained circuit surrounding H'. Join at these actual contacts in the observed exterior. The protected alternative is chosen outside H' and replaces the entire free visit; otherwise that visit is kept free in the comparison and is not assigned a realized physical cost. Finite crossing networks and, when needed, loops wholly contained in this observed exterior join the finitely many actual curves. Compact lowerings of the permitted observed replacement make its added cost less than $L_0 U/4$. Any later physical realization inside H' is a conditional existence argument and is not part of this raw event. Choose the common negative floor before the remaining error allowances.

The forcing is zero on every unaffected baseline piece and on the coding buffers. Its walls enforce the protected alternatives; every new negative support outside the free hole is reached through the protected anchors. An uncaptured return repeats a protected ordinary point and has a removable positive loop. Thus an uncaptured route costs at least the frozen baseline, whereas the captured replacement deletes at least L_0U and adds less than its fixed fraction. The earlier $\delta U/2$ margin is retained. This comparison is with F_{n_0}, never with a free matrix recomputed after the final shift.

There are finitely many patterns and hidden alternatives. The required route lengths and profile energies are finite, the crossing and deleted lengths are positive, and the protected tolerances are positive on the ordinary geometric event. Countable rational tubes, smooth profiles, integer wall heights and rational amplitude radii give a countable measurable catalogue. Finite exhaustion fixes a finite catalogue and all cutoffs with total failure at most ε. The observation block, delay and address degree do not change during this exhaustion. Conditional projections for the fresh residual construction use the local rows of Lemma 9.3; no genuinely far row is made local by calling its conditional expectation a forecast. □

For completeness, here is the count and forecast calculation. After the fixed separation coloring, Lemma 4.3 gives, for a separated index set I,

$$\mathbb{P}\Big(\bigcap_{i\in I}\mathrm{Bad}_i\Big) \le \varepsilon^{|I|/K}.$$

Consequently

$$\mathbb{P}\{\#\mathrm{Bad} \ge \eta n\} \le C\,2^n \varepsilon^{\eta n/K},$$

with the coloring absorbed into fixed constants. Choose ε after K, the coloring and the address degree. A strict-buffer Gaussian comparison transfers this estimate from a GOOD start once.

Keep the possibly nonlocal saving event S_j separate. If the local raw event E_j implies A_j^0 for every hidden filling, put $q_j = \mathbb{P}(S_j \cap A_j^0 \mid \mathcal{F}_j)$. Then

$$\mathbb{E}[\mathbf{1}_{S_j}\mathbf{1}_{E_j}\mathbf{1}_{\{q_j<p_0\}} \mid \mathcal{F}_j] \le p_0.$$

The delayed residue-class exponential-moment calculation applies to these indicators under their actual filtration. It never asserts a uniform raw probability under an arbitrary rare exterior. Finally perform the one whole-root Campbell transfer with the original normalizer. These estimates, the retained annular crossing/circuit tests and the prediction estimate prove the count assertion in Lemma 9.6.

Lemma 9.6 (Forecast source and GOOD states)**.** *There are a fixed finite delay L_{ref}, raw parent source events A_j^0, and raw retained events E_j^0, with the following properties. Each E_j^0 is revealed by level $j + L_{\mathrm{ref}}$, implies A_j^0 for every filling of that hidden descendant, and its ordinary failure count has any prescribed exponential rate after the numerical cutoffs are chosen. If S_j is the exterior saving flag, set $A_j = S_j \cap A_j^0$ and $q_j = P(A_j \mid \mathcal{F}_j)$. Use precisely (9.5) for GOOD. The constants can be chosen so that:*

1. *On GOOD, $cU_e \le Z_e \le CU_e$ and, for each fixed $p < \infty$, $E_e(L_e/U_e)^p \le C_p$.*
2. *Along an occupied rooted address, the number of non-GOOD levels in a prefix has any prescribed exponential tail rate, both for an aligned initial capture and after a GOOD starting exterior, with a fixed prefactor.*

Proof. Fix the two separated hidden alternatives and the delay before numerical cutoffs. The local port patterns retain their true far exterior edges as symbolic costs. Use the lower-support graph value F_m from Lemma 9.3, and set $S_e = \{T_{\mathrm{ext}} - F_m \geq \delta U_e\}$. An ordinary row visiting the capture region crosses a fixed buffered annulus on entry and exit. A compact lowering can make the retained intervening portion arbitrarily cheap while leaving its exterior legs unchanged. In particular the lower-support graph gains at least twice a retained separating-crossing cost $W_e > 0$. Thus, with the original row label retained,

$$\{L_e > 0\} \cap S_e^c \subset \{W_e \leq \delta U_e/2\}. \tag{9.6}$$

This argument uses a portion of an actual finite ordinary curve. It does not approximate arbitrary boundary endpoints. The finite local-pattern event in Lemma 9.5 makes the saving physical by its fixed-baseline replacement and supplies a captured crossing of length at least $\ell_0 U_e$. The event is valid for all hidden fillings and is local on the previously fixed block. Its finite profile catalogue has a finite Cameron–Martin energy bound H_0.

If $q_e \geq p_0$, one of its finitely many source pieces has conditional probability at least p_0/J. Cameron–Martin Cauchy–Schwarz gives

$$Z_e \geq \ell_0 U_e (p_0/J)^2 e^{-H_0^2}.$$

The source coordinates retain their original product windows by Lemma 10.1; their fixed probabilities are included in the constant. The upper frame moments are proved before this lower bound is used for normalization. On GOOD, the centered prediction and unit cutoffs give the strict-buffer ordinary density comparison of Lemma 8.1. Apply one Hölder inequality to the finite frame bound (9.4), the all-order ordinary trace estimate of Proposition B.15, and the fixed annulus circuit moments. It follows that $E_e(L_e/U_e)^p \leq C_p$ for each prescribed finite p, including $Z_e/U_e \leq C_1$.

For the counts, the raw events have arbitrarily small failure probability inside the same fixed block. The Gaussian-overlap and subset count following Lemma 9.5 give any prescribed ordinary exponential rate. Put $Z_j = \mathbf{1}_{S_j} \mathbf{1}_{E_j^0} \mathbf{1}_{\{q_j < p_0\}}$. It is revealed by $j + L_{\mathrm{ref}}$ and $E[Z_j \mid \mathcal{F}_j] \leq p_0$. In each residue class of the fixed delay,

$$E \exp\{t \textstyle\sum_j Z_j\} \leq [1 + p_0(e^t - 1)]^{n/L_{\mathrm{ref}}+1}.$$

Taking $t = \frac{1}{2}\log(1/p_0)$ makes this count rate arbitrarily large as $p_0 \downarrow 0$. The remaining occupied failures are contained in the raw failures, the low forecasts, the small annular ratio in (9.6), and the normalization/prediction failures. The annular ratio has probability tending to zero by the fixed crossing and circuit moments; the prediction count is Lemma 8.2. No far-dependent saving flag is counted as a raw local event.

For an aligned initial capture, union the at most Cb_0^{2n} deterministic addresses before changing measure. Its local capture occupation has the independent circuit moment and lower mean. A single Hölder inequality transfers the ordinary count to its original Campbell law. Choose the raw rate to absorb this fixed entropy and loss. After a GOOD exterior, run the forecast martingale under its actual conditional law and transfer only the local raw failures once on a strict buffer. The just established frame moments and lower normalizer allow the final Campbell transfer. This proves both assertions without a boundary-access network or a capped-radius endpoint approximation. □

10 The annular common submeasure

We give the measure calculation explicitly, since matching an ordinary annular law alone is not enough for an occupation root.

Lemma 10.1 (Residual-product acceptance of an original row)**.** *Retain an original row label ℓ and its original component mass Z_i. Suppose the two ordinary augmented laws have the form*

$$P_i^{\mathrm{res}}(da)\,\Phi(dA)\,P_{g(a)}(dY),$$

and on a residual event, throughout $I_1^{J_0}$, their captured occupations satisfy

$$\nu_{i,\ell}(a,A,Y,dx)|_C = e^{\kappa_i}\bar{\nu}_\ell(g(a),Y,dx)|_C, \qquad U_i = e^{\kappa_i}s_R, \qquad Z_i \le CU_i.$$

Let $\gamma_i \le P_i^{\mathrm{res}}$ have equal pushforwards under all matched residual data and original labels, and put $\Phi_0 = \mathbf{1}_{I_0^{J_0}}\Phi$. Then

$$\widehat{\Pi}_i = \frac{1}{Cs_R}\gamma_i(da)\Phi_0(dA)P_{g(a)}(dY)\bar{\nu}_\ell(g(a),Y,dx)|_C$$

is a common rooted subkernel. Its acceptance derivative relative to the original normalized component is

$$\frac{Z_i}{CU_i}\frac{d\gamma_i}{dP_i^{\mathrm{res}}}(a)\mathbf{1}_{I_0^{J_0}}(A) \le 1.$$

It is measurable before the hidden innovation Y is revealed. The occupation identity concerns the retained original row whose forcing event guarantees capture throughout the box. It is not an identity for every row of the port matrix irrespective of activation.

Proof. Substitute the occupation identity into the normalized component $Z_i^{-1}P_i^{\mathrm{res}}\Phi P_g\nu_{i,\ell}|_C$; the displayed quotient follows directly. Equality of the matched pushforwards gives equality of the common rooted pushforwards. Subtract this subkernel on rejection. When the row is part of a frame, first reattach its original weight Z_i/d; then sum the original rows and comparison cells and divide by the one original whole-frame normalizer d. No occupation-dependent row choice or renormalization of a realized captured mass occurs.

The common residual cut uses the conditional forecast $\psi(g) = \int \bar{\nu}_\ell(g,Y,C)P_g(dY)$, cut below by $z_0 s_R$ before the state-dependent score and meet cuts. The latter cuts are taken so their total variation loss is smaller than the retained common mass. They do not read A and hold on the enlarged box. A Cameron–Martin forcing translates (X,A) to $(X+f,A)$, so it leaves the product coordinate law unchanged. The residual covariance bound gives the required score estimates with its actual precision; orthogonality of the profile to one coordinate is not used as a substitute for that estimate. □

Lemma 10.2 (Pure annular minorization)**.** *Two GOOD states at the same physical scale and template admit a causal coupling which, with probability at least $p_* > 0$, matches a canonical normalized field and its occupation root on a strict inner disk. The accepted and rejected flags inspect only the revealed annulus and independent auxiliaries. After rejection each unrevealed core retains its own kernel (9.3). After acceptance, the two canonical normalized core kernels are identical, not merely close. All assertions about conditional states hold outside the null sets of their actual exterior marginals; arbitrary exceptional rough exteriors are not included.*

Proof. Set $s_R = R^{aQ}$, $\kappa_e = \log(U_e/s_R)$, and $X_e = F - \kappa_e$. The centered harmonic mean of X_e on the retained buffer has bounded cutoff energy on GOOD. In normalized coordinates use an outer disk D, a retained domain $R_0 \Subset D$, a hidden disk $H \Subset R_0$, a root cell $C \Subset H$, retained annulus $G = R_0 \setminus H$, and a separated outer boundary neighborhood $V \subset D \setminus \overline{R_0}$. All clearances are fixed.

The protected seed event in G has a positive ordinary pure probability. Lemma 8.1, including its reverse direction, gives a positive Dirichlet seed probability. Truncate the two bounded-mean shift scores on that seed. The two ordinary laws on G then dominate $c\rho_{\mathrm{old}}$ for one finite positive measure ρ_{old} of mass $p > 0$.

The outer part B_e of every parent source lies in V together with the old exterior. GOOD implies $P_e(B_e) \geq p_0$. If $q_e(g) = P_e(B_e \mid G = g)$, apply (8.6) to $E = \{g : q_e(g) < q_*\}$ and B_e to obtain

$$P_e(E) \leq (q_*/p_0^{\gamma})^{1/(\gamma-1)}\,.$$

Choose q_* so the two discarded ρ_{old} masses are small. Conditional Markov's inequality applied to separately chosen high-probability middle tests then leaves positive common mass on which the outer-and-middle source has conditional probability at least $q_*/2$.

Apply Lemma 9.4. It has two essential features. First, the portion of every macro geodesic visiting C between its first and last contacts with the normalizing loop has cost at most U_e; true wall plateaus and terminal neighborhoods with the retained access bound force it through both protected anchors, in opposite arms of the prescribed series network. It cannot leave the retained domain between those anchors without repeating one of them. Second, the outer saving forces a fixed deterministic-endpoint witness through the same two anchors for every compatible filling. Its exact distance is

$$D_{\mathrm{witness}} = A_{\mathrm{ext}} + J_g\,,$$

where J_g is the internal anchor-to-anchor distance in the retained domain. Concatenation of approximate competitors proves both inequalities; separate outside-leg attainment is not needed to obtain this identity. Once the actual witness geodesic is split at its anchors, equality shows that its outside factors and middle attain their respective infima. Any different minimizing middle could be concatenated with these actual outside factors. A positive loop in this concatenation would contradict minimality; otherwise it gives a different minimizing witness image. Simultaneous uniqueness for the original countable witness class, followed by Fubini over the annulus and core, therefore makes the middle unique for almost every ordinary conditional filling. Canonical descendants retain this original witness label; uniqueness is never invoked directly for field-selected anchor endpoints. Thus the path or union frame has exactly $\nu_{J_g}|_C$ on that filling. Exceptional tied fillings need no chosen occupation for the present measure argument.

Strict inner saving from exterior data. Work with the actual intrinsic central path problem between its two physical access anchors, as identified by the exact serial decomposition. The common serial source and connector costs lie outside this path problem and cancel from the entering-versus-avoiding comparison; no extra fixed exterior leg is included in the central cost. Choose two disjoint compact access envelopes $E_1, E_2 \Subset D$ with connected observed exterior portions and with positive-width windows meeting the smooth boundary of H. The anchors and the permitted local replacement paths lie in these envelopes. Every H-avoiding central route has a segment passing between the envelopes outside their interiors. Let wU be the intrinsic cost of this passage in the closed exterior of H, with the declared ordinary middle path class. The retained bypass gives finiteness; compact separation and lower semicontinuity give positivity. Closed-exterior locality makes w an exterior observation. Take a positive numerical lower cutoff before the access forcing.

All geometric templates are selected from a fixed countable catalogue. Conditional on the independent grid phase, a template contains one fixed boundary point b_i in each window and two fixed tangent disks there:

$$\overline{B}_i^{\mathrm{out}} \subset \operatorname{int} E_i \setminus H, \qquad \overline{B}_i^{\mathrm{out}} \cap \overline{H} = \{b_i\}, \qquad \overline{B}_i^{\mathrm{in}} \subset \operatorname{int} E_i \cap \overline{H}, \qquad \overline{B}_i^{\mathrm{in}} \cap \partial H = \{b_i\}.$$

Their centers and tangency points are deterministic once the template and phase are fixed; they are not contacts selected by seeing an inner curve. Choose finite observed paths, compactly contained in $\operatorname{int} E_i \setminus \overline{H}$, from each anchor to the corresponding outer disk center. Such connections are supplied by the retained ordinary reference paths, accessible loops wholly inside the observed exterior, and a finite network of fixed-point internal connections.

For each fixed disk and its specified boundary point, [4, Lemma 2.9] gives a finite internal distance, with a positive moment of every order $p < 2Q/a$. Only finiteness is used here. Applied to the outer tangent disk, it supplies a finite metric-continuous path from its center to b_i inside that closed disk. Stop at its first contact with b_i and concatenate with the observed anchor connection. After arclength parametrization the resulting prefix γ_i lies in $D \setminus \overline{H}$ before its terminal point and has no arclength atom at that endpoint. Its entire length functional is a closed-exterior observation. These assertions concern finitely many specified ports; there is no simultaneous access claim for every boundary point.

Choose a deterministic increasing cofinal sequence

$$f_n \in C_c^\infty\big((\operatorname{int} E_1 \cup \operatorname{int} E_2) \setminus \overline{H}\big), \qquad f_n \geq 0.$$

The template includes all required flat reference neighborhoods before this choice. The functions may be constant on those neighborhoods and still cofinal; the ordinary source sleeve is outside these central lowerings. If a common upstream reference floor is also changed, its serial contribution is identical in the entering and avoiding comparisons and cancels. For the fixed exterior prefixes,

$$a_i^{(n)}(e) = \int_{\gamma_i} e^{-f_n}\, d\nu_e \longrightarrow 0.$$

Indeed $f_n \to \infty$ at every point of the prefix except its terminal boundary point, that point has zero arclength, and the original prefix length is finite. Dominated convergence applies. These are exterior quantities. Every avoiding route still pays the unchanged passage wU outside the two envelopes.

We next prove the needed conditional lower-support comparison without observing an interior path. Apply the same specified-port result to the inner tangent disks. For almost every ordinary field it supplies finite paths from b_i to their inner centers, meeting ∂H only at b_i after truncation. Join the two fixed inner centers by a finite path compactly inside H, using a finite chain of fixed-point internal connections. The resulting bridge λ lies in $\overline{H}$ and has zero arclength on ∂H. The simultaneous assertion for the countable template collection holds on one probability-one event; by disintegration it holds for almost every exterior and almost every conditional filling. The bridge and its length are used only as witnesses in the present proof. They enter no exterior flag, recorded catalogue, normalization, or selected forcing profile.

Let $g_m \in C_c^\infty(H)$ be a fixed increasing cofinal sequence. Its lowered bridge length tends to zero by dominated convergence. Since each f_n vanishes on a neighborhood of $\overline{H}$, the conditional Dirichlet law inside H is unchanged by that exterior lowering. Conditional Cameron–Martin equivalence under each fixed g_m therefore gives, simultaneously for the countable pairs (n, m),

$$m_H(e - f_n) \leq K_{\mathrm{core}}(F - f_n - g_m) \leq a_1^{(n)}(e) + a_2^{(n)}(e) + \operatorname{len}_{F-g_m}(\lambda).$$

Here m_H is the conditional essential infimum of this same physical central row between the retained access anchors, with the declared local path class fixed. Letting $m \to \infty$ gives

$$m_H(e - f_n) \le a_1^{(n)}(e) + a_2^{(n)}(e) \longrightarrow 0.$$

Choose the first catalogue index with $a_1^{(n)} + a_2^{(n)} < wU/2$, or use the corresponding lower-support test. This choice is exterior-measurable. Finite exhaustion bounds the selected index and profile energy on the retained raw class; it does not revise the previously frozen free-row approximation index. We have obtained

$$m_H(e - f_n) < wU/2 < T_{\mathrm{ext,H}}(e - f_n).$$

This is a strict conditional lower-support saving. It implies positive conditional mass of entering ordinary rows. The subsequent ordinary inner forcing and its conditional forecast supply the captured mass; no assertion that every positive filling has small central cost is inferred from an essential infimum. The source coordinates and their whole enlarged box remain outside H, and the common serial terms cancel, so this saving and forecast are seed-free.

Measurability convention for the prefixes. Condition on the independent phase and fixed template. Use jointly measurable versions of the ordinary internal lengths and of lengths of curves in the closed exterior. An admissible prefix is a Euclidean continuous, metric-continuous rectifiable curve in the declared closed exterior path class, ending at the specified port, with all earlier points in the open exterior. It can be parametrized on $[0, 1]$ with its subcurve length bounded by $L|t - s|$ for an integer L. The subcurve conditions are imposed at rational times, and the stated endpoint and closed-set constraints are measurable in the usual continuous-path space. Length and the infimum over these fixed path classes use the completed exterior sigma-field from closed-exterior locality. Their existence relations are analytic; a universally measurable selection is sufficient on the completed Gaussian probability space. Choose from this relation using only exterior data. Selecting an arbitrary full-field geodesic and then recording its prefix would not by itself establish this convention. The functions $a_i^{(n)}$, the first successful index, and their finite numerical cutoffs are consequently exterior-measurable. The conditional interior witnesses above are never selected as part of this construction.

Choose all heights and boundary neighborhoods for the finite catalogue before making any score cuts. For each retained inner label, every old-port continuation has the same shift on a neighborhood of G; it varies only farther out. All shifts vanish on H and outside D. This common-restriction convention is obtained by fixing the seed network and its two outgoing stubs first, then routing the finite outer-port pairs outside a larger fixed boundary neighborhood. The graph has only the specified two terminal attachments at each protected seed. An arbitrarily cheap filling of H therefore cannot connect to an upstream arm.

Let P_i^A be the two ordinary laws on the full revealed annulus $A = D \setminus H$. For a source piece $E_{i\ell j}$ and its shift $v_{i\ell j}$, truncate the annular Cameron–Martin score at T. Its minimal-extension norm is at most the full shift norm H_0. Hence

$$\sigma_{i\ell j} = e^{-T - H_0^2/2}(T_{v_{i\ell j}})_*(\mathbf{1}_{E_{i\ell j}} P_i^A) \le P_i^A .$$

An independent finite seed label and division by the finite number of forcing pieces prevent overlaps of translated images from violating domination. Here is the projection loss caused by the score cuts. Before these cuts, the common-restriction property and the conditional source lower bound

give projection at least $c_0\rho_0$ for one translated, labelled seed measure ρ_0. Uniform Gaussian score tails make the total discarded mass at most $\epsilon_T \to 0$ in each state. Let b_i be the density relative to ρ_0 of the absolutely continuous part of this discarded projection. Then $\int b_i \, \mathrm{d}\rho_0 \le \epsilon_T$. Restrict the common seed to $\{b_1 \le c_0/2,\ b_2 \le c_0/2\}$, losing at most $4\epsilon_T/c_0$ of its mass. On this restriction both cut projections dominate $(c_0/2)\rho_0$; any singular part remains a nonnegative measure. Define ρ_{new} to be this restricted common translated measure. For T sufficiently large it still has positive mass. The resulting annular submeasures γ_i satisfy

$$\gamma_i \le P_i^A \otimes u_{\mathrm{seed}}, \qquad \gamma_i^{G,\ell} \ge c_g \rho_{\mathrm{new}} \,. \tag{10.1}$$

Here ρ_{new} is the common translated, labelled seed measure; it is not the unshifted ρ_{old}. Its label recovers the anchors from the common new annular field by subtracting the prescribed common shift.

Choose a fresh inner source after these annular shifts. Its ordinary failure probability can be made δ as small as needed. Condition that event on the full new annulus; the set where its forecast is less than $1/2$ has ordinary mass at most 2δ. On its complement, the compact inner forcing and Cameron–Martin Cauchy–Schwarz imply

$$Z_{i,a} := E_{P_i(\cdot|a)} \nu_{J_g}^{\mathrm{physical}}(C) \ge z_0 U_i \,. \tag{10.2}$$

The realized inner event is not an acceptance flag. Filter γ_i by this forecast. If b_i is the density, relative to ρ_{new}, of the discarded projection, then $\int b_i \, \mathrm{d}\rho_{\mathrm{new}} \le 2\delta$. Removing $\{b_i > c_g/2\}$ for both states loses at most $8\delta/c_g$ common mass. Additional strict-core prediction cutoffs are imposed in the same way, after translation. Thus, by taking their probabilities small enough, one obtains annular submeasures $\widetilde{\gamma}_i$ with exactly the same positive projection $\lambda(\,\mathrm{d}g,\,\mathrm{d}\ell)$ and satisfying (10.2). Thinning to that common projection depends only on the labelled annulus.

The pure Markov property now removes every residual mean discrepancy. Since g contains a full neighborhood of ∂H, it determines the harmonic extension of the normalized total field to H. Conditional on g, the remaining normalized field has one law P_g: that harmonic extension plus an independent a-Dirichlet GFF in H. Let $\nu_{J,g}$ denote its canonical middle occupation. It is the same functional for both states. The actual rooted transition, on the forced annulus, is

$$P_i^A(\,\mathrm{d}a)\,\frac{e^{\kappa_i}}{Z_i} P_g(\,\mathrm{d}X)\mathbf{1}_C(x)\nu_{J,g}(X,\,\mathrm{d}x) \,.$$

Since $Z_i \le CU_i = Ce^{\kappa_i} s_R$, it dominates

$$\frac{1}{Cs_R}\widetilde{\gamma}_i(\,\mathrm{d}a,\,\mathrm{d}\ell) P_g(\,\mathrm{d}X)\mathbf{1}_C(x)\nu_{J,g}(X,\,\mathrm{d}x) \,. \tag{10.3}$$

The projection of (10.3) onto (g,ℓ,X,x) is identical for $i=1,2$. Write $m_\lambda > 0$ for the total mass of λ. The common submeasure has mass at least $z_0 m_\lambda / C$, because $E_{P_g}\nu_{J,g}(C) = e^{-\kappa_i} Z_{i,a} \ge z_0 s_R$.

The acceptance density relative to the original rooted transition equals

$$\frac{\mathrm{d}\widetilde{\gamma}_i}{\mathrm{d}(P_i^A \otimes u_{\mathrm{seed}})}(a,\ell)\frac{Z_i}{CU_i} \,,$$

which depends only on the annulus and the old state. Therefore rejection does not condition the core on any future event. Couple the common projection of (10.3), sample its one common core and root, and sample each remaining annular field from its own disintegration. This gives the asserted causal coupling and exact equality of the accepted normalized cores.

The preceding calculation is the empty-menu case. For a retained nonempty source menu, apply Lemma 10.1. Tensor the original law of the entire coordinate vector, including its nonselected coordinates, before making any cut or thinning. All catalogue, forecast, score and common-projection cuts are residual-measurable and hold on the whole enlarged box I_1^J. Enlarge g to retain the normalized residual annulus, the template, original row label, reserve basis, selected indices, and every other residual datum used by the canonical middle. The source directions vanish near H, so its conditional kernel P_g does not depend on these coordinates. On acceptance couple the same whole vector in the two states and retain the probability of its entire original window I_0^J. Strict activation makes the captured occupation independent of that vector for this retained original row. It makes no assertion about rows that are not activated. Before summing rows, restore their original weights Z_i/d; sum the actual subkernels and use only the one original frame normalizer d. Rejection subtracts this actual subkernel. These are precisely the residual-product identities of the cited lemma, so causality and equality of the accepted normalized kernels persist. □

Remark 10.3. The exact-core assertion here is for the drift-free pure references. If a known nonharmonic drift v_e is present, matching $F + v_e - \kappa_e$ on a boundary neighborhood need not match its conditional extension inside. A separate coherent-mean estimate is needed in that case. A scalar or a harmonic drift causes no such problem.

Corollary 10.4 (Fixed exterior wiring)**.** *The same minorization and reference comparison apply to an ordinary query with a fixed positive exterior wiring, including compulsory outer arms. Precisely, roots lie in a strict buffer inside an open domain H_0; inside that buffer each row path minimizes the ordinary metric and every replacement between two of its points is admissible for the same row. The complete ordinary frame retains its row label, its one normalizer, and the loop-cap moments proved above. A row such as $S_i + k_{ij} + S_j$ need not be the unconstrained global distance between its outer ports.*

Proof. The crossing counts, loop shortcuts, source savings, and two-anchor forcing use only the stated admissible local replacements. The exterior wiring is retained in the old state and is never modified by a comparison in the strict buffer. The forced anchors have unchanged protecting brackets on both sides. Consequently their middle is a strict interior subsegment of an actually minimizing row fragment compactly contained in H_0. The internal version of simultaneous strict-interior uniqueness proved above makes this middle unique almost surely, also after the finite Cameron–Martin shifts and annular disintegration. This substitutes for the deterministic-endpoint witness in the bare path formulation; it does not assume uniqueness for arbitrary field-selected endpoint pairs. The normalized conditional core and its occupation in the root cell are therefore precisely the common kernel used in (10.3). Carry out the coupling with the row label retained. When summing rows, retain their original conditional masses before dividing by the one complete-frame normalizer. All count and moment arguments then apply to this labelled frame, with its fixed finite row cutoff. □

11 Initial capture and uniform restart

The pure law modulo constants is invariant under dilation. Pinning before dilating is pinning by ℓ afterward. By (7.1) and deterministic endpoint uniqueness,

$$\left(F^{[H]}(H\cdot), H^{-aQ}\nu_H(H\,\mathrm{d}y)\right) \overset{\text{law}}{=} \left(F^{[1]}, \nu_1\right). \tag{11.1}$$

To prove positivity of the unit capture mass, choose nested tubes $T_0 \Subset T_1 \Subset T_2$ from p to q, through a smaller capture disk $C_* \Subset C$. Require $T_2 \setminus C_*$ to separate the two endpoint components. Choose $s > 0$ with $T_1^{2s} \Subset T_2$, and use the entire wall $\{s \le \operatorname{dist}(x, T_1) \le 2s\}$. The internal T_0 route is finite; its wall crossing and the C_*-to-∂C crossing are positive. Hence with positive ordinary probability their costs are at most A_0, at least $B_0 > 0$, and at least $C_0 > 0$, respectively. A smooth shift with floor $-M$, equal to $-M$ on T_0 and to K on the full wall, vanishes near the pin. Choose $e^K B_0 > e^{-M} A_0$. Every shifted minimizer hits C_* and has at least $e^{-M} C_0$ occupation in C. Gaussian quasi-invariance proves a positive ordinary capture mean. An exterior rectifiable loop caps all capture occupation and has every positive moment. Equation (11.1) proves all normalizer and initial moment assertions uniformly in H.

For restart counts first use a larger capture consisting of a fixed finite union of aligned initial grid cells, and the same endpoints and pin. All descendant root cells are then wholly in that capture. Call this law $\mathsf{Q}_H^{\square}$. Choose the union to contain $\operatorname{supp} w_H$ and a fixed buffered enlargement. The preceding forcing and loop bound give

$$\mathsf{Q}_H \le C \mathsf{Q}_H^{\square} \tag{11.2}$$

as measures on the field and root, with a constant independent of H. The initial count argument in Lemma 9.6 applies directly to $\mathsf{Q}_H^{\square}$. There is no endpoint or capture-boundary exception for these aligned starts.

Lemma 11.1 (Restart without absolute-depth entropy)**.** *For every prescribed count rate, the numerical GOOD cutoffs can be fixed so that, under $\mathsf{Q}_H^{\square}$ or a GOOD-started conditional path/frame law,*

$$P(\text{no GOOD in } [i, i+n]) \le C e^{-cn} \tag{11.3}$$

uniformly in H and in the absolute depth i. For any fixed $\eta > 0$, the same is true of the event that more than ηn levels in this window are non-GOOD. Both conclusions transfer to Q_H by (11.2). For two causally coupled histories, the probability of no successful common annulus in a comparison window of n separated scales is at most $Ce^{-c_* n}$, provided its cells are eligible for their bare Campbell kernels.

Proof. Let T_0 be the first GOOD time. Lemma 9.6 gives $P(T_0 > n) \le Ce^{-cn}$ and, on $G_k = \{k \text{ is GOOD}\}$,

$$P(\text{no GOOD at } k+1, \dots, k+n \mid \mathcal{F}_k) \le Ce^{-cn}.$$

If there is no GOOD in $[i, i+n]$, either no GOOD has yet occurred, or some last GOOD $k < i$ occurred. Drop the future last-GOOD restriction before conditioning:

$$P(\text{no GOOD in } [i, i+n]) \le Ce^{-c(i+n)} + \sum_{k<i} Ce^{-c(i+n-k)} \le C' e^{-cn}.$$

There is no union over the b_0^{2i} possible addresses of the old history. For the bad-proportion estimate, find the first GOOD in the first $\eta n/4$ levels. Its absence has an exponential bound by the preceding calculation. Conditional on its occurrence, apply the conditional prefix count to the rest of the window. The original prescribed rates absorb the fixed fractions of the window.

For two histories, outside two exponentially small events, at least $(1 - 2\eta)n$ levels are GOOD for both. A greedy collection separated by the fixed reveal delay still has a positive multiple of n members. At each predictable opportunity Lemma 10.2 gives conditional success probability at least p_*. Sequential conditional expectation bounds failure at m opportunities by $(1 - p_*)^m$. Each rejection is exterior to the hidden core, so the prescribed marginal transitions and the GOOD estimates remain valid. Stop at the first success. □

12 Finite-chaos moments at the actual root

The following identity uses the pure Dirichlet precision. If v is supported in a fresh disk D_e, then (8.1) and harmonicity give

$$I_v(F) = I_v(X_{D_e}), \qquad E[I_u(X_{D_e})I_v(X_{D_e})] = (u,v)_{\mathcal{H}_a}\,. \tag{12.1}$$

Thus the entire collection of compact source scores in that disk has exactly its ordinary pure Gaussian law, independently of the exterior. No bound on an uncentered physical scalar is involved.

Lemma 12.1 (Rooted chaos tail)**.** *Let Z_c be a translated local Gaussian-chaos envelope, of finite degree at most $d \geq 1$, for an observation cell of size D. Assume its scores have a fixed strict buffer and its ordinary Gaussian Sobolev norm is bounded by M. Then*

$$\sup_{H \geq CD} \mathsf{Q}_H\{|Z_{c(x)}| > tM\} \leq C_d \exp(-c_d t^{2/d}), \qquad t \geq 1\,. \tag{12.2}$$

The same conclusion holds under an eligible GOOD-started conditional path or frame law. Constants may depend on the fixed chaos and its shape, but not on the reference start depth.

Proof. Gaussian hypercontractivity and Sobolev embedding give

$$P\{\sup_{x \in c} |q(F,x)| > tM\} \leq C_d e^{-c_d t^{2/d}}\,.$$

One may obtain the vector-valued bound directly by Minkowski applied to the finitely many output derivatives and scalar chaos hypercontractivity. Finite feature limits preserve it. By (12.1), the same ordinary bound holds below a GOOD ancestor whenever the complete observation support lies in its fresh disk. The one-path Campbell density is bounded; in a frame, (9.4) and the lower normalizer give every fixed finite density moment. A fixed Hölder inequality therefore preserves the displayed tail, with a smaller c_d.

An ancestor k levels earlier has at most Cb_0^{2k} possible descendant observation cells. Consequently

$$P\{G_{n-k},\ |Z_{c(x)}| > tM\} \leq Cb_0^{2k} e^{-c_d t^{2/d}}\,. \tag{12.3}$$

Do not condition this estimate on that ancestor being the last GOOD one. The depth since the last GOOD state is separately bounded by Ce^{-ak} from Lemma 11.1. Put $K = \lfloor \varepsilon t^{2/d} \rfloor$, with $\varepsilon > 0$ sufficiently small. If $K \leq n$, sum (12.3) over the possible GOOD ancestors in the last K levels, and bound their absence by Ce^{-aK}. This yields (12.2). If $K > n$, use the initial capture moment and one Hölder inequality before the union over at most Cb_0^{2n} cells. Since $n < \varepsilon t^{2/d}$, its entropy is absorbed in the same tail. Bounded domination (11.2) transfers the result from the aligned capture to the smooth one. Integrating this tail proves (7.4). □

13 Boundary cells, Cauchy convergence, and scalar descent

13.1 Capture labels at the comparison window

An initial capture indicator must not be silently discarded from a child Campbell kernel. At observation mesh R, retain the radius label t, and call a child cell a boundary cell if it meets $\partial C_{H,t}$. Its diameter is at most $C_0 R$. For each fixed spatial point x, the set of t for which its cell is a boundary cell has length at most CR/H. Therefore Fubini gives

$$\mathsf{Q}_H(\text{boundary cell at mesh } R) \leq \frac{CR}{Hd_H} \mathbb{E}\nu_H(C_H^+) \leq C'R/H\,, \tag{13.1}$$

where C_H^+ is one fixed enlargement containing all these cells. Its occupation is loop-capped, and its mean divided by d_H is bounded. On every remaining root cell the initial capture indicator equals one throughout the cell. All later refinements are contained in that cell. Conditional on its exterior and the retained radius label, their core laws and transitions are exactly (9.3). The radius label imposes no further condition on the core.

Proof of Theorem 7.1. The normalizer and initial moments were proved in Section 11; the local mark moments follow from Lemma 12.1. Compare starts $H_k \geq H_j$ in the common window with top and bottom radii

$$R_+ = D^{1/3} H_j^{2/3}, \qquad R_- = D^{2/3} H_j^{1/3},$$

rounded to compatible grid radii. This window has a positive multiple of $\log(H_j/D)$ separated scales. By (13.1), the probability that either initial capture cuts its top cell is at most $C(D/H_j)^{1/3}$. On its complement each marginal has the correct bare Campbell transitions through the window. Its GOOD count has a uniform absolute-start bound by (11.2) and Lemma 11.1. The annular coupling therefore fails to match a canonical normalized core with probability at most

$$C(D/H_j)^{\eta_0}$$

for some $\eta_0 > 0$.

Every accepted core contains the full D-neighborhood of its root, because its scale is at least $R_- \gg D$. The common field is determined modulo the two old scalars, which do not affect compact scores. For equal phases, grid-center translations are integer multiples of each finer mesh, so the matched roots have the same fractional position in the original source grid. More generally compare phases τ_1, τ_2 by using the nested grid with origin τ_i in the ith exploration. At a common radius R, its cell center is $z_i \in \tau_i + R\mathbb{Z}^2$. Centering at z_i sends each finer source lattice to the same lattice: R/r is an integer, and the residual displacement $\tau_i - z_i$ belongs to $r\mathbb{Z}^2$. The pure Gaussian comparison and minorization are invariant under these translations. Their constants, the initial address count, and the boundary estimate are uniform in the phase. Thus the same coupling compares arbitrary two phases, without averaging either fractional phase. The complete mark (7.3) is identical on success. For $p > 1$, Hölder gives

$$|Eq_1 - Eq_2| \leq (\|q_1\|_p + \|q_2\|_p) P(\text{failure})^{1-1/p} \leq C_q (D/H_j)^{\eta_0(1-1/p)}.$$

This proves (7.5), also uniformly for two different phases of the same translated prototype. Taking both reference sizes to infinity proves phase independence of the limit. Changing the initial template or smooth radius mixture leaves the same intermediate window, GOOD counts, minorization, and boundary estimate. The limit is therefore the same. The reference is generated independently with its copied physical phase; no independent fractional phase is averaged. □

Corollary 13.1 (Periodic whole-owner masks)**.** *Let $P = \sum_{c \in \tau + r\mathbb{Z}^2} P_c$ be a translated complete orbit, and let R_-/r be an integer. Fix a deterministic declared carrier for each translate, and let $m_\theta(c) \in \{0,1\}$ indicate that this entire carrier avoids the boundaries of the mesh-R_- grid with coset θ. The coset ranges uniformly over fine-lattice translations modulo R_-. Every coarser grid in the nested hierarchy uses this same coset. Then*

$$\Pi\left(\sum_c m_\theta(c) P_c\right) = p_{\text{good}} \Pi(P), \qquad p_{\text{good}} = E_\theta m_\theta(c). \tag{P2}$$

In particular, masking both terms of a centered orbit gives another centered orbit for each fixed coset. If the carrier has diameter at most $D \ll R_-$, then $1 - p_{\text{good}} \leq CD/R_-$.

Proof. For fixed R_-/r, the masked expression is a finite translated prototype repeated on the coarser lattice $R_-\mathbb{Z}^2$. Theorem 7.1 applies with a declared prototype radius enlarged to $O(R_-)$. Its phase-independence assertion shows that the center is the same for all fine-coset translates of this masked prototype. Averaging those finitely many cosets gives $p_{\mathrm{good}}P$ exactly, so linearity gives (P2). The probability of a grid line meeting a fixed diameter-D carrier is at most CD/R_- in each coordinate. A coarser nested grid has no additional boundary lines. On an accepted core at radius $R_{\mathrm{acc}} \geq R_-$, the matching translation is a multiple of R_{acc} and hence preserves this mask exactly.

The auxiliary coset remains averaged inside the estimating representation. Its bad branch can be integrated before any field estimate and equals $(1 - p_{\mathrm{good}})P$. No part of a signed current is cut off. For the surviving whole owners every original label and retained read has one trial address, so a revealed failure flag is one exterior scalar on the whole primary. The auxiliary finite-prototype reference limit is taken after r, R_- are fixed; its qualitative convergence does not assert a raw-chaos constant uniform in R_-/r. □

Corollary 13.2 (Pairing after a designated past)**.** *Let $\mathcal{B}$ be a designated-past sigma-field. Conditional on $\mathcal{B}$, let e be an eligible GOOD exterior state for the pure field, with its actual path or union-frame Campbell kernel Π_e. Suppose the complete finite current to be compared is contained in a strict central disk of radius D, and there is a source-free exploration window of n separated levels between that disk and the boundary of the current state. Every original observation used by the current is in that complete disk. The past $\mathcal{B}$ reveals no additional variable from the unrevealed Dirichlet core. The physical grid phase Θ may be $\mathcal{B}$-measurable.*

For each sufficiently large H, generate an aligned finite reference independently of $\mathcal{B}$, with the copied phase Θ. There is a conditional coupling preserving both prescribed marginals, with

$$P(\text{failure} \mid \mathcal{B}) \leq Ce^{-cn} + CR/H, \tag{P1}$$

where R is the top window radius. On success the normalized fields and roots agree on the complete disk. The flags at each attempted level are measurable from its revealed annulus and independent auxiliaries. In particular none of them inspects the complete unrevealed payload. The constants are uniform over the eligible GOOD states and copied phases.

Write q_H^{ref} for the reference current and $\beta_H(\theta) = E_{\mathsf{Q}_H} q_\theta$, and write $\Pi(q)$ for their common limit from Theorem 7.1. For every integrable $\mathcal{B}$-measurable complex variable Y,

$$E[Y q_H^{\mathrm{ref}}] = E[Y \beta_H(\Theta)], \tag{13.2}$$

$$\left|E[Y q_H^{\mathrm{ref}}] - \Pi(q) EY\right| \leq C_q (D/H)^{\eta_q} E|Y|\,. \tag{13.3}$$

For each fixed finite-chaos prototype and $p > 1$, its original and reference currents satisfy

$$\left|E[Y(q - q_H^{\mathrm{ref}})]\right| \leq C_{q,p} E|Y|\,[e^{-cn} + R/H]^{1-1/p}\,. \tag{13.4}$$

The same assertion holds for a family having one common conditional L^p envelope, with its conditional envelope bound retained on the right.

Proof. Disintegrate over $\mathcal{B}$ before generating the fresh reference. Each ordinary conditional core is the independent Dirichlet field plus its known harmonic prediction. Apply Lemma 10.2 whenever both histories are GOOD, and otherwise keep the prescribed marginal transition. The acceptance density displayed in that lemma depends only on the revealed annulus; its rejection does not condition a future core. Consequently induction preserves both complete conditional marginals.

Lemma 11.1 supplies the exponentially many-opportunity bound, and (13.1) supplies the reference capture error CR/H. Every accepted disk is chosen outside the entire old payload. This proves (P1) and exact equality of compact pure Wick marks on success.

Marginal preservation is essential here: after coupling, the conditional law of the second sample is still its freshly generated reference law. It follows that $E[q_H^{\mathrm{ref}} \mid \mathcal{B}] = \beta_H(\Theta)$, which proves (13.2). The phase-uniform Cauchy estimate and phase independence proved above give $|\beta_H(\Theta) - \Pi(q)| \le C_q(D/H)^{\eta_q}$ uniformly; this proves (13.3). Finally apply conditional Hölder to the failure event, using the conditional GOOD-started version of Lemma 12.1, and then integrate $|Y|$. □

This corollary also explains the order of an array calculation. For the ordinary conditional frame, $P_e(dF)\nu_e(dx) = Z_e\Pi_e(dF, dx)$. The exterior factor Z_e is retained when pairing each array entry. Thus a past-measurable coefficient Y_e is replaced by Y_eZ_e in (13.2)–(13.4). Sum the ordinary array entries first, and divide only by the one original whole-frame normalizer. One does not normalize each term to have unit mass and then sum. For positive assignment weights the same argument uses their sum one before applying a triangle inequality. A known nonharmonic mean or a covariance likelihood remains part of the full integrand; the corollary does not turn such a field into a pure Markov field.

Remark 13.3 (A selected past need not be GOOD). For an array selected below an initial aligned capture or a GOOD-started whole conditional frame, its selected exterior need not itself be GOOD. The same conditional marginal-preserving construction is still defined: run the two prescribed transitions without coupling until paired GOOD opportunities appear. Equations (13.2) and (13.3) remain exact and uniform, respectively. The conditional probability estimate (P1) is then not asserted. Instead average over the original whole labelled root law and apply the absolute-start count estimate of Lemma 11.1 to that marginal, and the same estimate to the independent reference. This gives the unconditional failure bound $Ce^{-cn} + CR/H$. A bounded designated-past multiplier costs its supremum norm; an integrable random envelope is handled by one joint Hölder inequality under this whole law. All old failure indicators remain in that joint envelope. This argument does not assume a uniform conditional estimate for an arbitrary selected past, and does not assert the count theorem for an arbitrary new exterior-weighted root law.

13.2 The complete scalar functional

Let $(P_c)_c$ be a translated finite physical expression of positive source degree, with original first-current identity

$$\mathbb{E}[P_c f(K_H)] = \mathbb{E}\int f'(K_H)q_c(F, x)\nu_H(\,\mathrm{d}x), \qquad q = \sum_c q_c\,. \tag{13.5}$$

All compact directions preserve the exterior pin for large H. For a term $P_c = D_uQ_c$, pure Gaussian integration by parts makes its current $u(x)W_{Q_c}(F)$, where W_{Q_c} is the Wick adjoint of the remaining compact directions. Finite feature integration gives the covariance-source version. This proves (13.5) for smooth bounded tests with bounded first derivative.

Choose smooth bounded functions $f_M : [0, \infty) \to \mathbb{R}$ such that $0 \le f_M' \le 1$ and $f_M'(t) \to 1$ for every finite t. The distance K_H is finite almost surely. For each fixed H, the absolute current on the finitely many relevant source supports is integrable: a surrounding circuit bounds the occupation of their union, and the local Gaussian current has moments of every finite order. Hence dominated convergence applies to the right side of (13.5). This defines the scalar limit below without a first-moment assumption on K_H.

Assume $q_c(F,x) = 0$ unless $|x - z_c| \le D$. Define

$$\beta_H(P) = d_H^{-1}\mathbb{E}\int w_H(x)\sum_c q_c(F,x)\nu_H(\,\mathrm{d}x), \qquad s_H(P) = \lim_{M\to\infty} d_H^{-1}\sum_c w_H(z_c)\mathbb{E}[P_c f_M(K_H)]\,.$$

Only finitely many anchors occur in s_H. Apply (13.5) before passing to the limit in M. The integrable local current gives the exact identity

$$s_H(P) - \beta_H(P) = d_H^{-1}\mathbb{E}\int \sum_c [w_H(z_c) - w_H(x)]q_c(F,x)\nu_H(\,\mathrm{d}x)\,. \tag{13.6}$$

The original label is x, before any estimating current is moved to a new arm. The absolute sum of the finitely many original local marks has uniformly bounded rooted moments by Lemma 12.1. Use the enlarged aligned capture to cover points within distance D of $\operatorname{supp} w_H$. Since $|w_H(z_c) - w_H(x)| \le CD/H$ on the support of q_c, we obtain

$$|s_H(P) - \beta_H(P)| \le C_P D/H\,. \tag{13.7}$$

Thus $s_H(P)$ has the same limit as the rooted expectation. Define this limit to be $\Pi(P)$.

Physical equality means equality of the complete unmarked scalar actions in the same anchored owner sector, for every admissible exterior and full Gaussian mean family. Such expressions have equal truncated scalar actions for every M, and hence equal s_H for every finite H. Equation (13.7) proves representation independence of Π. Linearity follows before either limit. For an assembled partition of unity $\sum_c \psi_c = 1$, its current is identically one, so $\Pi(\{D_{\psi_c}\}_c) = 1$. We denote this complete drift orbit by D_1.

If an exact unary reassignment has $P'_Q = \sum_c K(c,Q)P_c$ with $\sum_Q K(c,Q) = 1$, its original current can be chosen as $q'_Q = \sum_c K(c,Q)q_c$. The complete assembled marks are identical pointwise. Evaluate both in one phase-aligned reference larger than both read carriers and pass to the limit. This gives $\Pi(TP) = \Pi(P)$ and the usual intertwining of the owner-preserving drift lift with exact unary reassignment.

14 The precise interface with a structural activity norm

The reference construction does not manufacture an ordinary moment estimate for a legal activity record. That estimate is supplied by the joint likelihood and one-current structural theorem. Here is the exact transfer it needs.

Proposition 14.1 (One-envelope transfer)**.** *Let a normalized nonnegative envelope V have the following ordinary query bound: below every eligible GOOD ancestor, each prescribed carrier query has one common labeled-current envelope with a fixed $p_{\mathrm{ord}} > 1$ moment at most C, including its one loop/crossing restoration factor. At a separation of k levels there are at most Ce^{dk} possible carriers. The original Gaussian observation list, source contours, maximal opaque children, and old failure indicators remain in this single envelope. Fix $1 < p_0 < p_{\mathrm{ord}}$ and reserve the GOOD Campbell moment of order $p_{\mathrm{ord}}/(p_{\mathrm{ord}} - p_0)$. Then its actual rooted tail satisfies*

$$P\{V > t,\ \text{depth since the last GOOD state} = k\} \le \min\{Ce^{-ak}, Ce^{dk}t^{-p_0}\}\,. \tag{14.1}$$

Consequently V has every moment $p < p_0 a/(a+d)$, uniformly in the reference start. The same conclusion holds after a GOOD starting exterior. The count rate a can be prescribed before fixing the final norm constants, so it can be chosen to leave one final exponent strictly greater than one.

Proof. The first bound in (14.1) is the restart estimate. For the second, drop the future last-GOOD restriction, use the stated one-envelope moment under the actual ordinary query, and apply Hölder with exponents p_{ord}/p_0 and $p_{\mathrm{ord}}/(p_{\mathrm{ord}} - p_0)$ to the envelope power V^{p_0} and the GOOD Campbell density. The latter moment is bounded by the ordinary frame moment and the lower normalizer bound. This gives a bounded rooted p_0th moment for each candidate. Markov's inequality and the at most Ce^{dk} candidate carriers give the second bound. The initial no-GOOD term is treated by the initial loop moment and the same joint ordinary observation estimate, before a future failure is imposed.

Split the sum over k at $k_* = p_0 \log t/(a+d)$. Below k_* sum the second geometric bound; above it sum the first. Both sums are bounded by $Ct^{-p_0 a/(a+d)}$. Integrating the tail proves the claim. The proof has not conditioned a Gaussian or current moment on the future event that its ancestor is the last GOOD one. □

Corollary 14.2 (The quadratic owner loss)**.** *Suppose every legal exact estimating representation $\mathcal{R}$ of a complete owner-X expression has the one-envelope query bound in Proposition 14.1, with its assigned cost* $\mathrm{cost}(\mathcal{R})$*, uniformly on the permitted conditional mean range. Let $n = |X|$ and suppose the complete read carrier has radius $D \le Cnr$. Then*

$$|\Pi_X(P)| \le Cn^2 \,\mathrm{cost}(\mathcal{R}), \qquad |\Pi_X(P)| \le Cn^2 a_X(P)$$

after taking the infimum over legal representations. This is a bound on the complete scalar functional; it does not assert equality of two different marked current primitives at finite H.

Proof. Apply the chosen exact representation to each bounded scalar action $P_c f_M(K_H)$ before taking absolute values. The bound $|f'_M| \le 1$ makes every current estimate below uniform in M. The coefficient $w_H(z_c)$ is nonnegative, at most one, and independent of the Gaussian variables. The current is supported in a fixed enlargement of the radius-D carrier. Split the fine anchor lattice into m^2 residue classes, with $Cn \le m \le 2Cn$. Within one class these enlarged carriers have bounded overlap. At an ancestor of radius De^k there are at most Ce^{2k} candidate anchors in that class. Proposition 14.1 gives the required uniform actual-root first moment for the common spatial envelope. Lemma 21.3 below records this translated-array estimate explicitly; its proof uses the spatial invariant established in the same finite construction induction as Proposition 21.2. The enlarged capture occupation divided by d_H is bounded by the circuit estimate and the initial lower normalizer. Thus the absolute truncated scalar action in one residue class is at most $C\,\mathrm{cost}(\mathcal{R})$. Sum the m^2 classes, let $M \to \infty$, and then let $H \to \infty$. Finally take the infimum over exact legal representations. This proves the bound on Π_X. □

15 Positive refinement of complete owners

We fix one physical lattice phase for every source, reference, and scale. The coordinates in the formulas below are measured from that phase. Private phases used in a local geometric selection are auxiliary variables and do not change this phase.

Choose an integer m larger than every spatial derivative order used below, and let B_m be the cardinal spline obtained by convolving m copies of $\mathbf{1}_{[0,1]}$. Put $\psi(x) = B_m(x_1)B_m(x_2)$ and $\psi_{r,c}(x) = \psi(x/r - c)$. A fixed translation of the template is harmless. These functions form a nonnegative partition of unity. If additional smoothness is required in an initial identity, mollify the finite set of physical directions and pass to the limit in the required Cameron–Martin and finite-order spatial norms. The actual infinite-scale partition remains the one fixed here.

For an integer $b \geq 2$, let $N_{m,b}(k)$ be the coefficient of t^k in $(1+t+\cdots+t^{b-1})^m$, and set $h_b(k) = b^{1-m}N_{m,b}(k)$. The one-dimensional refinement identity is

$$B_m(x) = \sum_k h_b(k) B_m(bx-k)\,.$$

It follows by writing each independent uniform variable on $[0,1]$ as $(J+U)/b$, where J is uniform on $\{0,\dots,b-1\}$ and U is uniform on $[0,1]$. In two dimensions define

$$K_b(c,Q) = h_b(c_1 - bQ_1)h_b(c_2 - bQ_2)\,.$$

Then

$$K_b(c,Q) \geq 0, \qquad \sum_Q K_b(c,Q) = 1, \qquad \psi_{br,Q} = \sum_c K_b(c,Q)\psi_{r,c}\,. \tag{15.1}$$

The residue sums of $N_{m,b}$ are b^{m-1}: after the first $m-1$ residues are chosen, the last is determined. This proves the middle identity. Successive refinements compose to K_{b^L}.

A complete owner records all ranges already used by its physical expression, including both covariance legs, all original observations and coordinate supports, its protecting domains, and the whole support of its current labels. An assignment acts on the entire owner. It does not select a subcell of an already estimated signed current. A future observation enlarges this declaration only when it is actually introduced.

Lemma 15.1 (Whole-owner geometry)**.** *There are fixed geometry constants C_0, C_1 such that, after choosing b_0 sufficiently large, the following assignment is possible. A connected owner of n cells at scale r can be positively assigned by K_{b_0} to an owner at scale $b_0 r$ with*

$$n_{\mathrm{out}} \leq \theta n + C_0, \qquad \theta < 1/4\,. \tag{15.2}$$

Its relative output type is independent of the fine anchor residue and of the geometric trial outcome. For k compatible inputs grouped into one connected component,

$$n_{\mathrm{out}} \leq \theta \sum_{i=1}^{k} n_i + C_0 k\,. \tag{15.3}$$

The number of immediate backward anchor choices is bounded by a constant depending on b_0. Unread assignment prefixes have total mass one.

Choose N_* large enough that $\theta N_* + C_0 \leq N_*$ and put $\Phi(n) = n\mathbf{1}_{\{n>N_*\}}$. A genuine unary coarsening of an owner with $n > N_*$ decreases this potential by at least cn. More generally,

$$\Phi(n_{\mathrm{out}}) \leq \theta \sum_i \Phi(n_i) + C_1 k\,. \tag{15.4}$$

Same-scale grouping has the separate bound $\Phi(|\bigcup_i X_i|) \leq \sum_i \Phi(|X_i|) + N_* k$.

Proof. A spanning tree through the n cells has Euclidean length at most Cnr. Cover it at spacing comparable to $b_0 r$. Its fixed-width neighborhood meets at most $Cn/b_0 + C$ coarse cells. Include every source and protecting template already declared; by the definition of the current owner these are contained in the same fixed enlargement. The finitely many future domains actually read at this coarsening add a fixed number of coarse cells.

For an anchor c assigned to Q, include the union over every residue $c - b_0 Q$ with $K_{b_0}(c, Q) > 0$. These residues lie in a fixed-radius ball in coarse coordinates. Taking their union therefore changes only the fixed enlargement constant. It produces one relative output type, independent of which residue or trial outcome was realized. Choosing b_0 after these fixed shape constants proves (15.2).

For a connected family of inputs, join their spanning trees along a tree of the overlap graph. Each additional joining segment has length $O(r)$ times the fixed template constant. The same covering gives (15.3). An immediate coarse anchor has only $O(b_0^2)$ possible fine anchors because the refinement mask has fixed scaled support. The last assertion about weights follows from the middle identity in (15.1), iterated before any absolute current estimate.

If $n > N_*$, increase N_* so that $C_0 \le (1-\theta)N_*/2$. When the output remains large, its potential decrease is at least $(1-\theta)n - C_0 \ge (1-\theta)n/2$; when it becomes small, the decrease is n. For several inputs, replace each small n_i by N_* in (15.3). This proves (15.4). The same-scale inequality follows directly from $|\bigcup_i X_i| \le \sum_i |X_i|$ and $n_i \le \Phi(n_i) + N_*$. □

Internal potentials telescope along a finite genealogy. After suppressing unary vertices, a forest with m leaves has $\sum_{\deg(v)\ge 2}(\deg(v) - 1) \le m - 1$. Thus total large-owner coarsening work is bounded by the initial potential plus a fixed multiple of the active leaves. Empty reassociation performs no Gaussian evaluation and has no large-owner charge.

15.1 Current labels and exact scalar transport

For an original output carrier K at scale r, use the separable test space

$$\mathcal{T}_r(K) = C^1(K), \qquad \|f\|_{\mathcal{T}_r} = \|f\|_\infty + r\|\nabla f\|_\infty .$$

The current label lies in its dual. At a coarser scale R, the unit test ball restricts to the finer one. Freezing a label over a set of diameter d therefore costs at most d/R times the absolute current mass. The original output point remains the label after its current is represented on a fresh physical arm.

The spline mask also permits smooth coarse localization. For fixed m, finite differences of the truncated-power formula for B_m give

$$\sup_{c,Q} |K_b(c, Q) - \psi_{br,Q}(rc)| \le C/b . \tag{15.5}$$

For completeness, in one dimension use $N_{m,b}(k) = \sum_{\ell=0}^{m} (-1)^\ell \binom{m}{\ell} \binom{k - b\ell + m - 1}{m-1}_+$ and $B_m(x) = (m-1)!^{-1} \sum_{\ell=0}^{m} (-1)^\ell \binom{m}{\ell} (x-\ell)_+^{m-1}$. Uniform polynomial approximation on the fixed support gives the bound; products give the two-dimensional version.

If an old current μ_c has output within Cr of its anchor, then

$$\left| \int \{K_b(c, Q) - \psi_{br,Q}(x)\} \, d\mu_c(x) \right| \le C b^{-1} \|\mu_c\|_{\mathrm{TV}} .$$

For an owner of bounded cardinality the same estimate has the corresponding fixed cardinality factor. It remains valid after coherent summation under one ordinary occupation measure: the positive assignment and root integral are summed before taking the common absolute-current bound. There is no additional count of the fine cells.

In contrast, scalar drift transport is exact:

$$\sum_c K_b(c, Q) D_{\psi_{r,c}} = D_{\psi_{br,Q}} .$$

The complete translated orbit is assembled before its scalar center is assigned. Summing over coarse anchors then gives the same formal physical constant shift D_1. This is an identity of local physical operators on bounded observations, not a statement that the whole-plane constant belongs to the Cameron–Martin space.

16 A source-free atlas of positive serial arms

Fix a positive pure base a. All constants in this section may depend on a. After translation and dilation, choose smooth disks

$$H_0 \Subset H_c$$

of radii comparable to R, with fixed positive buffer ratios. The entire old module H_{old}, including all its source supports, readers, retained records, and output labels, is contained in H_0. Choose a connected annulus $W \Subset H_c \setminus \overline{H}_0$ and three closed annular bands with successive strict buffers in W. All origins of the kernels to be represented lie outside the outermost band. A compatible filling may change the length functional or the allowed-path body only in H_0. In the unchanged annulus it permits every ordinary local replacement of a subpath. This includes positive directed allowed-path bodies inside H_0 and singular limiting fillings with ties; it does not impose uniqueness inside such a filling.

An exterior context is a finite positive graph attached outside the outer band. It is used only after the local construction, as a scalar distance test. It cannot select a local metric, path, template, or avoiding kernel. Actual applications retain the genuine outside accesses and a surrounding circuit, so their fixed exterior graph connects all port vertices in both directions. This supplies the compactness needed for the finite potential valuation. An artificial edge with an arbitrarily prescribed cap cost is not inserted.

Choose a rectifiable surrounding loop in an untouched outer annular region of H_c, and denote its length by U. It is formed from a fixed finite collection of transverse metric crossings. Its definition precedes all reserve coordinates, and every reserve vanishes on a neighborhood of the loop. All geometric cutoffs below are in units of U.

Theorem 16.1 (An outer arm atlas at one pure base)**.** *For every $\varepsilon > 0$ and integer $k \geq 3$, there are finite constants M, J, numerical constants $0 < c < C < \infty$, a number $\sigma > 0$, a finite catalogue of templates, and deterministic product intervals $I_0^J \Subset I_1^J$, with the following property. After an independent finite or uniform grid phase is chosen, reserve a finite deterministic menu of nonnegative smooth functions $g_1, \dots, g_J$ supported in W and write*

$$F = X + \sigma \sum_{j=1}^{J} A_j g_j, \qquad A_j \overset{\mathrm{iid}}{\sim} N(0,1), \quad A \perp X, \qquad \operatorname{Cov} X \geq \tfrac{1}{2} \operatorname{Cov} F\,. \tag{16.1}$$

There is a completed-measurable event $\Gamma(X, \mathrm{phase})$ and a first successful template selected from those data only such that

$$P\{\Gamma,\ A \in I_0^J\} \geq 1 - \varepsilon\,.$$

Every selected template has at most M ports and pairwise separated private arms. Each arm consists of k serial intrinsic pieces and fixed connectors. For every $A \in I_1^J$ and every compatible filling, its body kernel between outer ports satisfies

$$q_{ij} = S_i + k_{ij} + S_j \quad (i \neq j), \qquad S_i = c_i + \sum_{v=1}^{k} S_{iv}(A_{iv})\,. \tag{16.2}$$

The diagonal is $q_{ii} = 0$. The complete inner matrix k is finite, positive or nonnegative, and satisfies the directed triangle inequalities. It contains every inner body route. The inner body and the connectors are independent of the selected arm coordinates. The row q_{ij} is the kernel of the compulsory-arm body. The unrestricted ordinary distance between the same outer ports can be smaller because of an exterior bypass; that distance is recovered only after the retained exterior wiring is glued and its minimum is taken. For local rooted estimates, an ordinary query means this fixed exterior wiring together with free ordinary minimization in a strict interior buffer of H_0. Every local path replacement there is permitted. No global ordinary-geodesic identification of a restricted body row is used. The private supports are disjoint from every nonincident arm and from the inner body, and every selected direction is constant on each complete anchor coding buffer. They vanish on the outer attachment windows and on the retained unchanged exterior accesses. Each piece has a mandatory positive plateau passage and satisfies

$$cU \le S_{iv} \le CU, \qquad c\sigma U \le \partial_{A_{iv}} S_{iv} \le C\sigma U$$

almost everywhere on its enlarged interval. Conditional on the residual and nonselected coordinates, all selected coordinates are independent Gaussians restricted to their stated intervals.

All kernels, including the avoiding kernels, are intrinsic to H_c. Gluing this body to its retained avoiding graph recovers every admitted original scalar distance exactly. The local body, its actual arm curves, and any later forest replacement lie in one declared modification region $C \Subset H_c$. Every route avoiding C is retained. Its protecting region C^+ contains the coding buffers and a surrounding loop. The selected outer ports have genuine unchanged exterior accesses to that loop, recorded in the atlas.

The density of the total cost of each arm has the fixed C^1 and derivative bounds at scale U obtained by convolving three BV seed densities. Widening the Gaussian coordinate intervals supplies a positive cost-image margin cU. The constants and exceptional probability are uniform over physical radius and additive field scalar. They also transfer, with the fixed Hölder loss, to any retained conditional GOOD class with a uniform local Gaussian L^p comparison bound. No uniformity in a varying real base parameter is claimed here.

Proof. 1. A universal finite graph in the unchanged annulus. Apply Lemma 6.1 in finitely many buffered observation disks covering the middle annular band. Include all W-internal geodesics whose endpoints lie outside the corresponding larger disk. Their complete local traces have finitely many representatives. Retain the full representing geodesics in W. Interior uniqueness implies that two representatives cannot split and rejoin at points strictly interior to both. Consequently their intersections have finitely many components: any remaining component endpoint is an endpoint of one of the full representatives. Subdivide at these finitely many points and shared-interval endpoints. Their union is a finite embedded graph in W carrying all the required traces in the middle band. No subdivision at every contact with a band boundary is made.

This is an exact list of complete local traces. Its finiteness and moment bounds are supplied by Proposition B.15 and Lemma 6.1. The compactness used there is obtained by trimming at finite-length surrounding circuits, whose images are compact for the internal metric. It does not require a bound for the metric diameter of $\partial D_{\mathrm{out}}$ or of a Euclidean neighborhood.

Every whole minimizer under a compatible filling contributes only traces from this graph. Indeed trim it inside an unchanged buffered disk, around each visit to its smaller observation region. The trimmed fragment is a W-internal geodesic, since a shorter W competitor would shorten the whole path. Its endpoints clear the smaller observation disk. This argument is independent of the hidden filling and applies to every tie. It concerns local fragments; a whole path leaving W and returning need not be one catalogue representative.

The same family controls sufficiently small uniform Weyl perturbations. Around each observation disk choose a fixed finite-length surrounding circuit whose filled Euclidean interior is compact in W and disjoint from H_0. Trim a passage at its last contact with the circuit before the observation and its first contact afterwards. The trimmed path stays in that compact interior, and replacement along the circuit bounds its length uniformly. Its endpoints lie in the internal-metric-compact image of the circuit. After taking a subsequence they therefore converge in the internal metric. For perturbations tending uniformly to zero, Weyl comparison gives convergence of the corresponding endpoint distances as well. Lemma B.1 identifies the limit as a base internal geodesic. Its trace belongs to the fixed list. This proves the required stability using only unchanged local fragments.

2. A finite separating set of private edges. Mark disjoint finite closed subgraphs covering all inner-band and outer-band contacts of the finite graph. Such subgraphs exist by subdividing each finite arc finely enough; the two contact sets have positive separation. Set a graph potential equal to zero on the inner marked subgraph and one on the outer one, extend it harmonically at the remaining vertices with arbitrary positive conductances, and make it affine along every remaining edge. Ignore components meeting neither marked subgraph. Choose a level in $(0, 1)$ different from every vertex value. Its inverse image is a finite set of points in interiors of private edges and separates every inner-to-outer graph passage. The finite-graph maximum principle ensures that each retained point has a monotone outward graph continuation to the outer marked subgraph. This continuation later gives its genuine unchanged access to the surrounding loop.

On each selected private edge choose two outer bracketing anchor neighborhoods and k disjoint seed subintervals between them, each with its own two smaller bracketing neighborhoods. All closures have positive separation from nonincident graph pieces. The complete local trace lists in these neighborhoods give exact anchors: choose a point of the reference arc outside the finitely many isolated contacts and shared-interval endpoints of all representatives. Every representative containing that point then contains a common interval around it; every other representative has positive compact separation from a smaller interval. A sufficiently small core therefore forces every eligible visitor to contain its selected anchor. Retain the complete coding buffer used for this conclusion.

3. Ordered passage and the intrinsic arms. Base fragments meeting a selected private support pass its two bracketing windows in order. Local compactness in the preceding unchanged buffers makes this stable under a sufficiently small uniform Weyl change. Choose the perturbation directions constant on all coding buffers. In each such buffer the new internal metric is exactly a scalar multiple of the old one. Trim the new passage again inside a nested coding buffer, with endpoints outside its smaller observation core. Its intrinsic minimizing property makes this trimmed piece a geodesic for the old internal metric of that buffer. The fixed complete trace list and the chosen anchor core therefore force the same exact anchor. This second trimming, rather than metric closeness, supplies exact inclusion. Repeated crossings would repeat an anchor and can be removed from a positive minimizing path in the unchanged annulus.

A short reference segment between neighboring anchors has cost smaller than leaving a slightly larger private tube. The anchors are fixed at this step. Interior uniqueness and compactness therefore give a positive excess for any competitor that leaves the tube, and uniform Weyl comparison preserves that excess on a sufficiently small amplitude box. Hence its intrinsic tube kernel agrees with the W-internal kernel for these endpoints. Thicken the inner component and the private stumps so each arm has exactly its two specified attachment windows and no lateral attachment. Nonincident tubes and all selected supports have positive clearance. The remaining exterior pieces and their true avoiding kernels are retained.

Any minimizing traversal of an arm splits at its exact anchors. The lower bound in (16.2)

follows by applying the intrinsic kernel lower bound to every piece. The upper bound follows by concatenating arbitrarily accurate intrinsic competitors. This also proves the identity when an open-domain infimum is not attained. A nonincident stump cannot supply an out-and-back shortcut, which would repeat its attachment anchor at positive cost. The inner kernel is therefore independent of all selected arm coordinates. Completing its distances gives the directed triangle inequalities. The untouched inner annular routing network makes every entry finite, even if the filling in H_0 is a directed allowed-path body.

Infima for arbitrary compatible fillings. The preceding passage and splitting statements also apply when the full query infimum is not attained. To make this precise, take $W_0 = \overline{W}$, still a compact ordinary annular buffer disjoint from H_0, and a closed middle band $A \Subset W$ containing all selected supports and their complete coding neighborhoods. Use the finite-episode tightening argument from Theorem B.16, in the actual reconstructed metric. The uniform Weyl bound on the whole retained amplitude box gives one positive lower bound c_* for the cost from ∂W_0 to A, and one finite bound for the relevant query cap. The latter is supplied by the unchanged surrounding loop, or by the retained physical exterior accesses for a port frame. A vanishing-defect sequence therefore has only finitely many components of its visit to $\operatorname{int} W_0$ which reach A: every nonendpoint component costs at least $2c_*$. At most two endpoint components are allowed when the endpoints are in W_0 but clear A. Infinitely many raw contacts with ∂W_0 cause no difficulty.

Replace each such complete episode by a shortest path with the same endpoints in the compact closed set W_0. Its infimum is finite and attained by compact path convergence and lower semicontinuity. This replacement is admissible and does not increase cost for every compatible filling, since W_0 is unchanged by all changes in H_0. Every strict fragment of a replacement near A is a W-internal geodesic: a shorter W competitor would also be admissible in W_0. Thus its local traces and its amplitude-stable exact anchors are precisely those of the universal catalogue already constructed, in the same open comparison domain W. Split these tightened representatives at their actual anchors and delete repeated ordinary-anchor loops. Taking the defects to zero proves the lower finite-graph comparison. Concatenation of arbitrarily accurate representatives of the genuine intrinsic and exterior kernels proves the reverse comparison. These operations take place before a subsequent forest restriction; they never tighten a path inside an opaque filling.

For a complete scalar test, every tightened representative meeting a private support traverses its bounding anchors. Split it at successive anchors; segments avoiding the private supports belong to the retained true avoiding kernels. After ordinary-anchor loop deletion, a traversal uses each of the finitely many anchors at most once. The two comparisons just proved give exact equality with the admitted scalar infimum, including directed or tied compatible fillings and unattained hidden or exterior infima. Define C as a sufficiently thin neighborhood of the inner region together with the sublevel graph and the selected stumps, ending at their outer port windows. The strict-superlevel outward continuations are disjoint from this closed set except at their prescribed port attachments. Their finite compact pieces have positive separation from its nonincident parts. Shrink the neighborhoods and supports using these separations. Strict ascent of the piecewise affine graph potential never crosses the chosen level again, so none of these continuations meets another selected support. Follow them to the outer marked graph and then the unchanged surrounding loop. They give finite physical exterior cap routes between all outer ports, avoiding the complete modified body except at the prescribed attachments.

4. Deterministic menus and quantitative cutoffs. At the fixed pure base, graph size, trace counts, candidate costs, and modulus constants are finite almost surely. The private lengths, separations, positive exit-cost margins, and local Weyl tolerances are positive almost surely.

Exhaust these quantities by deterministic numerical bounds and rational buffered tube choices. A finite subcatalogue succeeds outside an arbitrarily small exceptional event.

Choose a sufficiently fine grid in W and smooth cell functions equal to one on smaller plateau squares and zero near cell boundaries. Reveal an independent grid phase after choosing the finitely many private reference points. The relative width of the transition strips can be chosen so small that all reference points lie in prescribed plateau or zero regions with the required probability. The grid mesh is smaller than the retained private separations. Choose the complete anchor coding buffers only now, inside regions where every menu direction is constant. Selected seed supports lie between their own bracketing buffers and miss all others. Within each plateau retain a mandatory positive-length crossing interval. Further finite numerical exhaustion bounds these lengths below by cU and all seed costs above by CU. For the auxiliary port frames, retain also the positive unchanged crossings in disjoint neighborhoods of the outer ports. Their numerical lower cutoff in units of U is justified in Section 20; it is selected here with the other finite exterior cutoffs.

Reserve the entire grid menu, not just the ultimately selected coordinates. Its supports have disjoint interiors and hence are Dirichlet-orthogonal. Choose σ so that its finite covariance subtraction is at most one half of the pure covariance. Choose the old box with Gaussian probability arbitrarily close to one, then a fixed larger box, and decrease σ until every reconstruction in that larger box differs by less than the retained Weyl tolerance. Couple $F = X + \sigma \sum A_j g_j$ with its residual. Whenever F satisfies the previously retained cutoffs and A lies in the old box, one finite template works for every reconstruction in the enlarged box. Define Γ to be existence of such a template and choose the first. This proves the probability assertion by a union bound. Independence between F and A is not asserted or needed. Given the residual and template, the selected coordinates retain their original independent laws.

The selection is completed-measurable from W. Here are the quantifiers in this assertion. On each deterministic length and modulus cutoff, the trimmed geodesic family is a compact subset of a Polish continuous-path space. The geodesic relation is Borel, since its distance equalities can be tested at rational times. Measurable compact-set selection gives a countable dense sequence of fragments. Internal strong confluence implies that this sequence contains every exact local trace class. Use its indices, rational arclength points, and rational buffered tubes to form a countable candidate catalogue. For a fixed candidate, failure of a claimed universal local passage property is witnessed by one amplitude vector in the compact box and one eligible geodesic fragment. The existence of this witness is an analytic event. Its complement is therefore universally measurable. Selecting the first valid candidate in the finite retained subcatalogue is completed-measurable; this does not require uniformizing an uncountable family of universally quantified fillings. All tests concern unchanged W-internal fragments. The all-fillings conclusion is a deterministic consequence of their local replacement property. Any needed approximately minimizing intrinsic curves are selected from their analytic existence relation. No projection is asserted to be Borel merely because the trace list is finite.

5. Densities, scale, and the conditional class. The plateau passage and first variation give the displayed slope bounds. Restricting to uniformly bounded-cost competitors shows semiconcavity of each seed cost. The monotone change of variables consequently gives a BV density; three independent pieces give the C^1 convolution and derivative bounds of Lemma 17.2. The enlarged intervals give a uniform positive image margin, separately on each fixed residual fiber.

Every construction is intrinsic to H_c and is unchanged by an additive physical scalar after division by U. Exact pure scaling therefore transfers it to every R. For a retained GOOD exterior, use the strict-buffer Gaussian L^p comparison to transfer the exceptional probability after choosing the ordinary error sufficiently small. The joint reserved-coordinate comparison has the same bound:

means and covariance perturbations are placed in the residual, while the reserved coordinates keep their centered independent law. At a test annulus all other original reserve menus vanish there, by the disjoint active-shell convention, so only the present finite covariance subtraction is used for its geometric inputs. No assertion about an infinitely depleted residual follows merely from covariance comparability. This completes the proof. □

16.1 The retained exterior class is an actual high-probability event

The additional exterior cutoffs used by the reset can be added to the atlas without changing its annular span. The tests use an unchanged surrounding loop and strict-interior Gaussian predictions. All fields in this paragraph are in rescaled coordinates: at radius R use $F(z + R\cdot) + aQ \log R$, and normalize that field by subtracting $\log U$.

Lemma 16.2 (Strict-interior prediction and an unchanged loop cap)**.** *Choose an unchanged ordinary loop Γ surrounding H_0, strictly inside the outer arm annulus, with all inner ports on its exterior side. The additional exterior tests can be chosen outside an arbitrarily small ordinary event so that its length is at most M_0U, the centered harmonic prediction has bounded smooth norm on the required strict interior buffers, and the retained middle mean score has bounded dual norm. All these tests are exterior to H_0 and use the same fixed observation block. On the resulting class,*

$$m_i \le K_i^{\mathrm{old}} \le B_i \le m_i + M_0U$$

for every physical old row retaining the genuine avoiding routes. A row with $m_i = B_i$ is constant. No access to arbitrary boundary points is part of the statement.

Proof. Construct Γ using a fixed finite network in an untouched annulus, before selecting the outer reserve supports. All inner ports and their attachment neighborhoods can be placed in the arm annulus, on its exterior side. The ordinary positive loop moments and inverse moments of the unit U, obtained from [3, Proposition 2.9], make $\mathrm{len}(\Gamma)/U$ tight. Take its deterministic upper cutoff.

For a compact lowering f_n in H_0, take an ordinary competitor of cost within η of $K_i(F - f_n)$. If it avoids H_0 it is already an avoiding competitor. Otherwise its portion entering H_0 has first and last contacts with Γ. Replace the entire intervening part by a loop arc. The exterior field is unchanged, so this is a genuine avoiding competitor of cost at most $K_i(F - f_n) + \eta + \mathrm{len}(\Gamma)$. Let $\eta \downarrow 0$ and then $n \to \infty$ using Lemma 9.2. This proves $B_i \le m_i + \mathrm{len}(\Gamma)$. Physical lifting gives the other inequalities. The same argument may be applied directly to an actual free-hole competitor, but no identification of the essential infimum with a formal free boundary value is required.

On each strict compact subset the centered harmonic prediction is a smooth Gaussian field with uniformly bounded rescaled covariance seminorms. Gaussian supremum estimates and the positive/inverse unit moments give its finite smooth-norm cutoffs with probability tending to one. There is no assertion at every point approaching ∂H_0. For the middle mean score, the coherent-space inclusion into the physical H^1 space is Hilbert–Schmidt on a fixed two-dimensional buffer when its smoothness index exceeds two. Composition with the middle projection and the residual inverse square root preserves this bound. The dual score therefore has bounded covariance trace and operator norm and a small quadratic exponential moment. Its large deterministic cutoff is an actual exterior test. Intersect the finitely many tests after allocating their exceptional probabilities, and apply the fixed strict-buffer Gaussian comparison once for a retained GOOD starting class. □

17 Scalar capture and complete currents

17.1 Serial seeds in the original two-port component

The seeds in this argument are pieces of the original ordinary metric. They are reserved before the annular acceptance rule is evaluated. In particular, one must not add a new detour to an already sampled scalar kernel and identify the resulting kernel with the original one.

Lemma 17.1 (An ordinary serial component with isolated connectors)**.** *The finite source catalogue can be chosen so that, conditional on its seed-free residual, every relevant ordinary query visiting the comparison hole or any selected source support traverses one complete ordered serial component. Its three selected coordinate laws remain independent on a fixed box $I_0^3 \Subset I_1^3$. There are exact anchors and functions $S_j(A_j)$ such that*

$$J = S_1(A_1) + S_2(A_2) + S_3(A_3) + K, \tag{17.1}$$

where K is independent of these coordinates. For an ordinary symmetric filling every original row satisfies simultaneously

$$q_\ell = \min\{B_\ell, A_\ell + J\}. \tag{17.2}$$

Here B_ℓ avoids the comparison hole and all complete seed supports; other unchanged middle portions may be used by that bypass. For a directed estimating filling only the rowwise formula is asserted, with

$$K_\ell^{\mathrm{eff}} = \min\{A_\ell^+ - A_\ell + K^+, A_\ell^- - A_\ell + K^-\}, \qquad A_\ell = \min(A_\ell^+, A_\ell^-), \tag{17.3}$$

and $q_\ell = \min(B_\ell, A_\ell + S + K_\ell^{\mathrm{eff}})$, $S = \sum_j S_j$. If K^0 is the free-hole value and V an unchanged surrounding-loop cap, then $K^0 \le K^\pm, K_\ell^{\mathrm{eff}} \le K^0 + V$. There are retained numerical constants such that, almost everywhere on I_1,

$$cU_v \le S_j \le CU_v, \qquad c\sigma U_v \le S_j' \le C\sigma U_v. \tag{17.4}$$

Every selection and geometric assertion is residual-measurable and holds throughout the enlarged amplitude box and for every compatible positive hidden filling.

Proof. Start from one ordinary protected reference arc and its two delivery anchors, supplied by Lemma B.4. Choose one sleeve $U_* \Subset V_*$ containing the entire selected arc, every source support including its transition region, and all intermediate connectors. Its band has exactly two protected outlets. The permitted network has no side attachment to the sleeve. The hidden core is downstream of this ordinary arm; its alternatives do not meet the sleeve. Every visit to a source support crosses the band inward and outward. If both crossings used the same outlet, the path would repeat a protected ordinary point with positive length between the visits. Loop deletion excludes this. Thus each visit traverses both outlets and the full sleeve. The same argument confines the portion between the inward protected anchors to $\overline{V}_*$. An exit and return would add two further full band crossings.

Use one common floor $-M$ on the complete reference arc and on flat neighborhoods of all source supports. On a fixed ordinary comparison region E the forcing is at least $-M$ and equals $-M$ on the reference segment $P[a,b]$. If a further Weyl change v has supremum norm at most ϵ, a minimizing segment Q with those fixed anchors satisfies

$$e^{-M-\epsilon}\,\mathrm{len}_D\, Q \le \mathrm{len}_{D_{f+v}}\, Q \le \mathrm{len}_{D_{f+v}}\, P[a,b] \le e^{-M+\epsilon} D(a,b;E).$$

For an approximate minimizer the corresponding additive error tends to zero with its original error. Therefore $\mathrm{len}_D\, Q \le e^{2\epsilon} D(a,b;E) + o(1)$. The anchors are fixed, so their endpoint distances converge exactly. Lemma B.1 and uniqueness of the protected reference segment imply avoidance of any prescribed closed set disjoint from that segment, when ϵ is sufficiently small. Otherwise a convergent sequence of near-minimizers meeting that closed set would be a different minimizer or the same reference segment meeting the set.

Choose disjoint chambers $\operatorname{supp} g_j \Subset U_j \Subset V_j \Subset V_*$ in reference order. The exact chamber brackets lie outside $\overline{U}_j$, and all other source functions vanish on V_j. Choose their intervening reference segments at positive distance from every $\overline{U}_k$. Apply the last closed-forbidden-set argument to all chamber and connector anchor pairs: chambers stay in V_j, and connectors stay in $\overline{V}_* \setminus \bigcup_k \overline{U}_k$. There are finitely many pairs, so one positive tolerance works for all. The surrounding coding buffers, on which the perturbations vanish, retain the exact anchors. This controls partial returns to a source support as well as whole chamber visits. It proves amplitude independence of the connector infima.

Now split a minimizing query at these actual anchors. Each chamber piece has cost at least its intrinsic infimum S_j, each connector at least its amplitude-independent intrinsic infimum, and the remaining middle at least K. Conversely concatenate arbitrarily accurate competitors at the same anchors. These inequalities prove the series equality, even if an intrinsic infimum is not attained. The ordinary exterior is symmetric, so both orientations have the same middle. In a directed filling retain the two orientations and take their minimum; extracting A_ℓ gives the displayed effective middle. Its offsets are nonnegative and one is zero, which preserves the common cap interval. Every body retains the genuine hole-avoiding routes used for that cap.

Put a smaller crossing annulus in each plateau $\{g_j = 1\}$. The same confinement forces every relevant chamber competitor to cross it. Its ordinary crossing cost is positive, while a retained candidate gives a finite upper cost. Finite numerical exhaustion bounds both in units of U_v. Weyl comparison and difference quotients then give (17.4). The endpoints of the amplitude box are chosen only after this common confinement tolerance.

All supports, flat neighborhoods and coding buffers belong to the permitted network before the wall is chosen; the floor creates no extra attachment. Choose σ so the finite covariance subtraction leaves at least half the original covariance and the whole box stays within the Weyl tolerance. Select the first successful finite template from the residual, requiring its assertions for every amplitude in $I_1^{J_0}$. Conditionally, the selected coordinates are still the original independent truncated Gaussians. A first template chosen using the realized amplitudes would not have this property and is not used. □

Lemma 17.2 (Seed densities and a lower seed law)**.** *Conditional on the seed-free residual in Lemma 17.1, the old truncated law of $S = \sum_{j=1}^{3} S_j(A_j)$ has a density p satisfying*

$$\|p\|_\infty \le C/U_v, \qquad \|p'\|_\infty \le C/U_v^2, \qquad \|p''\|_\infty \le C/U_v^3 .$$

There are $\delta, c > 0$ and a widened truncated seed law p^-, obtained from coordinates in I_1^3, such that its Radon–Nikodym derivative with respect to the original untruncated reserved Gaussian law is bounded by C and

$$p^-(s-t) \ge c/U_v$$

whenever $0 \le t \le \delta U_v$ and $s \in \operatorname{supp} p + [0,t]$. The laws and constants depend on the residual, but neither law depends on a switching threshold, on the hidden filling, or on a dual test.

Proof. Uniform Weyl comparison permits one to restrict the competitors defining a seed minimum to curves with reference cost at most CU_v. On the enlarged interval every such curve has second

amplitude derivative at most $C\sigma^2 U_v$. Their infimum is semiconcave with this constant. Thus S_j' has bounded variation, with total variation at most $C\sigma U_v$, and stays between the two positive bounds in (17.4). The monotone change-of-variable formula gives the density $\varphi(a)/S_j'(a)$, composed with the inverse cost map. Its supremum and total variation, including the endpoint jumps, are at most C/U_v; the fixed σ is absorbed in the constant. For three independent seeds,

$$p' = Dp_1 * p_2 * p_3, \qquad p'' = Dp_1 * Dp_2 * p_3\,,$$

in the sense of distributions. The asserted bounds follow from the total variation and supremum bounds. In particular p is Lipschitz across its support endpoints.

Each widened seed image extends the old image by at least cU_v at both ends. Its density is at least c/U_v on that widened image. Convolution of the three lower interval densities is bounded below by c/U_v on the old sum image enlarged by a sufficiently small fixed margin. For example, at every point of this compact interior interval, a two-dimensional slice of the product of the three enlarged intervals has area at least cU_v^2. The coordinate law is a fixed Gaussian law restricted to a fixed larger box; its density relative to the untruncated law is bounded by the reciprocal of that box's Gaussian probability. This proves the claim. □

17.2 From scalar capture to the full typed current estimate

We specify the input norm used in the following argument. A finite paired representation has an ordinary labelled occupation frame

$$Q(dF, d\ell, dx) = d^{-1}P(dF)\,\nu_{q_\ell}(dx), \qquad d = E\sum_\ell q_\ell \asymp U_R\,.$$

Its finite estimating current representation, after ordinary occupation restoration, is dominated by AW times this one frame, where $\|W\|_{L^{p_*}(Q)} \le 1$ for some $p_* > 1$. The same domination applies to exterior flags, to a bounded scalar derivative test chosen measurably from the designated past, and to the retained dual norm of the original output label. It is domination of the displayed representation, not just a bound on one annealed scalar. The representation retains one chosen labelled primitive. Its original mark q_P is obtained by unfolding that same primitive; every conversion and complete history identity preserves its original output labels. The primary older input and its source identifier are selected before its histories are expanded and are followed even through a raw branch when other inputs are converted. Unmarked physical equality alone would not provide these labelled identities. The original centered current has mark $q_P - \beta$, and the pure-reference coupling supplies, for every bounded designated-past array,

$$\left|\sum_{C,\ell} E\,Y_{C,\ell}\int w_C(x)(q_P-\beta)\,d\nu_{q_\ell}(x)\right| \le CA\varepsilon_{\mathrm{mix}}\,d\,. \tag{17.5}$$

Here $w_C \ge 0$ is the comparison partition, $\sum_C w_C \le C$, and the centering β belongs to the complete assembled expression.

Lemma 17.3 (Cost-dependent original output labels)**.** *Let $\mathcal{T} = C^1(K)$ with the retained scaled norm, and suppose two finite representations of the same chosen labelled current primitive agree on $f'(D)w$ for every smooth bounded scalar f and every fixed $w \in \mathcal{T}$. Suppose each representation has an integrable positive envelope for its $\mathcal{T}^*$ norm. Then they agree for every bounded Borel map $H : \mathbb{R} \to \mathcal{T}$, with the scalar-cost argument evaluated at the corresponding represented distance. A*

single pointwise dual-norm envelope bounds this extended pairing by $\sup_t \|H(t)\|_{\mathcal{T}}$ *times its mass. The assertion is also valid conditionally on retained exterior data and for a finite list of current components.*

Proof. For a fixed w, push the current to its scalar distance. This is a finite signed measure, by the envelope. Every compactly supported smooth scalar function is the derivative of a bounded smooth primitive. The tensor identities therefore say that the two pushed measures paired with w agree. They consequently agree on every Borel subset of the scalar line. This proves the assertion for finite, and then countably valued, simple maps H: for a countable Borel partition only one value of H is used at any realized scalar cost, so the envelope dominates the entire sum, without the sum of the norms of its values. The space $\mathcal{T}$ is separable. Selecting the first point of a countable dense set within distance $1/n$ gives Borel countably valued approximants to any bounded Borel H, uniformly in $\mathcal{T}$ norm and with the same bound up to $1/n$. Dominated convergence proves the result. The argument uses weak-* scalar pairings only; it assumes no Radon–Nikodym property of $\mathcal{T}^*$. □

The comparison-frame invariant. Every scale-s comparison hole is enlarged by a fixed factor to contain the complete restoration regions of the old carriers meeting its partition support. All these holes lie in the strict ordinary minimizing buffer. Label transport has distance at most the complete radius-r carrier and costs Cr/s against one common current envelope. The short loops used in Lemma 17.5 are outside these enlarged holes and strictly inside the accepted capture core. They avoid the scalar seed supports. Captures and current restoration always keep the original row labels and original component masses; there is one final whole-frame normalization.

The three-seed bounds of Lemma 17.2 are interpreted in $W^{2,\infty}$, with a $C^{1,1}$ representative. A classical second derivative is unnecessary. The lower widened density holds on the entire old sum support plus its fixed translation margin, including both endpoint regions. Thus the scalar-translation calculation for arbitrary full covectors in Proposition 17.4 uses a genuinely global lower envelope.

After all old variables and histories have been integrated, convolve the actual typed vector primitive with the fixed old arm law. The resulting current density and its divergence are supported on the full Minkowski support in Lemma 19.2. Divide both by the same fixed real fresh density on that support. Its positive lower bound and the forest incidence right inverse give the paired value/current identities on every fresh fiber. Integrating the divergence gives zero total value mass on each fiber. The companion, denominator and physical basis are fixed independently of the incoming mean. A countable dense subspace of the retained separable C^1 test space supplies the weak-* densities and identities; a vector Radon–Nikodym theorem is not required.

Proposition 17.4 (Full covectors and one ordinary frame)**.** *Suppose a residual-measurable scalar capture as above is available at each retained comparison cell, outside a failure event exterior to the complete old payload. Let its seed scale be* $v_C \geq u$*, its hidden comparison scale be* s*, and its free-fill kernel and unchanged loop be* K_C^0, V_C*. Suppose*

$$K_C^0 \leq K_C \leq K_C^0 + V_C$$

for all fillings. Assume the rooted ratio and failure estimates

$$E_Q \sum_C w_C(x) \min\{1, V_C/U_{v_C}\}^{\theta p'_*} \leq C(s/u)^{\kappa}, \qquad Q(\textit{capture failure}) \leq C\varepsilon_{\text{fail}}\,,$$

with $0 < \theta < 1$ and $p'_* = p_*/(p_* - 1)$. Then, after integrating all old variables, the centered activity is a finite sum of typed current primitives ρ_τ satisfying

$$\sum_\tau \|\rho_\tau\|_{\mathrm{TV}} \le C_M A \left(\varepsilon_{\mathrm{mix}} + (s/u)^{\kappa/p'_*} + \varepsilon_{\mathrm{fail}}^{1/p'_*} + s/R + r/s\right) U_R \,. \tag{17.6}$$

The total variation is tested by arbitrary bounded covectors of all retained port costs, together with the retained dual output labels. It is not merely a supremum over gradients of scalar tests.

Proof. Fix an arbitrary bounded covector $H = (H_\ell)$ of the full port matrix, with values in the retained C^1 output-label space and norm at most one. At this reset use the fixed smooth coarse carrier containing the comparison centers; its tests restrict to the old carriers. If the source anchor is retained jointly, the carrier lies in $\mathbb{R}^2 \times \mathbb{R}^2$ and its scaled C^1 norm controls derivatives in both the original anchor and output variables. Freezing both at the comparison center therefore has the same Cs/R bound; arbitrary fine-scale oscillation in a discrete anchor is not part of this test class. If the owner sector fixes the anchor, only the output variable is needed. On a scalar capture every port output is the known function (17.2) of the one variable $J = S + K$. Hence the pairing with its ℓth current component is a C^1-valued scalar test $G_\ell(J)1_{\{J<T_\ell\}}$, where $T_\ell = B_\ell - A_\ell$ and $\sup_J \|G_\ell(J)\|_{C^1} \le 1$. Lemma 17.3 allows the existing cheap representation to be used for these joint tests. For the original ordinary matrix, on activation $q_\ell = A_\ell + J$, so the test is the Borel function $G_\ell(q_\ell - A_\ell)1_{\{q_\ell<B_\ell\}}$ of that scalar kernel; it is zero on its bypass branch. Replacing this one scalar test by an admitted converted body does not require its full port matrix to have the same common J. Its rowwise effective middle in (17.3) is sufficient below. Freeze the output-label argument at the exterior-measurable center of the comparison cell. The same positive envelope bounds its error by CAs/R in ordinary-frame units. Moving a comparison weight from the original point to the root within its complete radius-r carrier costs CAr/s. The original label remains in the resulting coefficient; this freezing is only an estimate in its dual norm. After freezing, write the scalar test as $g_\ell(J)1_{\{J<T_\ell\}}$, with $\|g_\ell\|_\infty \le 1$. Coefficients of any fixed finite forest valuation can be absorbed into g_ℓ at a cost depending only on the port cutoff. Conditional on the core and residual, the seed has a continuous density, so switching ties have probability zero. Bounded approximation gives this statement for all measurable covectors, not only continuous ones.

Suppress the cell and port labels. Set

$$F_g(k) = \int g(t+k)1_{\{t+k<T\}}p(t)\,dt, \quad E_* = \{S^- < T - K^0 - V\}, \quad q_* = P^-(E_*)\,.$$

On $V \le \delta U_v$, the lower seed law gives

$$|F_g(K^0)| \le Cq_*, \qquad |F_g(K) - F_g(K^0)| \le C(V/U_v)q_*\,. \tag{17.7}$$

Indeed, after writing $d_0 = K - K^0 \in [0, V]$, change variables to obtain

$$F_g(K) - F_g(K^0) = \int_{t<T-K^0} g(t+K^0)[p(t-d_0) - p(t)]\,dt\,.$$

The global Lipschitz bound and support margin imply $|p(t-d_0) - p(t)| \le C(V/U_v)p^-(t-V)$. The same lower envelope bounds $p(t) \le Cp^-(t-V)$ on its support. Integration proves (17.7). There is no assumption that the activation probability is bounded below.

Define

$$Y = \mathbf{1}_{\{V\le\delta U_v\}}\frac{F_g(K^0)}{q_*}, \qquad Y = 0 \quad \text{when } q_* = 0.$$

It is a bounded designated-past test. The following frozen and remainder calculations are on $\{V \le \delta U_v\}$; the complementary large-ratio branch is estimated separately below. On E_* the lower-seed parent route is active for every hidden filling, so its capture occupation is exactly the canonical middle occupation. Restore the lower coordinate law to the ordinary physical field before summing the cells. If R_C^- denotes its bounded density relative to the original Gaussian coordinates, the frozen expression is

$$\sum_{C,\ell} E \int w_C(x) R_C^- 1_{E_{*,C,\ell}} Y_{C,\ell}(q_P - \beta)\, d\nu_{q_\ell}(x)\,.$$

All factors multiplying the centered mark form one bounded designated-past array, so (17.5) applies.

Different cells can use different augmented seed menus. This does not change the preceding conclusion. Sample the ordinary field first and sample each menu from its Gaussian conditional law given that field. Because its score directions are supported outside the hidden payload, this conditional kernel is measurable from the designated past. Averaging its bounded coefficient preserves both its bound and past measurability. The resulting expression is still on one ordinary physical field and one labelled occupation frame. No product of the cellwise density bounds is taken.

For the remainder, divide $F_g(K) - F_g(K^0)$ by q_* and realize q_* by the same event E_*. The quotient is a scalar function of K, with norm at most CV/U_v. Conditional on the lower seed, K equals $q_\ell - A_\ell - S^-$ on E_* for the original ordinary row. For a directed converted body the same event activates its row for every filling, and the corresponding quantity is K_ℓ^{eff}, still in $[K^0, K^0 + V]$. The remainder quotient is therefore an admitted scalar derivative test with the same uniform bound on every such body. No common middle across its directed rows is used. The input estimating-current bound applies with the exterior flag $R_C^- 1_{E_*}$; no pointwise bound on the original Wick mark is required. Summing w_C before estimating yields one expectation of $W \min\{1, V_C/U_{v_C}\}$ under Q. Since $\min\{1, x\} \le x^\theta$, Hölder's inequality and the stated rooted ratio bound give $CA(s/u)^{\kappa/p'_*} d$. The branch $V_C > \delta U_{v_C}$ is estimated in the same way, using its exterior indicator and $1_{\{x>\delta\}} \le \delta^{-\theta p'_*} x^{\theta p'_*}$.

On capture failure retain the old estimating representation. Its pointwise current envelope and ordinary occupation restoration give $CA\varepsilon_{\mathrm{fail}}^{1/p'_*} d$ by Hölder's inequality. This step does not identify its current with a raw vector measure on the original port matrix. It produces separate typed estimating primitives. The new outer atlas includes the entire old modification region; its untouched annulus provides finite positive avoiding connections between all new inner ports. Thus each such term has a complete finite directed shortest-path matrix to which the same finite forest valuation applies. Directed triangle inequalities suffice for the source–sink active-graph argument and for the arm-independent face lattice; symmetry is not needed.

Finally push each successful full-vector current and each failed estimating current through its finite forest row map. Multiplication by the bounded valuation coefficients and pushforward increase the sum of total variations by at most a constant depending only on the retained port count. As every bound above is uniform over the arbitrary covector H, this gives (17.6). All old comparison residuals, seed coordinates, labels, and histories are integrated before this pushforward is declared complete. None is subsequently made part of the fresh companion's conditioning sigma-field. □

Scales. A fixed finite seeded forcing catalogue has a uniform positive success probability at each eligible GOOD annulus. Iterating the predictable trials between radii u and R gives a failure

probability $C(u/R)^\eta$. If the accepted scale is $v \in [u, R]$, the rooted length-ratio estimate is summed over the geometric list of possible v; its sum is bounded by $C(s/u)^\kappa$. This geometric sum does not cost the number of comparison cells. For example, $s = r^{1/2}R^{1/2}$ and $u = r^{1/4}R^{3/4}$ turn these errors into positive powers of r/R. An incoming-mean innovation is measured on the fixed smaller comparison hole; it must not be estimated by v/R when a trial near R can be accepted. The latter ratio need not be small.

Lemma 17.5 (A truncated loop ratio at occupied comparison holes)**.** *Use a GOOD ordinary row-labelled frame at scale v, with whole-frame normalizer $d \geq cU_v$ and one fixed moment $p > 1$ of its total length L/U_v. All roots in question lie in its strict freely minimizing buffer. Let w_C be a nonnegative scale-s partition with bounded overlap, $s \ll v$. Enlarge each comparison hole once so that it contains the complete old restoration region for every carrier meeting* $\operatorname{supp} w_C$.

There are $\chi \in (0,1)$ and, for every fixed $t > 0$, constants $C, \kappa > 0$ such that one can choose an unchanged surrounding-loop cap V_C in the annulus from radius Cs to $v(Cs/v)^\chi$, and

$$E_Q \sum_C w_C(x) \min\{1, V_C/U_v\}^t \leq C(s/v)^\kappa.$$

The same bound holds for the accepted canonical suboccupation, with its original component masses retained. Consequently, summing accepted scales $v \geq u$ gives $C(s/u)^\kappa$. No assertion is made at unoccupied holes.

Proof. First use one ordinary physical field and the original row label. A nonzero restored current forces an ordinary visit to its modification region by Lemma 20.4. The entire protecting region lies in the enlarged comparison hole. Original and restored labels move by at most the complete carrier diameter Cr, so moving w_C costs Cr/s against the common current envelope and involves only a bounded number of neighboring cells. This deterministic enlargement is fixed before the numerical source cutoffs.

Trim a freely minimizing ordinary episode through the hole at its first and last contacts with the unchanged scale-v loop. The loop arc is an allowed competitor in the same ordinary buffer, so this episode has length at most U_v (or its fixed retained multiple). Its endpoints have fixed positive clearance from the smaller annulus. Apply [4, Lemma 2.13] with $f = 0$ to this internal geodesic. The lemma is simultaneous over centers, endpoints and geodesics; hence it applies to the occupied, field-dependent comparison holes. On an event $G_{s/v}$ of ordinary probability at least $1 - C(s/v)^b$ it yields

$$V_C \leq C(s/v)^\theta U_v$$

for every such occupied hole. A near-optimal surrounding loop is enough. If the construction uses a nonnegative lowering supported deeper in the hole, the pre-entry version of the same lemma applies instead. The reference loop and the short loop are measured in the unchanged field.

The statement of that primary result is for the pure metric. On the retained strict buffer the ordinary conditional law has a fixed higher Gaussian-density moment relative to that law. A single Hölder comparison therefore preserves a positive exponent b. It is taken before any future last-success restriction. Under the ordinary Campbell frame,

$$Q(G^c_{s/v}) = \frac{E[L\mathbf{1}_{G^c_{s/v}}]}{d} \leq C\|L/U_v\|_p \, P(G^c_{s/v})^{1-1/p}.$$

The sum of the weights is bounded. On $G_{s/v}$ use the short-loop bound; on its complement use the truncation by one. This proves the conclusion with any $\kappa \leq \min(t\theta, b(1-1/p))$.

For an accepted component, condition on its matched annular data. The common forecast cut and the ordinary frame moment give the same normalized Campbell estimate. Lemma 10.1 requires that its original unnormalized mass be retained. Summing these subprobability components therefore introduces no factor equal to the number of cells. For each possible accepted v, discard further selection indicators only inside this original positive comparison. Sum the geometric series $\sum_{v\ge u}(s/v)^{\kappa}$.

Finally choose the minimum accepted age so that $v(Cs/v)^{\chi}$ stays strictly inside the capture core and away from every seed support. The comparison-hole enlargement and this choice change neither the hidden alternatives nor the fixed observation block. □

18 Finite port bodies

The polyhedral argument must retain one representation while the independent arm lengths vary. Attaching an arm separately to every edge of an arbitrary directed forest would count an intermediate arm twice along a directed path of length two. We prove that completing the inner port distances and attaching the positive arms first excludes those paths from every active constraint graph.

A body with at most one port cannot affect an outside-to-outside distance: every visit returns to the same port and can be removed at nonnegative cost. Its internal first variation is zero. We therefore fix $n \ge 2$. Let $k = (k_{ij})$ be a finite directed distance matrix on the ports $\{1, \dots, n\}$, and let $s_i > 0$. Thus $k_{ii} = 0$, $k_{ij} \ge 0$, and $k_{ij} \le k_{i\ell} + k_{\ell j}$. In $V = \mathbb{R}^n/\mathbb{R}\mathbf{1}$, define

$$P(k,s) = \{y \in V : y_i - y_j \le k_{ij} + s_i + s_j \text{ for } i \ne j\}\,. \tag{18.1}$$

This is a bounded full-dimensional polytope. An active directed edge $i \to j$ at a face means that its displayed inequality is an equality throughout the face. A directed graph is called source–sink oriented if every vertex has only outgoing edges or only incoming edges, with isolated vertices permitted.

Proposition 18.1 (The positive arms determine fixed forest types)**.** *For fixed k, the face lattice of $P(k,s)$, labelled by its active directed inequalities, is independent of $s \in (0,\infty)^n$. Every active graph is source–sink oriented. There is a finite collection of oriented forests τ and integer functions $c_\tau(k)$ such that*

$$\mathbf{1}_{P(k,s)} = \sum_{\tau} c_\tau(k)\mathbf{1}_{P_\tau(k,s)}, \qquad P_\tau(k,s) = \{y : y_i - y_j \le k_{ij} + s_i + s_j,\ (i,j) \in \tau\}\,. \tag{18.2}$$

Every forest in this identity is source–sink oriented. The number of terms and $|c_\tau(k)|$ are bounded by a constant depending only on n. The coefficients can be chosen constant on each member of a finite polyhedral stratification of the directed triangle cone, and

$$\sum_{\tau} c_\tau(k) = 1\,. \tag{18.3}$$

The empty forest is allowed, with $P_\varnothing = V$.

Proof. Step 1: Active edges. If $i \to j$ and $j \to \ell$ were active, with $i \ne \ell$, then

$$y_i - y_\ell = k_{ij} + k_{j\ell} + s_i + 2s_j + s_\ell > k_{i\ell} + s_i + s_\ell\,, \tag{18.4}$$

contradicting feasibility. Opposite active edges are excluded by positive costs. Hence no vertex has both an incoming and an outgoing active edge.

Step 2: Vertices and faces. At a vertex the active incidence vectors span V, so their underlying graph is connected and contains a spanning tree T. Write its source–sink bipartition as $U \sqcup W$. Solve, modulo constants,

$$z_u - z_v = k_{uv} \quad (u \to v \in T), \qquad y_i = z_i + \epsilon_i s_i, \qquad \epsilon_i = \begin{cases} 1 & i \in U, \\ -1 & i \in W. \end{cases} \tag{18.5}$$

This solves all active tree equalities. It is feasible exactly when

$$z_u - z_v \le k_{uv} \qquad (u \in U,\ v \in W)\,. \tag{18.6}$$

Necessity follows by substituting (18.5). To prove sufficiency, choose an active neighbor v of $u' \in U$. Then

$$z_u - z_{u'} \le k_{uv} - k_{u'v} \le k_{uu'}\,. \tag{18.7}$$

This proves all same-source inequalities, with a strict margin after the arms are added. The same argument with an active neighbor of a sink proves $z_v - z_{v'} \le k_{vv'}$. Finally, choose an active source u' neighboring v. The same-source bound gives

$$z_v - z_u = z_{u'} - z_u - k_{u'v} \le k_{u'u} - k_{u'v} \le k_{vu}\,. \tag{18.8}$$

Thus reverse cross inequalities also hold strictly after arm addition. Only forward cross inequalities can be active, and their slacks are $k_{uv} - z_u + z_v$, independent of s.

All tree-defined vertices and their complete active sets are consequently independent of s as labelled data. Different active bipartitions cannot give the same vertex: their connected union would then have a vertex with both incoming and outgoing active edges. Vertex active sets determine all faces and incidences. Since tree solutions are linear in k, their feasibility and equality conditions also give a finite polyhedral stratification on which the labelled face data are fixed.

Step 3: Forest cones. Apply the Brianchon–Gram identity [9, Theorem 3.7.1]:

$$\mathbf{1}_P = \sum_{\varnothing \ne F \text{ face of } P} (-1)^{\dim F} \mathbf{1}_{T_F P}\,. \tag{18.9}$$

If y_F is in the relative interior of F, then $T_F P = y_F + C_F^\circ$, where C_F is the cone generated by its active incidence vectors and the polar uses the inequality $\langle v, x\rangle \le 0$. Triangulate C_F using its existing extreme rays, with a fixed ordering of the possible incidence vectors. The choice depends only on the active set.

For completeness, polarity gives the required signed indicator identity as follows. Let $C_1, \dots, C_m$ be the maximal cones in the triangulation. Inclusion–exclusion applies to their intersections. For a fixed v, intersect each of these cones with $\{x : \langle v, x\rangle \ge 1\}$ and with a sufficiently large common closed ball. Choose the ball to meet every nonempty intersection in this finite family. Each resulting nonempty compact convex set has Euler characteristic one. The relation

$$\mathbf{1}_{\{v \in C^\circ\}} = 1 - \mathbf{1}_{\{C \cap \{\langle v, x\rangle \ge 1\} \ne \varnothing\}} \tag{18.10}$$

and inclusion–exclusion therefore express $\mathbf{1}_{C_F^\circ}$ as a finite signed sum of polars of simplicial fan faces. The constant coefficients sum to one because every cone intersection contains the origin.

A set of incidence vectors $e_i - e_j$ is linearly independent if and only if its underlying undirected graph is a forest. One direction follows by summing oriented vectors around a cycle; the other

follows by successively removing leaves. Each translated simplicial polar is therefore exactly a set $P_\tau(k,s)$ in (18.2). Its rows are active at F, so the translation has the required right-hand sides. There are only finitely many possible incidence subsets and triangulations for fixed n, which gives the stated uniform finite bound. The choices are constant on the preceding strata. Finally, 0 belongs to every set in (18.2); evaluating there proves (18.3). □

This statement requires the complete finite inner distance matrix. Disconnected inner bodies require a componentwise treatment or an explicitly retained finite intrinsic completion. Arbitrary additional box constraints are not part of $P(k,s)$. Only routes avoiding the complete modified body must remain among the fixed exterior constraints; a route avoiding a smaller inner hole need not remain in every forest term. A collection of unrelated edge costs need not satisfy the triangle inequalities. For example, along a directed chain $1 \to 2 \to 3$, attaching arms to its two rows gives a cost $k_{12} + k_{23} + s_1 + 2s_2 + s_3$, whereas attaching arms only at the external ports gives $k_{12} + k_{23} + s_1 + s_3$. Proposition 18.1 excludes such a directed chain from the after-attachment active graphs.

18.1 Exterior tests and first variations

An estimating forest in (18.2) is directed. Its edges may be actual physical paths with restricted allowed traversal; it is not asserted to define a symmetric metric. Every genuine route avoiding the complete modified body remains in every term. An avoiding route inside that body may instead be used only as an upper bound for an old kernel; it is not automatically a parallel edge of every forest term.

Proposition 18.2 (Compatibility with exterior gluing)**.** *Fix a finite positive exterior network with the n ports and two terminals. Assume its fixed directed paths connect every vertex to every other vertex. Let $D(k,s)$ be the terminal distance after adjoining all body edges of costs $k_{ij} + s_i + s_j$, and let $D_\tau(k,s)$ be the distance after adjoining just the oriented forest edges. For every absolutely continuous f on an interval containing all these distances,*

$$f(D(k,s)) = \sum_\tau c_\tau(k) f(D_\tau(k,s)) \,. \tag{18.11}$$

The same coefficients work for every exterior network satisfying the stated condition.

Proof. Represent directed distance by potentials: an edge $i \to j$ of cost b_{ij} imposes $y_i - y_j \le b_{ij}$, and the distance from the first terminal to the second is the supremum of their potential difference. To check this formula, sum the inequalities along any directed path for the upper bound, and take y_i to be the distance from vertex i to the second terminal for equality.

Pin one potential. Strong connectivity of the fixed exterior makes its potential set compact, since fixed exterior paths bound every potential difference in both directions. Lift (18.2) to all vertices, and intersect it with these fixed constraints and the half-space where the terminal difference is at least t. All intersections are compact convex sets, including possibly the empty set. Euler characteristic and its additivity give

$$\mathbf{1}_{\{D(k,s)\ge t\}} = \sum_\tau c_\tau(k) \mathbf{1}_{\{D_\tau(k,s)\ge t\}} \,. \tag{18.12}$$

Every distance is nonnegative and bounded above by a fixed exterior route. Extend f absolutely continuously to include zero if necessary. Integrate (18.12) against $f'(t)$ between zero and that bound, then use (18.3) for the constant $f(0)$. □

The connectivity condition in this proposition must be furnished by genuine retained exterior routes. Adding an artificial cap changes the observation. Compactness of the potential sets here is proved before Euler characteristic is applied.

For first variations, let $k = k(x)$ be a locally Sobolev family of pseudometric matrices in a finite-dimensional physical parameter x, and fix the arms and exterior context. Distances in these finite positive networks are minima over finitely many simple-path costs, each affine in k. Thus $F(x) = f(D(k(x), s))$ and $F_S(x) = \sum_\tau c_\tau^S f(D_\tau(k(x), s))$ are locally Sobolev for $f \in C^1$, where c^S is the coefficient list for one stratum S. The common exterior route bounds all distances, so f' is bounded on the interval in use. On $E_S = \{x : k(x) \in S\}$, identity (18.11) gives $F = F_S$. Locality of weak derivatives gives $\nabla F = \nabla F_S$ almost everywhere on E_S. There are finitely many strata, hence

$$\nabla f(D(k(x), s)) = \sum_\tau c_\tau(k(x)) \nabla f(D_\tau(k(x), s)) \quad \text{for almost every } x\,. \tag{18.13}$$

The same locality property handles ties of minimizing routes. These assertions concern each fixed exterior test. They do not replace the probability and integrability estimates needed to pass to a continuum Gaussian limit.

18.2 The complete vector current

Let $J = (J_e)_e$ be the finite vector measure on the complete inner port matrix induced by a physical first variation, with the preceding tangency convention on strata. The relevant norm is

$$\|J\|_{\mathrm{TV}} = \sup_{\max_e \|H_e\|_\infty \le 1} \left| \sum_e \int H_e(k)\, J_e(\mathrm{d}k) \right| . \tag{18.14}$$

For each forest, let π_τ retain its edge coordinates, and push forward $c_\tau(k)$ times the corresponding components of J. The resulting vector measures ρ_τ satisfy

$$\sum_\tau \|\rho_\tau\|_{\mathrm{TV}} \le C_n \|J\|_{\mathrm{TV}}\,. \tag{18.15}$$

Indeed, restriction of components and pushforward cannot increase the sum of their total variations, and the number and coefficients of terms are bounded by n. This is the measure estimate required before convolution by the independent outer arms.

The scalar capture formula can be used to prove (18.14) only with simultaneous control of the full vector. Suppose, on a comparison event, every entry has the form

$$k_e = \min\{B_e, A_e + J_0\}\,, \tag{18.16}$$

where the arrays A_e, B_e are measurable outside the entire source region and J_0 is one common captured scalar. Then an arbitrary bounded covector $H_e(k)$ becomes the bounded scalar function $g_e(J_0) = H_e((\min\{B_f, A_f + J_0\})_f)$. Its current also has the activation factor $\mathbf{1}_{\{J_0 < B_e - A_e\}}$. Thus a scalar estimate uniform over all such g_e controls the vector pairing.

If comparison regions are summed, their error must have the form ϵd_C with $\sum_C d_C \le CU$ for the actual occupation frame. A separate bound ϵU for each region would lose the number of regions. The scalar-capture construction and its current estimate in Proposition 17.4 supply these bounds before this finite-dimensional operation is applied.

18.3 Convolution and division on each fiber

For a forest with m edges, let $A_\tau : \mathbb{R}^n \to \mathbb{R}^m$ have a one at each endpoint of each edge. It maps arm lengths to the additions $s_i + s_j$. Root each component and set the root coordinate to zero. For prescribed row increments t, define recursively

$$(R_\tau t)_{\text{child}} = t_{\text{edge}} - (R_\tau t)_{\text{parent}} . \tag{18.17}$$

Then $A_\tau R_\tau = I$ and $\|R_\tau t\|_\infty \le (n-1)\|t\|_\infty$. The Euclidean coarea factor is

$$\mathcal{J}_\tau = \sqrt{\det(A_\tau A_\tau^{\mathsf{T}})} . \tag{18.18}$$

It is positive and bounded above and below at fixed n. For a three-vertex path,

$$A_\tau = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}, \qquad A_\tau A_\tau^{\mathsf{T}} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \qquad \mathcal{J}_\tau = \sqrt{3} . \tag{18.19}$$

Thus no Euclidean coisometry is available merely by changing orthonormal coordinates.

Lemma 18.3 (A fixed density preserves the value and current)**.** *Let $U > 0$. Suppose g is a compactly supported C^1 probability density on $\mathbb{R}^m$ satisfying*

$$\|g\|_\infty \le CU^{-m}, \qquad \|\nabla g\|_\infty \le CU^{-m-1} . \tag{18.20}$$

Let ρ be a compactly supported vector measure with $\|\rho\|_{\mathrm{TV}} \le \epsilon U$, and put

$$\sigma = \rho * g, \qquad \mu = -\operatorname{div} \sigma . \tag{18.21}$$

For each retained filling ω, suppose a real probability density $g_{+,\omega}$ satisfies

$$g_{+,\omega}(y) \ge cU^{-m} \quad \text{on } \operatorname{supp}\sigma \cup \operatorname{supp}\mu . \tag{18.22}$$

Suppose also that a differentiable physical map $Y_\omega : Z_\omega \to \mathbb{R}^m$, on a finite-dimensional Euclidean coordinate space with probability law π_ω, pushes π_ω forward to $g_{+,\omega}(y)\,\mathrm{d}y$. Assume there is a measurable family of linear right inverses $R_\omega(z)$ with

$$DY_\omega(z) R_\omega(z) = I, \qquad \|R_\omega(z)\| \le C/U \tag{18.23}$$

for π_ω-almost every z whose image belongs to $\operatorname{supp}\sigma$. Define $h_\omega = \mu/g_{+,\omega}$ and $a_\omega = \sigma/g_{+,\omega}$ on their respective supports and zero elsewhere. Then

$$\|h_\omega\|_\infty \le C'\epsilon, \qquad \|a_\omega\|_\infty \le C'\epsilon U , \tag{18.24}$$

and the physical direction $R_\omega(z)a_\omega(Y_\omega(z))$, taken to be zero where $a_\omega = 0$, has norm at most $C'\epsilon$. Here C' depends only on C, c, m. For every Lipschitz ϕ,

$$\int h_\omega g_{+,\omega} = 0, \qquad \int \phi h_\omega g_{+,\omega} = \int \nabla\phi \cdot a_\omega g_{+,\omega} . \tag{18.25}$$

Proof. Convolution gives $\|\sigma\|_\infty \le C\epsilon U^{1-m}$ and $\|\mu\|_\infty \le C_m \epsilon U^{-m}$. Divide by (18.22) and apply the assumed right inverse. Since σ is compactly supported, integration by parts gives $\int \mu = 0$ and $\int \phi\mu = \int \nabla\phi\cdot\sigma$. Substituting the definitions proves the identities on each individual filling. □

The quotient density of a smooth independent arm density is given by

$$g(y) = \mathcal{J}_\tau^{-1} \int_{A_\tau s = y} p(s)\, \mathrm{d}\mathcal{H}^{n-m}(s)\,. \tag{18.26}$$

If $p \in C_c^1(\mathbb{R}^n)$ satisfies $\|p\|_\infty \le CU^{-n}$, $\|\nabla p\|_\infty \le CU^{-n-1}$ and $\operatorname{diam}(\operatorname{supp} p) \le CU$, parameterize the fibers by $R_\tau y + \ker A_\tau$. Their intersection volume is $O(U^{n-m})$, which proves the bounds on g in Lemma 18.3. This retains the Jacobian and the chosen right inverse.

Lemma 19.2 supplies (18.22) for the fresh physical realization, separately on every retained filling. An averaged lower density would not suffice. For example, two densities with disjoint supports can have a positive average on their union while each vanishes on half that union. The fresh density and physical right inverse are fixed throughout the incoming holomorphic mean family. With those quantities fixed, convolution, division, and lifting are bounded linear operations on the integrated numerator.

A bypass contact. A fixed exterior bypass gives $D = \min(y, B)$. For a compactly supported C^1 current density r, its value measure is

$$\mu = -r'(y)\mathbf{1}_{\{y<B\}}\, dy + r(B)\delta_B\,.$$

Integration by parts gives $\int f(D)\, d\mu = \int_{y<B} r(y) f'(y)\, dy$, and $\mu(\mathbb{R}) = 0$. Thus the contact at the active cap is part of the value identity. Omitting it would already fail in this one-edge example.

19 A fresh physical realization

19.1 Fresh rows near their conditional essential infima

Let D be the new reset hole, containing the complete old restricted body, and let $U > 0$ be the retained exterior unit. For its finitely many ordinary local rows use m_i from Lemma 9.2 and the genuine avoiding kernels B_i.

We check the law under which an old restricted row is evaluated. Given the original field outside a finite family of disjoint maximal child domains, the ordinary conditional Dirichlet kernels in those domains form a product. A child's fresh companion is obtained from its own finite augmented conditional Gaussian kernel by restriction to its retained residual event and coordinate windows, with a finite density bound. Pushing forward by reconstruction preserves that bound. Replacing finitely many disjoint child kernels therefore gives a joint law absolutely continuous with respect to the original product kernel: its density is the product of the child densities. Integrating the common exterior preserves absolute continuity. Repeating this argument along the finite genealogy proves the same assertion for the one reconstructed physical field in which the whole old estimating body is drawn. The shared coherent mean is kept in its original coefficient likelihood; it is not placed on the newly introduced companion coordinates.

Thus, conditional on the retained exterior e, the ordinary lower-support null sets of Lemma 9.2 are still null for the actual old body. Every allowed estimating traversal is a physical curve in this same reconstructed field, so, almost surely,

$$m_i(e) \le K_i^{\mathrm{ord}}(e, F) \le K_i^{\mathrm{old}}(e, F). \tag{19.1}$$

The new reset hole contains every old modification region. A route avoiding this hole consequently avoids each old modification region, and is retained in every old term. Hence $K_i^{\mathrm{old}} \le B_i(e)$. The

retained exterior cap gives the remaining upper bound in (19.2). Each fixed real Cameron–Martin translation inside the reset hole is equivalent to the conditional ordinary law, so these support bounds also hold after the complete-numerator translation. Only this almost-sure assertion is used to contain the support of the integrated numerator.

Physical lifting and the retained cap give

$$m_i \le K_i^{\mathrm{old}} \le \min\{B_i, m_i + C_0 U\}. \tag{19.2}$$

The fixed exterior Gaussian class is exhausted by the local events of Lemma 9.5.

Proposition 19.1 (Fresh physical rows with positive private increments)**.** *Fix an accuracy $\varepsilon > 0$, a finite row cutoff, and a compact positive interval $[t_-, t_+]$. After a finite local exterior and residual exhaustion, there are disjoint private supports, nonnegative smooth g_i, and an augmentation*

$$X = X^{\mathrm{res}} + \sigma \sum_{i \in I} A_i g_i, \qquad A_i \sim N(0,1), \quad A \perp X^{\mathrm{res}}, \qquad \operatorname{Cov} X^{\mathrm{res}} \ge \tfrac{1}{2} \operatorname{Cov} X,$$

where $I = \{i : m_i < B_i\}$, with the following property. A residual event Γ has conditional probability at least $p_0 > 0$ uniformly on the retained exterior class. On each retained residual fiber there are actual curves, chosen before A, whose costs satisfy

$$L_i(A_i) = \kappa_i + T_i(A_i), \qquad m_i \le \kappa_i \le m_i + \varepsilon U, \qquad T_i(0) = 0. \tag{19.3}$$

Curve i avoids every other private support. On a fixed interval $[0, A_*]$ the increments cover $[t_- U, t_+ U]$, their independent pushforward densities are at least c_0/U, and on the inverse image of this interval

$$c_1 U \le T_i' \le C_1 U. \tag{19.4}$$

If $m_i = B_i$, the old row equals this constant and is kept exactly. All constants may depend on the retained finite cutoffs.

Proof. Fix disjoint private balls strictly inside D, with smaller plateau crossing regions and surrounding finite circuits. Reserve their bumps with σ sufficiently small. By Lemma 9.2, after a sufficiently large common compact lowering each variable row has an actual competitor within the prescribed accuracy of m_i and strictly below B_i. Replace all its visits to private supports by surrounding circuit arcs. A further compact lowering makes these finitely many arcs as cheap as required. The detoured row still has cost below B_i and therefore has a nonconstant portion inside open D.

Attach the private excursion by an actual intersection. A rational annulus crossed by that interior portion has a surrounding circuit which meets it. A finite chain of overlapping rational corridors, avoiding the other private balls, connects that circuit to the selected plateau; transverse crossings meet at every join. Add the plateau crossing and the reversed connector. These are finitely many compact physical paths. Increasing the compact lowering makes their added length arbitrarily small, without increasing any retained piece. The final curves satisfy (19.3), avoid the other supports, have private mass at most εU, and have strictly positive mass on their own plateaux. They are allowed estimating traversals; no ordinary-minimality claim is made for the excursions. Intersections of different labelled traversals do not add switches to their allowed path class.

All shifts used so far belong to one fixed countable smooth catalogue. Their Gaussian translations are equivalent. Thus the event that suitable curves exist has positive conditional residual

probability for almost every exterior: otherwise all its countably many fixed translates would be null, contradicting the construction. Use the local rows of Lemma 9.3. Conditional probabilities of the finite catalogue events have versions measurable in the fixed observation block. Increasing finite subcatalogues and decreasing positive cutoffs therefore give exterior classes exhausting full probability on which the retained event has probability at least a fixed $p_0 > 0$. This is exactly the local numerical exhaustion already included among the raw events. It is not a uniform assertion over arbitrary exteriors. Choose the final curve and profile first; then cut its positive plateau mass below by $q_*U > 0$. No earlier approximation index is changed.

For the residual length measure ν_i of the chosen curve,

$$T_i(A) = \int_{\gamma_i} (e^{\sigma A g_i} - 1)\, d\nu_i. \tag{19.5}$$

Since $0 \le g_i \le 1$, the private mass is at most εU, and the plateau mass is at least q_*U, for $A \ge 0$,

$$T_i(A) \ge q_*U(e^{\sigma A} - 1), \qquad T_i'(A) \ge \sigma q_* U e^{\sigma A}, \qquad T_i'(A) \le \sigma(\varepsilon U + T_i(A)). \tag{19.6}$$

Choose finite A_* covering t_+U. The Gaussian density has a positive minimum on this interval, and the change of variables gives both density and derivative bounds. The coordinates are independent conditionally on the residual. Their original window probabilities remain in the formula. □

Lemma 19.2 (Density on the whole old convolution support)**.** *Let A be the unsigned incidence matrix of a nonempty forest with m rows and n ports, and let R satisfy $AR = I$ and $\|R\|_{\infty\to\infty} \le n - 1$. The old arm support is a box $\mathcal{B}_0$; the fresh arm law has density at least cU^{-n} on a box $\mathcal{B}_1$ containing its bU neighborhood. Suppose (19.2) holds. Choose $0 < \delta < b/(4(n-1))$, then $\varepsilon < \delta/4$, and fresh increments on a window strictly containing $[\delta/2, C_0 + \delta]U$. On each retained fresh residual fiber, the law of $\kappa + T + AS$ has density at least $c'U^{-m}$ on the complete old support*

$$\{k + As : k_i \in [m_i, \min(B_i, m_i + C_0U)],\ s \in \mathcal{B}_0\}.$$

A row with $m_i = B_i$ is fixed at m_i.

Proof. For an old point $y = k + As$, put $t_i = \delta U$ on variable rows and zero on constant rows, and $s' = s - Rt$. Then s' is in the interior of the widened arm box and

$$y - \kappa - As' = k - \kappa + t.$$

Each variable component lies in $[(\delta - \varepsilon)U, (C_0 + \delta)U]$; the constant components vanish. A ball of radius c_2U in the affine arm fiber fixing the constant rows preserves all these margins. If there are c constant rows, coarea contributes an $(n-c)$-dimensional volume of order U^{n-c} and a fixed Jacobian. Multiplying the arm density U^{-n} and the $m - c$ increment densities $U^{-(m-c)}$ gives U^{-m}. Only finitely many forest matrices occur. This calculation includes the endpoints of the old support by the strict fresh margins, and uses no lower bound for $B_i - m_i$.

The fresh forest does not keep B_i as a parallel edge of each variable row: its required increment can exceed that smaller-hole cap. It retains all genuine routes avoiding the complete modification region. Every edge is still an actual physical curve. Empty forests have no current and need no right inverse or density construction. □

The same zero-current convention handles a one-port body.

The cap convention is essential. If an inner avoiding-D route of cost B_i is retained in parallel with a fresh variable row and shares the same outer arms, the actual row becomes $\min(B_i, \kappa_i + T_i) + (AS)_i$. The arm shift in the last proof may require $\kappa_i + T_i = K_i + \delta U > B_i$, invalidating that displayed density map. Accordingly, the forest replacement restricts paths throughout the complete coarse body; only its genuine exterior bypasses are fixed in every term. This preserves the required physical-curve property and every route avoiding the complete modification region. A formulation requiring all smaller-hole bypasses to survive needs a different realization theorem.

20 The paired local replacement

20.1 The query state and its mean convention

Fix the pure real base a, a finite port cutoff M, and deterministic buffered regions

$$H_{\mathrm{old}} \Subset H_0 \Subset H_c\,.$$

These coherent projection domains are chosen from the fixed countable source-grid catalogue before any private atlas phase is sampled. A continuous private phase moves only reserve supports and local tubes inside their fixed annulus. The random modification body C is not a coherent projection domain; its complete declaration is contained in the selected H_c. Comparison holes are fixed enlargements of the physical scale-s partition cells, and the auxiliary grid coset below changes only the reference exploration, not an opaque record's H_0 or H_c. The complete old carrier has diameter at most r, while H_0, H_c have diameter comparable to $R \gg r$. Every old source support, old observation, old reserve that has already been used, and original output label belongs to that carrier. For a translated assembly this means each translated copy $H_{\mathrm{old}}(c)$; all copies retained by the smooth whole-owner weight lie in one fixed strict inner buffer of H_0. The parameter r bounds each complete carrier, not the diameter of the union of all its translates. The new arm atlas is in a fixed annulus of $H_c \setminus \overline{H}_0$. Its complete modification region is denoted by C, and its protecting region by $C^+ \Subset H_c$. The outside query terminals lie beyond C^+; the corresponding genuine avoiding paths and surrounding loop are retained. For an auxiliary port-only frame the atlas instead retains its positive whole-port lower bounds and finite physical cap bounds. No separating-crossing assertion is made for a terminal already on C.

For port-only redistribution we retain an actual pointwise lower bound, not only a lower bound for an expectation. Here the nontrivial rows are actual physical path infima between distinct outer ports of the selected atlas; abstract exterior graphs enter only as scalar tests. At zero amplitude choose, for each of these finitely many ports p_i, nested Euclidean neighborhoods $p_i \in B_i^- \Subset B_i^+$ whose closures lie in the unchanged observed annulus W, miss $\overline{H}_0$, and are pairwise disjoint. They can be selected from a countable rational catalogue after the port template is selected. Let ℓ_i be the intrinsic crossing distance between the two boundary components of $B_i^+ \setminus \overline{B_i^-}$. This is an exterior observation and is strictly positive almost surely: the ordinary metric is lower semicontinuous and separates points, and the two boundary components are disjoint compact sets. Every physical traversal from port i to a different port must pay this unchanged crossing before it can reach that other port or the hidden filling. The same lower bound therefore holds for an unrestricted bypass row, for a compulsory-arm row, and for any restriction of their physical path classes, uniformly over all compatible fillings of H_0.

Include, in the atlas's finite numerical exhaustion and before choosing the later contraction constants, the exterior condition $\min_i \ell_i \geq c_M U$. Since there are finitely many ports and U is

positive and finite, these conditions exhaust the admissible exterior states as $c_M \downarrow 0$. Consequently every nonempty actual frame satisfies

$$L_\Sigma(X, 0) \geq c_M U$$

for every compatible filling. Diagonal and other identically zero rows have zero current and are omitted; if the frame is empty, no occupation normalization is made. The subsequent Gaussian Weyl comparison gives the lower bound at arbitrary values of the original amplitude variables.

Let e_0 consist of the residual field outside H_0, the nonselected outer coordinates, the original source phase, the independent atlas phase, and the selected finite template. Selected outer arm coordinates are unexposed. The full original success indicator is the residual event Γ together with the window on the entire original outer menu. Its nonselected-coordinate factors are fixed in e_0, and its selected-coordinate factors are integrated with their original probability; no window probability is discarded. The surrounding loop length $U = U(e_0) > 0$ is known at this time.

All physical covariance perturbations and incoming means are placed in the *residual* Gaussian variable. In an augmentation $F = X + GA$, the law of A remains its original centered independent Gaussian law throughout the mean family. Thus F has the intended physical mean, but the old arm density and the scalar seed densities remain fixed. The residual covariance is at least half the pure covariance. Under the real reference law, before inserting the single residual perturbation likelihood, all outer reserve functions vanish near H_0. Conditional on e_0, the reference inner field is therefore exactly

$$F|_{H_0} = h_{e_0} + a\,\mathrm{GFF}_0(H_0)\,,$$

and is independent of the selected outer coordinates. The harmonic prediction h_{e_0} is unaffected by those coordinates.

The retained conditional class has deterministic upper and lower arm and whole-frame bounds, the strict-buffer GOOD prediction bounds, and the fixed-block exterior tests of Lemma 16.2 with the local fresh-row exhaustion. Its middle Gaussian mean score has bounded norm in the dual of the coherent mean space. The latter is an actual retained event; a bound on the Cameron–Martin norm of a mean does not by itself bound its random likelihood. These tests are functions of e_0 and are selected before any old inner variable is integrated.

An admitted scalar test is $f(D)$, where D is obtained by gluing the local body to a fixed positive exterior graph. Its edge weights may depend on e_0 but not on a hidden filling. The graph contains genuine fixed exterior routes connecting all its vertices, as required in Proposition 18.2. This includes every actual buffered query in the argument. The local atlas, companion, and avoiding kernels are intrinsic to H_c; a neighboring hidden field is never read to select them. An ordinary local query may have fixed compulsory-arm exterior wiring. Its paths minimize the ordinary physical metric freely in a strict interior buffer of H_0, where every local replacement is allowed and every root used below lies. Its body row need not equal the unrestricted ordinary distance between the same outer ports: the latter also minimizes over the retained exterior bypasses. The reference argument uses only these genuine local minimizing fragments and the stated whole-frame normalizer and moments.

A finite complete physical expression P has an original first-current mark $q_P(F, x)$ obtained by unfolding its chosen labelled primitive after its full original history partition has recombined. Every identity used below preserves that same original-output primitive; equality of unmarked physical operators alone would not justify a replacement under a nonconstant label. Any finite auxiliary labels in the unfolded mark remain part of its complete declaration. The protected input is chosen before its selected histories are expanded. Its original source identifier and label are followed

through every branch and conversion, even if that branch is raw while another input is converted. Thus unfolding gives the same original finite Gaussian current, rather than a history-dependent choice of primitives. Its scalar center is $\beta = \Pi(P)$, the phase-independent pure-reference limit. No individual selected history or source mask is given a new center. In addition, $P(v)$ has an exact finite paired estimating representation throughout a fixed complex neighborhood B^+ of a bounded real coherent-mean ball B. Its current and value coefficients have the stipulated common positive envelope AW_v, with one fixed spare moment $p_* > 1$ under every admitted ordinary labelled frame and compatible reference. This is an estimate on the complete representation, including its old history factors; the reset does not recharge those factors. The envelope is valid for exterior flags and for the retained original output-label norm. That label has test norm $\|\phi\|_\infty + R\|\nabla\phi\|_\infty$ at output scale R.

Write the ordinary whole-port frame as

$$d(e_0) = E_{e_0} \sum_{\ell=1}^{L} q_\ell, \qquad Q_{e_0}(dF, d\ell, dx) = d(e_0)^{-1} P_{e_0}(dF)\nu_{q_\ell}(dx), \qquad L \le M^2 . \tag{20.1}$$

The retained zero-amplitude arm and cap bounds and finite Gaussian Weyl comparison give

$$c_M U \le d(e_0) \le C_M U, \qquad E_{e_0}(\textstyle\sum_\ell q_\ell/U)^p \le C_{M,p} .$$

These are whole-frame bounds. An individual activated middle is not assigned a uniform positive normalizer. This frame is a sum, not a union occupation: the row label ℓ is retained in every conditional state, and the one-path common-component coupling is applied separately at that label. Several active rows may each contribute the same canonical middle occupation; their sum is not identified with one copy of it. The finite row count is paid once before division by $d(e_0)$. Only roots in the strict inner support of the whole-owner weight need the GOOD-count and endpoint-clearance estimates.

20.2 The past-array estimate

Lemma 20.1 (Exterior marginal weights on coupling failure)**.** *Consider a finite annular coupling of two prescribed rooted laws. Its exploration stops at a deterministic bottom mesh before it can read any part of either complete protected payload. A payload includes the original source range, every retained old reader and reserve coordinate, and all regions used by its estimating current and ordinary restoration. The explored cell address is constant on each such whole payload. Choose the coupling once per actual rooted cell, original row, and retained layer label, independently of the fine source anchor. At each trial the common-submeasure acceptance density and the rejection decision are functions of the two revealed annular histories and auxiliary coins. The marginal annular transitions are the prescribed ones. The conditional rooted kernel uses the whole occupation of that row in the cell, summed over all its excursions. No root-dependent excursion or trace-component label is introduced. The canonical accepted component is a submeasure of this whole row kernel.*

On the branch on which no trial is accepted, complete the remaining hidden payloads independently using their respective original conditional rooted kernels. Denote by E_1, E_2 the marginal annular fields and retained row, layer and cell addresses down to this fixed bottom mesh. These data are revealed on failure. On success they may be defined afterwards from the completed sample solely for disintegration; they do not enter the earlier stopping decision. Thus the marginal cut is deterministic in scale. Denote the remaining hidden payloads by X_1, X_2, and their original kernels by $\mathsf{K}_i(E_i, dX_i)$. There is a subprobability measure Λ on the pair of histories such that the failed

joint law is

$$\mathbf{1}_{\mathrm{fail}}\, d\mathbb{P}_{\mathrm{pair}} = \Lambda(dE_1, dE_2)\mathsf{K}_1(E_1, dX_1)\mathsf{K}_2(E_2, dX_2). \tag{20.2}$$

If the ith prescribed marginal is $\pi_i(dE_i)\mathsf{K}_i(E_i, dX_i)$, the projection of this failed law onto that marginal is

$$b_i(E_i)\,\pi_i(dE_i)\mathsf{K}_i(E_i, dX_i), \qquad 0 \le b_i \le 1. \tag{20.3}$$

In particular the marginal failure weight is exterior to the complete payload. If a multiplier Y of modulus at most one is measurable from the joint revealed histories, its corresponding signed marginal weight $\widetilde{b}_i$ is also exterior and satisfies $|\widetilde{b}_i| \le b_i$.

For a fixed marginal physical field and retained row/layer labels, b_i is constant as the root varies inside its final bottom cell. Consequently it is unchanged when the represented current is moved within a complete good carrier contained in that cell. This statement uses one cell coupling for all fine source anchors; it does not follow from constructing a separate exterior weight for each anchor.

Proof. Use the usual prescribed marginal annular transitions until a trial is accepted or the deterministic bottom mesh is reached. An acceptance density that is annular-measurable removes an annular subkernel and leaves the conditional unread-core kernel unchanged on rejection. Induction therefore defines the failed annular-history measure Λ. Completing the two final cores by the stated product of their original conditional kernels gives (20.2); it does not inspect either core to decide failure. All retained old augmentation coordinates are completed from their original conditional regression kernels as part of K_i. They are not deleted.

Let Λ_i be the projection of Λ onto E_i. Marginal preservation implies $\Lambda_i \le \pi_i$. Hence $b_i = d\Lambda_i/d\pi_i$ exists, is between zero and one, and yields (20.3). Projecting $Y\Lambda$ instead gives a signed measure whose variation is bounded by Λ_i and proves the last assertion. These arguments also apply to random pasts by first disintegrating over that past. Joint annulus-measurability alone, without the completion in (20.2), would not establish the marginal assertion.

The own marginal history records the annuli, original row, retained layer, and cell addresses down to the bottom mesh, but no position or excursion label within its final cell. The kernel K_i includes that root position. The formula for $b_i(E_i)$ therefore also proves the asserted constancy within the final cell. The whole-carrier mask puts the original and restored current supports in the same such cell. □

Lemma 20.2 (A bounded array outside the comparison holes)**.** *Let χ be the smooth whole-owner weight, supported in the retained strict inner buffer of H_0, with $0 \le \chi \le 1$ and $\|\nabla\chi\|_\infty \le C/R$. Write $w_Q = \chi\psi_Q$, where (ψ_Q) is a nonnegative scale-s partition of unity with bounded overlap. Thus $\sum_Q w_Q = \chi$ and $\|\nabla w_Q\|_\infty \le C/s$. The comparison hole for Q is a fixed enlargement of its support, so it contains a strict buffer around every finer cell meeting that support. Suppose each $Y_{Q,\ell}$ is bounded by one and measurable from the designated past outside that comparison hole, together with its retained labels. Then, under the one conditional law P_{e_0},*

$$\left|\sum_{Q,\ell} E_{e_0} Y_{Q,\ell} \int w_Q(x)(q_P - \beta)\, d\nu_{q_\ell}(x)\right| \le C_M A (r/s)^\eta d(e_0) \tag{20.4}$$

for some $\eta > 0$. Original output labels may first be transported within their complete old radius-r carrier, with the corresponding Cr/s error. The same conclusion holds on each bounded real mean ball, with constant $C_{M,B}$ and the additional positive-power coherent-innovation error in s/R.

Proof. We first work at zero incoming mean. Let $D \le Cr$ be one deterministic radius containing the complete original and estimating carrier, including every retained reader, reserve support and restoration region. Use an intermediate mesh $t \asymp (Ds)^{1/2}$ and a deterministic bottom mesh $R_- \asymp (Dt)^{1/2}$, with compatible nested grid rounding. Increase the fixed minimum accepted age so that R_- contains the required strict buffer around the whole radius-D carrier. The forgetting coupling stops at R_- and never reads this carrier. In the following whole-owner boundary estimates use D in place of the bare source radius r.

There is a separate whole-owner issue before any root comparison. Choose one common radius-D bounding carrier for every translate of each term in the complete centered orbit, and give the drift lift the same declaration. Choose an auxiliary grid coset uniformly among the fine-lattice translations modulo R_-. It is independent of the old record, the scalar-capture data, and the array Y. Declare a whole carrier bad if it meets a boundary of this bottom grid. Its probability is at most CD/R_-. Every coarser trial boundary is a bottom-grid boundary. Average the entire bad-owner branch over this private coset before estimating fields: it is the corresponding deterministic probability times the whole centered expression and has cost at most CAD/R_-. No signed output-label measure is cut into pieces.

The good-owner mask is R_--periodic. For a fixed coset, the masked complete orbit has pure-reference center zero. To check this exactly, regard it as a finite prototype repeated on $R_-\mathbb{Z}^2$. Translating the coset by a fine lattice vector translates that prototype, so phase independence of Π makes its center the same for every coset. Averaging the cosets gives a constant multiple of the original centered orbit, including its identically masked drift lift. Its center is zero, hence the common masked center is zero. Reference convergence for this auxiliary prototype is needed only at fixed r, s, t, R_- before the reference size tends to infinity; its quantitative constant need not be uniform in R_-/r.

On a good carrier every original output point and every read in the carrier has one trial address at every scale down to R_-. Acceptance and failure flags are consequently single exterior scalar flags for that whole old record. Moreover, an accepted-cell translation is a multiple of its mesh, which is at least R_-, so it preserves the bottom-periodic mask. Thus the masked original rewards agree on matched cores. The auxiliary coset remains averaged inside this estimating calculation; it does not change the fractional physical source phase or create an additional read-address count.

A smooth weight is not itself designated past merely because its coefficient is. To retain an exact ordinary core law, use the layer-cake identity

$$w_Q(x) = \int_0^1 1_{\{u<w_Q(x)\}}\,du$$

and retain the auxiliary label (Q, u). Add a cemetery label with weight $w_* = 1 - \chi$ and coefficient $Y_* = 0$. The augmented weights now sum to one, so this augmentation leaves the marginal root law exactly Q_{e_0}. The cemetery label makes no contribution and needs no local source comparison.

Call a mesh-t cell a boundary cell for (Q, u) if u lies between the minimum and maximum of w_Q on that cell. For each fixed root position, the sum of the bad u-masses is at most Ct/s, by the Lipschitz bound and bounded overlap. Fubini with the original whole occupation gives boundary probability at most Ct/s, before any conditional renormalization. Away from these cells, the retained capture indicator is identically one throughout the cell and all its future subcells. Its conditional transitions are therefore the bare ordinary Campbell transitions of the pure-reference construction.

For each selected cell and its designated past, generate a sufficiently large aligned finite reference independently, copying its physical source phase. The reference grid is translated with the cell,

so every finer source lattice has the same relative phase on matched cores. Run the marginal-preserving annular coupling strictly *inside the comparison hole*, from scale t down to $R_- \gg r$. The scalar-capture seed annuli at scales at least $u \gg s$ are already designated past and are not resampled by this coupling.

On the branch where every attempted coupling rejects, complete the remaining cores exactly as in Lemma 20.1. An accepted pair is stopped at once; subsequent observations never determine whether it counted as a success.

The selected conditional starting state need not be GOOD. Run the prescribed marginal transitions through all bad states and use the common submeasure only at paired GOOD opportunities. After averaging over the original whole labelled root law, its absolute-start GOOD-count and initialization estimates, together with those of the independent reference, give a positive fraction of paired GOOD opportunities. The annulus-only acceptance density and successive rejection bound consequently give failure probability, reducing the positive exponent because $R_-/t \asymp \sqrt{D/t}$,

$$C(D/t)^{\eta_0} + Ct/s + Ct/H$$

when the finite reference has outer radius H. No uniform conditional failure estimate under an arbitrary rare past is asserted.

On success the normalized physical cores and root labels agree on the entire old module. Each original compact score has the form

$$I_v(F) = \langle F, -\Delta v\rangle/(2\pi a^2)\,,$$

so it is invariant under the scalar and harmonic differences removed by this matching. Deterministic Wick contractions agree as well. Hence the *original* complete mark is identical in the two cores. Estimating forest rewards are not used in this equality.

Conditional marginal preservation gives, exactly at finite H,

$$E[Y_{Q,\ell} q_H^{\mathrm{ref,good}}] = E[Y_{Q,\ell}\beta_H^{\mathrm{good}}(\Theta)]\,.$$

Here the reference carries the same fixed good-owner mask and its full centered prototype. The phase-independent reference theorem and the preceding coset argument give $\beta_H^{\mathrm{good}}(\Theta) \to 0$. At each fixed scale triple the coset family is finite, so the limit is uniform over its cosets. If the ordinary conditional cell mass is Z_Q, it multiplies both terms in this formula: one pairs $Y_{Q,\ell}Z_Q$, then sums the ordinary cell contributions, and divides only by the original $d(e_0)$. The positive partition weights have total one. This is the normalization which prevents a loss proportional to the number of comparison cells.

On failure, apply Lemma 20.1, retaining the original row, layer label and whole-owner address. Write $\widehat{Q}_i$ for the two complete augmented marginal root laws of this array coupling. They include the original ordinary cell masses, the cemetery label and the single whole-frame normalization. For each marginal the partner and its auxiliary coins give an exterior weight $\widetilde{b}_i$ with $|\widetilde{b}_i| \le b_i \le 1$ and $E_{\widehat{Q}_i} b_i = \mathbb{P}_{\mathrm{pair}}(\mathrm{fail})$. Each b_i is constant throughout its final root cell, and the same coupling and the same b_i are used for every good fine carrier in that cell. Moving a represented current within its complete carrier therefore retains the same failure multiplier. Here restoration also retains the same ordinary row: every contributing complete carrier and restoration region is in the strict central buffer of H_0, whereas the ordinary row terminals are outside H_0. The fixed scale delay leaves an unchanged surrounding loop and separating crossing around each old carrier in the freely replaceable ordinary region. Reference endpoints likewise lie outside their capture buffer. Thus the coordinatewise local case of Lemma 20.4 applies. The port-only estimate obtained by

redistributing current over other rows of a whole frame is not used for this row-dependent failure multiplier. Apply the exact labelled identity of the chosen cheap representation with this exterior multiplier. Its common spatial-current envelope, summed over the translated array by Lemma 21.3 before estimating, gives

$$\left|\mathcal{I}_i(\widetilde{b}_i)\right| \le C_M A\, d(e_0)\, E_{\widehat{Q}_i}[b_i W_i] \le C_M A\, d(e_0)\, \mathbb{P}_{\mathrm{pair}}(\mathrm{fail})^{1-1/p}, \qquad \|W_i\|_{L^p(\widehat{Q}_i)} \le 1.$$

This estimate is applied after summing the original cell masses; it does not assert a uniform conditional failure probability for each selected cell. For the reference marginal, first integrate its partner; the signed exterior weight is still bounded by its own failed projection. The same bound is uniform in the finite reference size and in the auxiliary coset. It uses no L^p estimate of the unfolded raw Wick mark. On success equality is used only for the original recombined mark on the fully matched core. In a cheap representation every retained read history remains in its coefficient, or in its posterior coefficient from Lemma 4.2.

The auxiliary masked prototype has center zero and its qualitative reference limit is taken at fixed D, s, t, R_-. Thus the uniform cheap bound can be passed to that limit. The unmasked whole-frame estimate also gives $|\beta| \le C_M A$. The carrier-boundary error D/R_-, the level-set error t/s, and the unsuccessful coupling probability are positive powers of D/s. Since the lemma's complete carrier has $D \le Cr$, they retain a positive power of r/s. In the later small-owner application the declaration is $D \le CN_*r$, and the same argument has a constant depending on the fixed N_*. Consequently the use of a deterministic bottom cutoff, and retention of the complete old history, do not change the assertion that the accepted coefficient tends to zero as $R/r \to \infty$.

Letting $H \to \infty$ in these estimates proves (20.4) at zero incoming mean.

For a real coherent mean $v \in B$, split its residual shift by the original Dirichlet projection on the comparison hole. Its harmonic part stays in the outside likelihood and does not affect the compact Wick mark. Its inner innovation has coherent norm at most $C_B(s/R)^{\eta_1}$. The input Cauchy bound on B^+ pays this discrepancy. A designated-past likelihood is kept in its original residual Gaussian variable. If it is not already bounded by the retained dual-score event, truncate it once: its supremum over B has fixed Gaussian moments because the coherent-to-Cameron–Martin map is Hilbert–Schmidt. For a threshold T, the bounded part costs $T\varepsilon_{\mathrm{mix}}$ and the tail costs $C_B T^{-c}$ by the common current envelope and one spare joint moment. Taking T to be a small negative power of $\varepsilon_{\mathrm{mix}}$ preserves a positive power. This is an estimate on the unchanged analytic numerator; it does not define a mean-dependent companion or an analytic selector. The middle likelihood at e_0 is paid once. This proves the bounded-real-mean assertion. □

20.3 The reset and its physical lift

Theorem 20.3 (A local paired reset on every fresh fiber)**.** *Fix a legal complete representation $\mathcal{R}$ of P, its chosen original first-current primitive, and the actual accepted history event Γ. Let χ be the smooth whole-owner weight at scale R. Write*

$$\mathcal{I}^{\mathcal{R}}_{\chi,\Gamma}(P - \Pi(P)D_1; v, f)$$

for the accepted, weighted action of this primitive: both centered legs carry the same weight, the same actual history predicates, and their original labelled row masses. The center is the center of the complete unmasked physical expression P; it is not recomputed for a selected history. The complementary history terms are retained in the complete representation. For the preceding query class there are $\eta > 0$ and $C_{M,B} < \infty$ such that this accepted action has the exact finite signed

paired representation

$$\mathcal{I}^{\mathcal{R}}_{\chi,\Gamma}(P-\Pi(P)D_1;v,f)=\sum_{\tau}E^{+}_{e_0,\tau}h_\tau(v)f(D^{+}_{\tau})=\sum_{\tau}E^{+}_{e_0,\tau}D_{u_\tau(v)}f(D^{+}_{\tau}). \tag{20.5}$$

Each D^{+}_{τ} is a positive directed allowed-path forest body whose edges are actual physical path costs. It retains every genuine route avoiding the complete modification region C. Its fresh companion law, body maps, reserve basis, and density division are fixed at real mean zero, independently of the old fine labels and of v. On every retained fresh residual fiber,

$$\sum_{\tau}\bigl(|h_\tau(v)|+\|u_\tau(v)\|_\infty\bigr)\le C_{M,B}A(r/R)^{\eta},\qquad v\in B\,. \tag{20.6}$$

The value measure has mass zero on each such fiber, and its value and current identities remain paired there. The full holomorphic mean identity on B^{+} is preserved. The original output label is retained after the current is moved to the fresh outer arms.

Proof. Choose

$$s\asymp r^{1/2}R^{1/2},\qquad u\asymp r^{1/4}R^{3/4}\,,$$

with fixed buffers and compatible grid rounding. For each comparison hole at scale s, seek the actual serial capture of Lemma 17.1 at eligible scales $v\in[u,c_0R]$, where $c_0>0$ is fixed so that every capture annulus and all its seed supports are contained in H_0. The fixed finite seeded catalogue has a uniform positive success probability. For the retained nonempty menus this is the residual-product extension in Lemma 10.1: all cuts use the residual data on the full enlarged box, and all original coordinate-window probabilities and row weights remain in the rooted submeasure. Reconstruct unread inner menus only in the complete original physical mark used by the coupling. In an estimating integrand retain every read multiplier, using its posterior coefficient under Lemma 4.2 if the physical field is recombined. A residual with an unread inner reserve is not itself a pure Dirichlet filling. The conditional kernel P_g is declared only after this reconstruction; every remaining past reader is supported outside its protected core. The ordinary GOOD-count argument and annulus-only acceptance give failure probability $C(u/R)^{\eta_2}$. All these trials lie outside the comparison hole. Its radius-s closure lies inside the larger forced capture core. After acceptance, expose the remaining middle down to that comparison hole. In its simultaneous scalar formula the bypass is the genuine infimum over routes avoiding the comparison hole and all seed supports, including their transitions; it may use other unchanged middle portions. This bypass is a past quantity. It is distinct from the genuinely C-avoiding routes retained in every final forest term. The trials are separate from the reference-forgetting trials inside s in Lemma 20.2.

The pure loop-ratio estimate, the derived past-array estimate, and Proposition 17.4 yield typed primitives satisfying

$$\sum_{\tau}\|\rho_\tau(v)\|_{\mathrm{TV}}\le C_{M,B}A(r/R)^{\eta}U,\qquad v\in B\,. \tag{20.7}$$

On scalar-capture successes the estimate is uniform over arbitrary covectors of all old ordinary, symmetric port costs. The simultaneous single-middle identity is used only for this original matrix. When a scalar test is replaced by a directed estimating body, the rowwise effective middle of (17.3) lies in the same free-hole cap interval and suffices for the lower-seed activation and remainder bound. The pure-reference coupling compares only ordinary locally minimizing rows; it never compares directed estimating fillings. Failed captures use their separate existing estimating primitives and the ordinary envelope. The complete inner distance matrices are finite, with directed triangle inequalities, because the untouched outer annulus supplies genuine finite avoiding routes.

Apply Proposition 18.1 to the completed matrix after attaching the positive arms. Its types and coefficients are independent of the arm lengths. The Sobolev-locality argument differentiates the full valuation on each stratum. It does not discard a derivative of a selected combinatorial label.

Integrate all old comparison residuals, seed coordinates, labels, and history variables inside H_0 before declaring (20.7) complete. The original old output label is frozen on comparison cells only for its C^1 testing error; it remains a coefficient label thereafter. No integrated old variable is later conditioned upon again.

For a forest τ, let A_τ be its unsigned port-arm incidence matrix, with the rooted right inverse R_τ. Its actual coarea Jacobian is $\sqrt{\det(A_\tau A_\tau^{\mathsf{T}})}$. Let g_τ be the fixed quotient density of the old original arm law. The three-piece arm bounds imply

$$\|g_\tau\|_\infty \le C_M U^{-m}, \qquad \|\nabla g_\tau\|_\infty \le C_M U^{-m-1},$$

where m is the number of forest rows. Set

$$\sigma_\tau(v) = \rho_\tau(v) * g_\tau, \qquad \mu_\tau(v) = -\operatorname{div} \sigma_\tau(v).$$

Thus $\int \mu_\tau(v) = 0$, and

$$\|\sigma_\tau(v)\|_\infty \le C_{M,B} A (r/R)^\eta U^{1-m}, \qquad \|\mu_\tau(v)\|_\infty \le C_{M,B} A (r/R)^\eta U^{-m}.$$

Use Proposition 19.1 inside H_0, on the retained local exterior class. Its residual-only event has probability bounded below uniformly in e_0. On each retained residual fiber ω, it supplies actual fixed row curves with costs $\kappa + T$. The curves avoid all outer arm supports. Constant rows retain their exact exterior kernels. Widen the original outer arm windows within the atlas's same enlarged box. The explicit arm-shift and coarea argument of Lemma 19.2 then gives the fresh output map

$$Y_\omega = \kappa + T + A_\tau S$$

a density $g_{+,\omega} \ge c_M U^{-m}$ on the entire old convolution support, on *every* retained residual fiber. Inner avoiding-H_0 costs are old row bounds and classifiers; they are not automatically parallel caps of each variable forest row. Genuine routes outside the complete modified body remain in every term.

Divide on that fiber:

$$h_\tau(v; y) = \mu_\tau(v; y)/g_{+,\omega}(y), \qquad a_\tau(v; y) = \sigma_\tau(v; y)/g_{+,\omega}(y),$$

with zero values outside the numerator supports. This gives bounds of order $A(r/R)^\eta$ and $A(r/R)^\eta U$, respectively, and gives the same signed measures μ_τ, σ_τ after multiplication by the fresh density. Therefore their integration-by-parts pairing and zero value mass hold on each individual fiber, not just after averaging fibers.

For the physical lift, differentiate one selected seed coordinate in each *original outer arm*. If its slope is $b_i = \partial S_i / \partial A_{i1}$, then $c_M U \le b_i \le C_M U$ almost everywhere on the enlarged box. Restricted to these coordinates,

$$DY_\omega = A_\tau \operatorname{diag}(b_i), \qquad \mathcal{R}_\omega = \operatorname{diag}(b_i^{-1}) R_\tau, \qquad \|\mathcal{R}_\omega\| \le C_M/U.$$

Use the coordinate direction $\mathcal{R}_\omega a_\tau(v; Y_\omega)$, multiplied by the corresponding physical smooth reserve functions. It is the actual physical field direction $u_\tau(v)$ and has the claimed supremum norm. The fixed fresh row excursions are not differentiated. The almost-everywhere chain rule suffices because the coordinate laws have densities. Empty forests have zero body current and are omitted.

For the retained labels, take the separable C^1 test space on the compact old carrier. Define the vector measures, convolutions, and density quotients by pairing with that space. The scalar identities on a countable dense set, together with the displayed uniform dual-norm bounds, determine the same weak-* densities and the fiberwise identities. No vector Radon–Nikodym theorem is being assumed.

The finite-dimensional change of variables and integration by parts give the accepted weighted identity (20.5) for smooth scalar tests. The common physical caps and current bounds extend it to Lipschitz tests. Every density, normalization, and physical basis used in the division belongs to the fixed real companion. Convolution, differentiation, division, and lifting are bounded linear maps on the integrated numerator. They preserve its holomorphy on the same mean neighborhood. The original residual mean likelihood is kept once when enlarging from H_0 to H_c; a new likelihood is not placed on the fresh physical field. The sum over finitely many types and the label-freezing errors is absorbed in $C_{M,B}$ and in a smaller positive exponent. This proves the theorem. □

20.4 Restoration to one ordinary occupation frame

Lemma 20.4 (One physical current and one loop ratio)**.** *Fix one real reconstructed physical field and an allowed-path estimating body whose every edge lifts to an actual physical curve. Fix its exterior wiring throughout, and let D_{ord} denote the ordinary interior metric query with this same exterior wiring. Thus the admissible exterior routes and their fixed costs agree for D_{ord} and D_{est}; only the interior path class is restricted in D_{est}. This convention includes the actual physical query frames. The unchanged exterior is freely spliceable. Suppose the allowed class retains every ordinary route avoiding C and permits the surrounding-loop replacement, and let U_C be the length of a retained surrounding loop. Query terminals lie outside its protecting region C^+. Suppose every ordinary visit to C pays a separating crossing $L_C > 0$ inside C^+. Let u be a physical current direction supported on the ordinary outer arm pieces, represented in their nonnegative private bump basis. Define its arm coefficient norm as $\max_i |b_i| \|\sigma g_i\|_\infty$ when $u = \sum_i b_i \sigma g_i$; disjoint supports and nonempty plateaux identify this with its physical supremum norm (up to the fixed finite overlap convention). At almost every retained amplitude,*

$$|D_u D_{\mathrm{est}}| \le C_M \|u\|_{\mathrm{arm}} \frac{U_C}{L_C}\, \nu_{D_{\mathrm{ord}}}(C^+)\,. \tag{20.8}$$

The norm $\|u\|_{\mathrm{arm}}$ includes the physical bump normalization, so there is no unrecorded factor from a small σ. It is bounded by C_M times the norm furnished in the preceding lift. For a port-only frame use its retained positive whole-port lower bound and physical cap instead of the loop/separating-crossing comparison.

The fresh companion can consequently be restored to this same ordinary physical field at one fixed density cost. For products of disjoint closed children, choose one current child and all other children in their value representations. The resulting occupation estimate uses one containing modification region, one ordinary labelled frame, and one loop ratio.

The same estimate holds for a smooth compact physical direction supported in C when the allowed path family, its wiring, and all its retained coefficients and readers are fixed under this variation. This includes a protected raw primary with other disjoint children in their value representations. Its direction need not be a fresh private-arm basis vector.

Proof. Every allowed curve is physical, so $D_{\mathrm{est}} \ge D_{\mathrm{ord}}$. Let D_{avoid} be the infimum over genuine C-avoiding routes. It is retained, so $D_{\mathrm{est}} \le D_{\mathrm{avoid}}$. If $D_{\mathrm{avoid}} = D_{\mathrm{ord}}$, increasing any nonnegative private bump supported in C leaves both infima equal to D_{avoid}: their lower bound is the original

ordinary distance and their upper bound is the unchanged avoiding infimum. Hence every such right derivative of D_{est} is zero. At the almost-everywhere differentiability points supplied by the absolutely continuous amplitude laws, its two-sided derivative is zero. This argument also covers an avoiding infimum that is not attained. When minimizers exist, the same conclusion follows because first-variation comparison makes all ordinary minimizing curves have the same integral of the private bump; an avoiding minimizer has integral zero. No uniqueness assertion for random port endpoints is needed.

Thus a nonzero arm current requires $D_{\mathrm{avoid}} > D_{\mathrm{ord}}$, and forces an ordinary visit to C. For an estimating minimizer, retain the portions before its first and after its last contacts with the surrounding loop. Joining those contacts by a loop arc is a genuine C-avoiding competitor and remains allowed. Minimality bounds the total cost of its intervening portion, hence the sum of all differentiated arm lengths, by U_C. The ordinary visit contributes at least L_C to its occupation in C^+. The first-variation formula gives (20.8). For an allowed infimum that is not attained, apply the loop replacement to an ϵ-minimizing competitor; the intervening cost is at most $U_C + \epsilon$. Directional comparison with these competitors and $\epsilon \downarrow 0$ gives the same derivative bound. If $D_{\mathrm{avoid}} > D_{\mathrm{ord}}$, all sufficiently accurate ordinary competitors visit C and pay the retained crossing in C^+; the ordinary metric frame is obtained from its confined minimizing limit. Thus the estimate is unchanged by use of intrinsic infima.

For this extension use the strict buffer of the protected source range to choose a nonnegative smooth bump $\sigma \geq |v|$ supported in that same protected range. It therefore avoids every reader and the other children, as does v. If $D_{\mathrm{avoid}} = D_{\mathrm{ord}}$, monotonicity and the unchanged avoiding competitors show that the right derivative in this bump is zero. The metric cost is locally Lipschitz along this compact direction. Fubini and equivalence under its Cameron–Martin translations give two-sided differentiability at the original amplitude almost surely. The derivative in σ is consequently zero, and monotonicity between the perturbations $\pm t\sigma$ gives derivative zero for the signed direction v. A right derivative alone would not suffice for this conclusion. In the other case the same first/last-loop replacement bounds the total physical length on which the direction acts by U_C. Its first variation is bounded by the supremum norm of the direction times this length. The argument with ϵ-minimizers is unchanged. Because the family and the readers are fixed, no derivative of a selector, path class or coefficient is introduced by this extension. For a source feature integral, perform this argument on the finite feature approximations and pass using the same positive Gaussian envelope and current bound.

For a nonempty frame put

$$B(A) = \sigma \sum_j |A_j| \|g_j\|_\infty.$$

The Weyl comparison, valid also outside the retained coordinate windows, gives

$$L_\Sigma\Big(X + \sigma \sum_j A_j g_j\Big) \geq c_M U e^{-B(A)}, \qquad U/L_\Sigma\Big(X + \sigma \sum_j A_j g_j\Big) \leq C_M e^{B(A)}.$$

The last factor has every fixed moment under the original finite Gaussian coordinate law. It is pointwise bounded on each fixed companion window. A physical current bounded by its retained cap $C_M U$ can therefore be distributed against the whole ordinary frame with density at most $C_M U/L_\Sigma$; include this factor in the one ordinary envelope. An expectation lower bound on L_Σ alone would not justify this construction. This port-only redistribution is still not used for the row-dependent forgetting flags.

The companion is obtained from the original finite augmented Gaussian law by conditioning one residual event of probability at least p_M and fixed amplitude windows of probability at least c_M.

Its density relative to that law is bounded by $(p_M c_M)^{-1}$. Pushing to the reconstructed physical field cannot increase this bound. Restore that density before estimating the ordinary occupation. The loop and inverse-crossing moment bounds then supply the single required ratio moment. Any already retained parent flags remain inside this one joint envelope.

For disjoint children, condition on their common exterior and shared coherent mean. Use the current identity for exactly one child and the value identities for the others. Their scalar coefficients multiply; their old histories are not reopened. After gluing there is one allowed-path body in the union of the declared regions. Apply the preceding pathwise argument to one containing region and its unchanged outer loop. There is only one differentiated path occupation. The finite port multiplicity is absorbed once in C_M; one does not multiply separate child occupation measures or separate loop ratios. The finite companion density costs are paid at their declared births or in the one common record envelope, rather than charged again at every later reset. □

21 Complete activities and their estimates

Fix a positive base parameter a. All constants in this section may depend on a. The real-pressure comparison on both sides of a means that uniformity in a varying base parameter is unnecessary. We use the geometry of Lemma 15.1 and the pure-reference comparison of Theorem 7.1. The source contours are chosen after the mean weights, in the order specified below. Every Gaussian source below is a source for this pure base.

The stage spaces at different absolute radii are identified by translation from their common physical lattice origin and by dilation. Under $x = ry$, the pure metric contributes r^{aQ}; the whole-frame length unit U contributes the same factor, and the additive field scalar is removed by that normalization. Compact covariance features ϕ_r, physical mean masks, Dirichlet energies, the rescaled coherent spaces $E(H)$, and the original reserve templates transform by their defining rescalings. The positive refinement coefficients depend only on the mesh ratio, and the translated pure reference has the same scalar center by phase independence. Every legal record, its cost, and its normalized query therefore has the same form after this identification. In particular all microstage maps, and the block map $\mathcal{F}$ below, are independent of absolute depth. This uses the exact pure scaling law; no fixed infrared cutoff law is used in the bulk construction.

A finite physical expression is a finite linear combination of coefficients of products of the original compact covariance and mean germs

$$\exp\Big(\frac{q}{2}D_K^2\Big)\,, \qquad \exp(tD_\psi)\,.$$

Repeated occurrences of a physical coordinate are combined before its coefficient is taken. Feature integrals, positive assignment integrals, and complete finite history partitions may be included. The expression is the resulting constant-coefficient physical operator. A selected history, a Gaussian adjoint, or a converted local formula is an estimating representation of this operator.

An owner is a connected finite set X of current mesh cells, with one distinguished cell and one of the fixed protecting templates. Its complete declaration contains both legs of every source, every retained observation and original coordinate support, all maximal protecting domains and unchanged surrounding route regions, and the support of the original current label. Compatibility uses this complete declaration. A future observation is outside the whole inherited payload; at a merger the original physical lists are united, not replaced by new independent copies.

Here is the precise test class. A local body is glued through its declared ports to a finite positive directed exterior graph having genuine fixed directed routes between every pair of its vertices,

and the observable is $f(D)$ for the resulting scalar terminal distance D. This is the exterior class of Proposition 18.2; every actual buffered query has these routes in the unchanged regions around its disjoint guarded complete bodies. No artificial edge or reverse traversal is added to supply them. Initially f is bounded smooth; for a current we also use smooth f with bounded first derivative, followed by the dominated Lipschitz limits. Every allowed internal edge is an actual physical path cost or an intrinsic infimum of such costs. Routes avoiding the complete modification region remain available. The exterior graph, its edge weights and its terminal choices are fixed before sampling the local body; they may depend on the retained exterior. They may not inspect a hidden port matrix through another observable. An unintegrated expression may instead act on an ordinary whole-domain distance before any finite port elimination. Physical equality means equality of these unmarked scalar actions and total first-current actions throughout every deterministic smooth physical mean family, and its finite-dimensional holomorphic extension. All statements are conditional at every exterior where the specified versions exist. We do not impose a norm bound at arbitrary rough exteriors.

For a compact original-output carrier K at processed scale r, use

$$\mathcal{T}_r(K) = C^1(K), \qquad \|w\|_{\mathcal{T}_r} = \|w\|_\infty + r\|\nabla w\|_\infty, \qquad \mathcal{B}_r(K) = \mathcal{T}_r(K)^* .$$

One may use a fixed smooth compact enlargement and restriction to define this space. It is separable. If the source anchor is not fixed by the owner sector, take K in the joint anchor/output space $\mathbb{R}^2 \times \mathbb{R}^2$ and let the displayed gradient control both spatial variables. The anchor is not an additional freely oscillating discrete test coordinate. A record retains this original label; replacing its current by a fresh-arm current does not replace either original point. At a coarser processed scale R, the new unit test ball is contained in the old one, so its dual norm cannot increase. Freezing the label over a set of diameter d costs at most d/R times the common absolute-current envelope. Only the protected primary carries this label. All other labels are evaluated at 1 before a product is formed. A future selector may multiply the whole represented activity but cannot split its fine-label functional.

A normalized ordinary query has a deterministic buffered domain from the catalogue, an eligible retained exterior e, and the original Gaussian augmentation of Proposition 4.1. The unexposed Dirichlet cores and the selected original outer coordinates have their original, untruncated Gaussian law P_e. Eligibility consists of the fixed outer template and the numerical prediction, length, local row, and middle-score bounds. It is a retained indicator; its probability is never divided out. If the template has whole-port kernels K_ℓ, their entire paths, including compulsory arms and unchanged bypasses, form the labelled ordinary frame

$$L_\Sigma = \sum_\ell K_\ell, \qquad d(e) = \mathbb{E}_e L_\Sigma, \qquad Q_e(dF, d\ell, dx) = d(e)^{-1} P_e(dF)\nu_{K_\ell}(dx) .$$

There is one exterior unit $U(e) > 0$ with $d(e) \asymp U(e)$ and the fixed moments of $L_\Sigma/U(e)$ furnished by the source-free atlas. Individual path labels are retained, even when traces agree. No path is normalized by its own activated or captured mass. The actual pure-reference capture kernels and GOOD-started path/frame kernels are also admitted queries; they retain their own whole lower normalizer and the corresponding loop moment from Theorem 7.1. Enlarging a capture region is permitted only with that theorem's bounded normalizer comparison.

For such a query, an exterior flag $Z \geq 0$ is measurable outside the complete old record, with its designated-past labels retained. We use the following stronger, spatial form of the current estimate. In an actual typed estimating body τ, retain its whole port-cost vector, its original output label and its physical current directions. The vector current, with cost-coordinate index j, has a lift

$\widetilde{\mathfrak{J}}_{\tau,j}$ retaining these data and a restored ordinary spatial point x. Its projection onto the cost coordinates is the actual estimating vector current. Let pr forget the additional estimating data and keep (F,ℓ,x). There is one nonnegative envelope $V_{\mathcal{R}}$, supported in its declared restoration region $C_{\mathcal{R}}^{+}$ (a fixed enlargement of the complete old carrier in the locally rooted case), such that

$$\sum_{\tau,j} \mathrm{pr}_{*}\,|\widetilde{\mathfrak{J}}_{\tau,j}| \le C_B A e^{W_r(\varphi)} V_{\mathcal{R}}(F,\ell,x)\, P_e(dF)\nu_{K_\ell}(dx), \qquad \|V_{\mathcal{R}}\|_{L^p(Q_e)} \le 1. \tag{21.1}$$

The ordinary row label ℓ is distinct from the cost-coordinate index j. In locally rooted queries, whose terminals are outside the complete restoration carrier, the lift preserves the original row: a current in row j is restored against ν_{K_j} on its own local containing region. In a purely port-only query one may instead distribute a cost-coordinate current over the whole ordinary labelled frame, using its retained whole-frame lower bound. Only in that case may the spatial carrier be the entire declared query region. This latter estimate is not used for row-dependent failure weights or for the translated local array.

Variation uses the retained dual output norm and the sum norm in the finite current coordinates. The same envelope is chosen before any scalar derivative test, bounded whole-vector covector, or exterior flag. Such tests may depend on the entire cost vector of their actual typed body. In the locally rooted case an exterior flag may retain the original row label. In the generic whole-frame case a flag used in the following displayed bound is common to the ordinary rows; arbitrary row-dependent flags are bounded first by their common pointwise maximum. After an admissible exterior flag is inserted, (21.1) gives in particular

$$|\mathcal{I}(P;Z,f,w)| \le C_B A e^{W_r(\varphi)} d(e) E_{Q_e}[Z V_{\mathcal{R}}].$$

The construction also provides, below every eligible GOOD ancestor, the uniformly bounded ordinary conditional moments of each prescribed carrier envelope needed by Proposition 14.1. The value estimate uses the corresponding ordinary Gaussian envelope without an occupation. Proposition 21.2 proves this invariant from the actual record data. It does not identify a failed estimating vector current with the original ordinary matrix current; only the identities already proved for the chosen labelled primitive are used.

A chosen marked primitive need not be intrinsic to the physical quotient. For example two valid first-current primitives of a second derivative can differ under a nonconstant spatial label. The representation retains its own labelled identities, while physical equality uses their unmarked value and total current. Only complete source expressions are assigned a reference center. When a cheap legal representation is used in a coupling argument, its original mark means the raw finite-current primitive obtained by unfolding that representation. All its estimating replacements preserve that same labelled identity. A merely unmarked physical equality with a different primitive is never substituted in a nonconstant-label test. Simultaneous common-scalar capture is used only for the original ordinary symmetric field and its complete raw source family. A directed estimating filling is tested one scalar distance at a time, using that row's effective captured cost and the common all-fillings lower and upper bounds from the paired reset; no vector reference coupling of directed fillings is asserted.

In particular, a selected accepted or failed labelled piece is not defined as an operator on the quotient. Its estimate is an estimate for the chosen representation and its retained label. The physical transfer is the sum of its complete history family. A zero physical expression has the zero representation available, and every bound is proved for all legal representations before their infimum is taken. No future operation can read a chosen internal label as an additional physical observable.

Lemma 21.1 (Congruence of complete expressions)**.** *Physical equality is preserved by composition with a finite physical source expression, by positive owner reassignment, and by complete history partition. Consequently finite source products are well-defined on the physical quotient.*

Proof. Let P have zero action on every admitted test Ψ throughout its mean family. For each Ψ the function

$$u \longmapsto \mathbb{E}[P\Psi(F+u)]$$

is identically zero. Differentiating this identity in finitely many physical directions proves the assertion for any finite constant-coefficient differential polynomial. A finite-rank covariance coefficient is such a polynomial. The compact-source bounds on the original common contours permit passage to the actual feature limit and hence to every finite covariance coefficient.

For a conditional identity, first condition on the exterior. The identity holds for every exterior graph and every value of its edge weights, so it remains true when these are the actual exterior data. Integrating gives the zero global mean family above. A later source meeting both the old body and its exterior is therefore covered by the same differentiation argument. Positive assignment and complete history summation preserve equality by linearity. The argument would not apply to an independently centered selected history. □

The persistent labelled primitive has a separate algebraic justification. Before any history partition, write a finite physical coefficient, with its chosen first leg, as

$$P = \sum_{\nu} D_{v_\nu} Q_\nu, \qquad P^{[w]} = \sum_{\nu} D_{wv_\nu} Q_\nu \,,$$

where Q_ν are constant-coefficient physical operators; a finite feature integral may replace the sum. For a joint anchor/output label, wv_ν means $x \mapsto w(z_\nu, x)v_\nu(x)$ with that original anchor z_ν retained. A covariance coefficient is written this way by resolving one covariance leg, and a mean coefficient by resolving its first mean direction. For every new constant-coefficient source polynomial S,

$$SP^{[w]} = \sum_{\nu} D_{wv_\nu}(SQ_\nu)\,.$$

This is the same primary and the same original label. Physical mean differentiation of the full estimating identity gives precisely this composition, including every contact with its other coefficients. After a conversion the differentiation acts on the stored full-mean coefficient family; the fresh body and physical direction basis are held fixed, and no new source is applied to its fresh variables. Other disjoint children may be used in their value identities because the protected primary range annihilates their local coefficients and readers. The compact-source estimates pass this calculation through feature limits, and the positive envelope passes it to the retained C^1 labels. A centering drift is a separate complete summand with its own primitive. This proves that finite induction preserves the chosen labelled identity, without asserting equality of labels for two merely unmarked-equivalent operators.

We retain the two stages of a representation. At scale r, a representation before mean insertion has every live source birth at scale at most r/b_0. A representation after mean insertion may additionally have one aggregated scale-r mean coordinate per physical cell. Its scale-r covariance coordinate is still absent. An opaque converted child carries the scale and list of consumed source slots of its actual conversion. A contour is never renewed by rewriting a direction or by hiding a newly occupied coordinate in an opaque child.

For a complex coherent innovation $\varphi = x + iy$, put

$$W_j(\varphi) = \alpha_j \|x\|^2 + \beta_j \|y\|^2 + \gamma_j \|x\|, \qquad j \in \{c, r\},$$

where the converted weight W_c is strictly smaller than the unintegrated weight W_r.

21.1 Order of constants for the weighted reset

The constants are chosen in the following order. First fix the macroscopic projection geometry, complete-carrier conventions, reference delay, Gaussian overlap constants, and all ordinary and transferred moment orders. The surrounding-loop and separating-crossing moments used here hold for every $a > 0$: apply [3, Proposition 2.9] to fixed separated connected compact sets that are not singletons, then use a finite crossing network. No point-distance first moment is used.

On an initial absolute contour neighborhood, with every residual covariance at least half the base covariance, the original imaginary likelihood has a quadratic bound $\beta_0 \|y\|^2$ uniform in the subsequently smaller contours. Choose $\beta_c > 4\beta_0$ and $\alpha_c > 0$. Choose the raw weights strictly larger than the converted weights, leaving fixed numerical margins for the joint Gaussian comparison and the spoiled-trial estimate. The converted linear weight may be taken to be zero; any retained linear real term is absorbed into a reserved positive real quadratic margin at a finite multiplicative cost.

Choose the stored failure rate and then the shifted-trial tolerance so its quadratic loss is smaller than that reserved real margin. Fix the local raw failure target and perform the finite numerical source exhaustions, including the middle-score cutoff, the common forecast cutoff, fresh residual probability and density lower bounds. These choices do not enlarge the predetermined observation block. Now shrink the original source contours so the joint Gaussian quadratic exponential with the chosen raw and converted weights is integrable. This is possible because the relevant joint covariance has operator norm $O(\rho)$ and trace $O(\rho k)$. Fix C_* large enough for the resulting one-record moment bound and its separately reserved terminal bound. No contour is renewed later.

Only now choose the minimum accepted age k_0. The coherent restriction to the actual complete old carrier tends to zero with age. Increase k_0 so that

$$C\alpha_r \|L_k\|^2 + \delta_0 < \alpha_c/2, \qquad C\beta_r \|L_k\|^2 + \beta_0 < \beta_c/2 \quad (k \geq k_0),$$

and so that every physical clearance in the local reset holds. The real middle-score term is absorbed in a fixed part of the remaining real margin. This proves the strict estimate using the same companion and fiber denominator as the bounded-real-mean estimate. The local evaluation induction is now performed with these fixed constants. Initial records have the independent raw Gaussian envelope; a converted child can be introduced only after its actual paired construction and pointwise fiber bound.

21.2 Legal records

A converted record has an actual conversion scale s, fixed buffered domains $H_0 \Subset H_c$ of that scale, and the declaration that all sources consumed at conversion were born at scales at most s/b_0. It has a real companion kernel $\mathsf{Q}_e(d\xi)$, including its retained middle variables, fixed positive directed body maps $D_\tau(e, \xi)$, and compact physical private directions $g_{\tau l}(e, \xi)$. They are the objects constructed by the fiberwise local-reset theorem. The kernel is dominated by the corresponding augmented

real physical law with the stated finite constant, which is included in the record's cost. No later source likelihood is applied to its newly introduced coordinates.

For every admitted exterior graph and every full physical mean v, set $\varphi = Q_{H_c} v$. The record has paired value and current identities, with all exterior harmonic mean factors kept in their original Gaussian block:

$$\mathcal{I}(P; f, w) = \int \sum_{\tau} \langle h_\tau(\varphi, e, \xi), w \rangle f(D_\tau)\, d\mathbb{Q}_e = \int \sum_{\tau, l} \langle u_{\tau l}(\varphi, e, \xi), w \rangle D_{g_{\tau l}} f(D_\tau)\, d\mathbb{Q}_e \,.$$

For $w = 1$ this is the unmarked physical action; for other labels it is the record's chosen original-output primitive. The densities take values in $\mathcal{B}_r(K)$ and obey, on each retained fresh residual/filling fiber,

$$\sum_{\tau} \|h_\tau(\varphi, e, \xi)\|_{\mathcal{B}_r} + \sum_{\tau, l} \|u_{\tau l}(\varphi, e, \xi)\|_{\mathcal{B}_r} \|g_{\tau l}\|_\infty \le A e^{W_c(\varphi)} \,.$$

The value measure has total mass zero on each such fiber for every shared mean. The companion, body maps, physical direction basis and original labels are independent of φ. They are jointly measurable in the retained exterior; only the coefficient family depends holomorphically on the incoming mean. The record is zero outside its retained eligibility indicator, with that indicator included in the full physical mean identity. After closing the module, no later selector may read its internal comparison or companion variables.

An unintegrated record consists of one finite original joint Gaussian coefficient integral, its original physical source identifiers and fixed coefficient contours, its maximal converted records, and its finite tree of unintegrated observations and read owner addresses. The protected primary is taken from a distinguished older input factor, chosen before expanding that factor's estimating histories. Its same original mark is followed in every branch. If that distinguished input is represented by a converted child, use its stored current; if its branch is raw, use its protected original compact source range even when another input happens to be converted. The choice is never changed merely because a selected history contains a different converted child. Thus unfolding the complete history family recovers one original finite-current primitive, with its original output label. All later selectors are outside that whole input and are annihilated by the primary. Other source derivatives may hit older selectors; these contacts remain in the original joint likelihood. Maximal child domains, including their unchanged surrounding regions, are disjoint. After the fixed radial thinning the full observation buffers have bounded overlap. The local failure is the complement of the actual residual-selected atlas event intersected with the entire original-menu window; a coupling forecast is not stored as a failure event.

If the numerical coefficient is c, there are k distinct original slots and q maximal opaque children of costs A_i, and the record retains N_f distinct failed observations and g read address steps, assign it the cost

$$|c| C_*^{k+q} \prod_{i=1}^{q} A_i \, \exp(-\lambda_f N_f + d_0 g) \,. \tag{21.2}$$

The coefficient c includes the original coefficient-contour factors and exact coefficient factorials. For an initial raw record these are the full original value-coefficient contour orders. The favorable factor $1/\alpha_s$ in the protected first-current formula is used only to prove its estimate; it does not reduce the assigned raw cost or replace the cost of the value coefficient by the smaller current estimate. Thus the same declared cost controls both original Gaussian value coefficients and their protected currents. An opaque child counts once, irrespective of the number of variables it consumed. The fixed constant C_* has separate multiplicative reserves for the structural query

estimate and for the single terminal Gaussian-and-routing evaluation. Both reserves are chosen from the fixed microscopic geometry and moment orders before the block length or spatial weights. No estimate may spend the terminal reserve at each intervening age. A genuinely finer unconsumed source is disjoint from that child; every remaining source intersecting it is coarse or in the fixed comparable-scale range of Proposition 3.2. Every retained unintegrated observation buffer is outside the maximal opaque child domains. If a previously read observation lies inside a newly closed child, it belongs to that child's stored numerator and is no longer an exterior trial. This is the geometry used by the joint exterior Gaussian angle estimate, including in the torus gauge. Union retains the union of all original lists, maximal children and read flags, counting each physical item once. It does not resample Gaussian variables, renew a coefficient circle, or introduce a separate occupation for each child.

Legal records are generated by a finite induction. The initial records are finite original source coefficients, including primitive compact mean drifts, with the protected raw current below. The allowed operations are literal union, positive whole-owner reassignment, insertion of an absent physical mean axis, and a complete finite history partition whose readers avoid the protected input. A converted record is introduced only by the proved paired reset at a strictly later admissible query scale; its companion and its pointwise cost are the actual output of that construction. One may then close finite linear combinations and the prescribed feature or auxiliary-phase integrals. Every such record has finite ancestry. The evaluation proof below is induction on this construction, using the already proved contract of its maximal converted children.

Individual selected records are estimating pieces. A legal representation of a complete physical expression is a finite family of these pieces, with their exact physical identity and their complete history sum. Its cost is the sum of their costs. Only these complete families are atoms of the completed physical space; a selected success, failure or conic summand is not added as an independent physical state and receives no separate center. Deterministic feature and auxiliary-phase integrals can be included inside one such atom with their actual total cost.

All choices in a legal estimating piece are made from normalized data: $F - \log U$, the metric divided by U, its ratios, the original coordinate amplitudes, and the retained Euclidean declarations. Thus adding a physical constant c leaves selectors and normalized coefficient functions unchanged and multiplies every physical body cost by e^c. A global torus source likelihood is kept separately; it need not have this invariance. For the converted formulas this assertion also follows directly from their construction. If a typed current on m cost coordinates is ρ, then under scalar multiplication of the metric

$$\rho_c = e^c(\mathrm{dil}_{e^c})_*\rho, \quad \sigma_c(y) = e^{(1-m)c}\sigma(e^{-c}y), \quad \mu_c(y) = e^{-mc}\mu(e^{-c}y), \quad g_c^+(y) = e^{-mc}g^+(e^{-c}y)\,.$$

Hence $h = \mu/g^+$ is invariant at the scaled physical cost, σ/g^+ scales by e^c, and the physical Jacobian right inverse scales by e^{-c}. The lifted physical direction is invariant. Compact pure scores annihilate constants, and products, history partitions, and positive assignments preserve these rules. A raw global torus adjoint has a scalar-precision contribution and is not classified as one of these invariant multipliers: its entire original source likelihood remains separate. Every typed estimating term is therefore degree one through its body and degree zero through its normalized retained multiplier, after its remaining global source likelihood is set aside. This is a homogeneity statement about the exact retained term, not a claim that it is separately a complete source polynomial or has a reference center.

21.3 Joint records and coherent means

The deterministic mean in a child $H_i \subset H$ is $Q_{H_i}\varphi$, with

$$\sum_i \|Q_{H_i}\varphi\|^2_{E(H_i)} \le \|\varphi\|^2_{E(H)} . \tag{21.3}$$

The random source shifts are projections of the same original auxiliary Gaussian variable. They are not resampled in different children. An arbitrary child innovation in $E(H_i)$ is never extended by zero and treated as an element of $E(H)$.

Proposition 21.2 (Uniform evaluation and spatial current bounds)**.** *With the constants chosen as below, a record of cost A in (21.2) has the value and protected first-current estimates*

$$\|\textit{evaluation at } \varphi\| \le CAe^{W_r(\varphi)}$$

in every admitted normalized ordinary query. It uses one whole-frame occupation and one original output label. Union retains the old source and observation lists and charges only newly occupied slots or newly opaque inputs. Insertion of an absent aggregated mean coordinate maps the before-mean class to the after-mean class with its original coefficient-contour cost.

Every finite legal representation also satisfies (21.1). In a locally rooted query its restored spatial carrier is a fixed enlargement of the complete declared owner, including all current supports, retained readers, and restoration regions. Below every eligible GOOD ancestor, each prescribed carrier envelope has the uniform ordinary conditional moment required by Proposition 14.1. These statements hold throughout the same full mean family.

Proof. We prove the value, current, and spatial assertions together by induction on the finite construction depth. All body maps below are actual real physical bodies; complex means and source contours occur only in the retained coefficients and the one original Gaussian likelihood.

First fix one reconstructed physical field and one actual body with port costs $(K_\ell)_\ell$. At common differentiability points for the finite physical directions, a direction v gives

$$|D_v K_\ell| \le \|v\|_\infty \nu_{K_\ell}(\operatorname{supp} v). \tag{21.4}$$

Consequently, for every bounded covector $H = (H_\ell)$ of the entire cost vector,

$$\left|\sum_\ell H_\ell(K) D_v K_\ell\right| \le \sup_\ell |H_\ell(K)| \, \|v\|_\infty \sum_\ell \nu_{K_\ell}(\operatorname{supp} v).$$

The same inequality holds with the retained dual output norm. Thus the estimate is made before cost pushforward and does not require a scalar test representation of H. Finite sums of directions and types are estimated by summing their coefficient norms first. The absolutely continuous retained amplitudes, or an auxiliary compact Cameron–Martin shift, give a common full-measure differentiability set for the finite frame. The physical first-variation statement and its infimum version in Lemma 20.4 justify this calculation for the actual bodies used here.

For an initial raw record, the positive joint Gaussian envelope in Lemma 3.3 dominates the absolute coefficients of its protected directions. Multiplying it by the positive spatial currents in (21.4) gives a single dual-norm envelope, independent of the matrix covector. It is supported in the declared source carrier. This is stronger than a bound on a single scalar pairing.

For a converted current child, its actual stored bound gives on every retained fresh residual/filling fiber

$$\sum_{\tau,a} \|u_{\tau a}(\varphi)\|_{\mathcal{B}_r} \|g_{\tau a}\|_\infty \le A_i e^{W_c(\varphi)}.$$

Apply Lemma 20.4 separately to each port coordinate of the actual glued body in a locally rooted query, using the same ordinary exterior wiring, one containing modification region C, and its single surrounding loop and separating crossing. Sum the displayed coefficient norms before taking a Gaussian or occupation moment. A common ratio $\Theta_C = U_C/L_C$ then bounds the sum of all restored current components by the same ordinary labelled frame on C^+. Only the fixed finite port multiplicity is paid. In a port-only query use its physical cap and whole-frame lower bound instead, distributing over the whole labelled frame as described below; no coordinatewise lower bound is inferred from a bound on the whole sum. The companion's density with respect to the original augmented real law is restored at its declared cost. All these steps are positive estimates, valid on every retained filling on which the fixed geometric contract is used; no covariance or mean change is placed on the fresh field.

For a union, use exactly one primary current. If it belongs to a converted child, take that child's stored current and every other child's value representation. If it is raw, retain its protected compact direction and take all closed children in their value representations. The protection rule makes that physical derivative annihilate their coefficients and all retained readers. The extension of Lemma 20.4 above covers this raw case as well. After gluing, there is one actual allowed-path body, one containing modification region and one ordinary occupation; the other children contribute only their positive coefficient envelopes. Their type sums factor into the stored pointwise child bounds. In particular no product of child occupations or of child loop ratios appears. This proves the spatial domination step of the construction induction.

Here is an explicit choice of restored spatial lift. For each fixed physical field and each coordinate, a scalar current whose absolute value is bounded by $B\nu_{K_\ell}(C^+)$ can be distributed against $\nu_{K_\ell}|_{C^+}$, dividing by $\nu_{K_\ell}(C^+)$ when it is positive and assigning zero when it vanishes. Its total cost-coordinate current and its original row label are unchanged. In the port-only case the same construction uses $\sum_\ell \nu_{K_\ell}$ and the pointwise zero-amplitude lower bound, with its Gaussian inverse factor from Lemma 20.4; keep the cost index j as well as the distinct ordinary row label ℓ in that lift. Its finite cost-coordinate multiplicity is paid once. Such redistribution is not used in the locally rooted coupling calculation. Apply this to the finite coefficient-valued direction sum before forgetting the actual body data. Integrating the original augmented variables and restoring their real density gives the lift in (21.1). Conditional expectation contracts its positive ordinary moment. For the retained dual space, perform this construction first on a countable dense subset of the separable C^1 label space; the common domination extends it to the whole unit ball. This avoids assuming a vector Radon–Nikodym property for that dual space.

The following joint calculation is applied to this positive spatial envelope, taking absolute coefficient norms before any test is inserted. Below a prescribed eligible GOOD ancestor all complete old readers and restoration regions are in its fresh protected disk. Use the same original residual law, the same source contours and the same joint Gaussian auxiliary variables. The conditional Gaussian bounds, the local loop/crossing moments and the whole-frame lower normalizer give the same fixed higher ordinary moment for each prescribed carrier. These bounds are uniform in its position and the ancestor radius: scalar metric normalizations cancel from Θ_C, and the retained centered harmonic prediction is controlled by the stated GOOD query bounds. No conditioning on the future event that this ancestor is the last GOOD one is made. In particular the ordinary calculation is available before selecting the carrier containing the actual occupied root.

The full original likelihood disintegrates into the exterior and child likelihoods at the same shared mean. After integrating each child through its full-mean identity, its argument has the form

$$Q_{H_i}\varphi + w_i + Z_i\,.$$

Proposition 3.2 gives the joint direct-and-query bounds

$$\sum_i \|w_i\|^2 \le C\rho^2 k, \qquad \|\mathrm{Cov}(\Re Z, \Im Z)\|_{\mathrm{op}} \le C\rho, \qquad \mathrm{Tr}\,\mathrm{Cov}(\Re Z, \Im Z) \le C\rho k\,.$$

Apply (21.3) and Young's inequality to the sum of the converted weights. For each fixed small $\eta, \varepsilon > 0$, its deterministic part is at most

$$((1+\eta)\alpha_c + \varepsilon)\|x\|^2 + (1+\eta)\beta_c\|y\|^2 + C_{\eta,\varepsilon} q\,,$$

plus a fixed quadratic form of $\Re(w+Z), \Im(w+Z)$. The Gaussian determinant formula bounds the exponential moment of that form by $e^{C(k+q)}$ when the original contours are sufficiently small. One exterior likelihood handles the remaining original Gaussian variables. The compact-source direct estimate bounds it in the same fixed higher moment. This proves an ordinary joint envelope bound $C_p^{k+q}\prod_i A_i$.

The source and restoration calculation has so far been under an ordinary conditional law. Retain the failed observations in that law. Proposition 4.4, applied to the same original residual fields and whole original menus, gives

$$P_e\left(\bigcap_v \Gamma_v^c\right) \le Ce^{-\tau N_f}$$

at any prescribed fixed rate τ, after choosing the ordinary numerical cutoffs. Form the single positive envelope

$$W = e^{\lambda_f N_f}\mathbf{1}_{\cap_v \Gamma_v^c} V\,.$$

Fix ordinary and transferred exponents $p_{\mathrm{ord}} > p_0 > 1$ among the original moment reserves. Apply one ordinary Hölder inequality to this event and the joint envelope V, using its fixed higher moment. Taking τ larger than the stored rate after this loss gives an ordinary $L^{p_{\mathrm{ord}}}$ bound $C_{p_{\mathrm{ord}}}^{k+q}\prod_i A_i$ for W. The old observations and their indicators are still present in this common envelope; none has been separately conditioned on or regenerated.

For an ordinary finite whole-frame query, direct Hölder with $L_\Sigma/d(e)$ and its fixed moment proves the required occupation estimate from this ordinary envelope. This includes roots on compulsory exterior arms; no reference crossing-count assertion is needed at such a root. For an actual pure-reference query, or a GOOD-started query rooted in the strict locally minimizing buffer of Corollary 10.4, apply Proposition 14.1 to W. The complete old owner contains every fresh arm and unchanged restoration region, so the restoration carrier remains within a fixed multiple of its declared diameter from the original current carrier. The same carrier catalogue and backward-address count therefore apply even when the stored primary has been converted. The ordinary exponent p_{ord} leaves the exponent p_0 after the single Campbell-density comparison: Hölder uses the density moment $p_{\mathrm{ord}}/(p_{\mathrm{ord}} - p_0)$. These moment orders were fixed before the cutoffs and C_*. The last-GOOD estimate has the form

$$\mathbb{P}\{W > t, G = \ell\} \le \min\{Ce^{-\alpha\ell}, Ce^{\beta\ell}t^{-p_0}\}\,,$$

after normalizing its ordinary moment. It yields every rooted moment strictly below $p_0\alpha/(\alpha+\beta)$. Choose the original ordinary moment and the GOOD cutoff so that the final required exponent lies

strictly below this threshold. This fixes one spare rooted moment throughout the construction. The same argument handles the initial no-GOOD term using the whole lower normalizer and its initial loop moment. The factor C_*^{k+q} pays the Gaussian and restoration moments once, while $e^{-\lambda_f N_f}$ is retained in the assigned cost. Read backward addresses have bounded degree and are paid by $e^{d_0 g}$; unread assignments retain positive total mass one. This order of proof never combines a rooted moment with a failure probability under a different ordinary law.

For a real mean x, the centered change in an observation of radius r_v inside a radius-R query is at most $C(r_v/R)\|x\|$. Bounded overlap gives $\sum_v r_v^2 \le CR^2$. At most $C\|x\|^2/M^2$ trials therefore exceed a tolerance M. Keeping the stored history factor and using a larger available failure rate on the unspoiled trials loses at most

$$\exp(C\lambda_f \|x\|^2/M^2)\,.$$

Choose M after λ_f, before the ordinary cutoff; this quadratic coefficient is as small as desired. The shifted whole-frame normalizer changes by a fixed linear Weyl factor. For an imaginary mean, the original Gaussian likelihood contributes its fixed quadratic factor; no new likelihood is placed on a fresh companion. The previously reserved raw quadratic margins dominate these coefficients; the remaining linear factor is absorbed in the reserved real quadratic margin.

The same proof with all old contours retained permits an absent physical mean axis on its fixed contour. Coefficient extraction includes $1/m!$, so the cost of $D_\psi^m/m!$ is at most C^m. The source-axis absence must hold in the whole input, including opaque source lists. All estimates concern one evaluation of the flattened record; union does not recharge its old lists.

Every remaining legal operation preserves the invariant. A positive whole-owner reassignment multiplies the same measure by its scalar weight and retains its total mass bound. A complete finite history partition retains all read flags; its selected pieces use the positive envelope and the stored failure gain just described. Such flags cannot inspect or split the protected payload. Inserting an absent physical mean axis keeps the original coefficient contour and its $1/m!$ factor; the same positive likelihood calculation costs C^m. It changes no fresh companion or physical direction. Finite sums and the prescribed feature or auxiliary-phase integrals preserve the domination by the triangle inequality and Minkowski's inequality with their actual costs. A bounded multiplier of the actual whole cost vector, a restriction to an exterior failure event, and a measurable typed forest pushforward preserve variation domination. The finite forest coefficients cost their fixed sum of absolute values. None of these operations asserts a new equality of raw and estimating matrix currents.

Conditional reconstruction uses Lemma 4.2 with the single original likelihood. The body maps, physical bases, and spatial containing regions stay fixed throughout the mean family. □

Lemma 21.3 (A translated array at one occupied root)**.** *Use the spatial-current domination proved in Lemma 21.2, as part of the induction of Proposition 21.2. Let one complete owner prototype have n cells at mesh r, legal cost A, and complete restoration carrier of radius at most Cnr. Translate this same prototype over its fine lattice. Use a locally rooted ordinary frame whose original row terminals lie outside every complete restoration carrier under consideration, or a pure-reference capture with this same clearance. Retain the original row label and use the same-row loop/crossing restoration. This is the query class in the strict inner buffer of the paired reset and in the reference capture; a generic port-only frame does not imply this property. In this class the sum of the translated absolute currents, with any coefficients in $[0,1]$, has one common positive rooted envelope of L^p norm at most Cn^2A for a fixed $p>1$. The bound is independent of the ratio between the outer query radius and r. It also bounds tests chosen separately for the translated*

records, provided their absolute values are at most one and the spatial domination applies before those tests are chosen.

Proof. The spatial invariant gives, after division by each record's assigned cost and the stipulated mean weight, a nonnegative envelope $V_c(F, \ell, x)$ supported on a fixed enlargement D_c of its complete carrier. It dominates the retained-label current as a measure against the same ordinary occupation. In particular it is independent of the bounded cost covector subsequently tested. Color the lattice into $O(n^2)$ residue classes so that the sets D_c in one class have uniformly bounded overlap. At an occupied point, put

$$V(F, \ell, x) = \sum_{c:\, x \in D_c} V_c(F, \ell, x)$$

for one class, with the original row ℓ unchanged. This is a bounded number of locally selected envelopes, not the sum of their global L^p norms.

We verify its rooted moment by the argument of Proposition 14.1. If the last GOOD ancestor is j levels above the local carrier scale, at most Ce^{dj} carriers can occur below that ancestor. Conditional on its state, the proved ordinary query estimate for each prescribed carrier and the single Campbell-density Hölder inequality give a uniform p_0th moment. Drop the subsequent last-GOOD restriction before applying that estimate. The union bound over the candidate carriers, followed by averaging the ancestor state under its probability law, gives

$$\mathbb{P}\{V > t, G = j\} \le Ce^{dj} t^{-p_0}.$$

No sum over the possible positions of that random ancestor is made. The GOOD-count bound, applied to the same root marginal, gives

$$\mathbb{P}\{V > t, G = j\} \le \min\{Ce^{-aj}, Ce^{dj} t^{-p_0}\}.$$

The initial no-GOOD branch has the same estimate using its initial loop moment and whole-frame lower normalizer, as in Proposition 14.1. Summing at $j = p_0 \log t/(a + d)$ proves every moment $p < p_0 a/(a + d)$. The moment orders and GOOD rate were chosen to leave one such $p > 1$. These constants depend on the fixed local carrier geometry, not on the number of fine cells in the outer query.

Minkowski's inequality over the $O(n^2)$ residue classes proves the claim. Positive coefficients at most one can only decrease the dominating spatial measure. An unread positive assignment prefix has total mass one, so averaging it also preserves this bound. □

21.4 A strict accepted estimate

We explain why the complex mean norm does not require an exponential mixing rate. The geometry, mean weights, numerical cutoffs, original contours, and C_* are now fixed in the stated order. Choose one finite minimum accepted age k_0. It is large enough that $R = b_0^k r$ contains all required private-clearance buffers and that the strict coherent restriction below holds for every $k \ge k_0$. This finite delay is relative to the input scale and is the same at every absolute depth. No new trial is read at an age $k < k_0$, and no artificial failure factor is charged there. Let P be a centered complete input on an owner of at most N_* cells. Below $T_k P$ denotes the accepted converted representation obtained from a chosen legal representation of P; its norm is its proved pointwise density/current cost. It is not asserted to be an intrinsic nonconstant-labelled operator on the physical quotient. The bounds are uniform in the choice of legal representation and will be used only after the complete last-success family is restored. For each fine translate assigned to a

new cell at age k, its actual old carrier has radius $O(N_* r)$, while the query radius is $R = b_0^k r$. The coherent restriction bound is

$$\|Q_{H_{\mathrm{old}}}\varphi\|_{E(H_{\mathrm{old}})} \le C_{N_*}(r/R)^{\varepsilon/2}\|\varphi\|_{E(H)}\,.$$

Apply the local reset separately to these old carriers, with the same retained exterior, same fresh companion and directions, and same density denominator. Do not replace their union by one old carrier of radius R.

For clarity, positive reassignment at age k gives the exact expression $\sum_c K_{b_0^k}(c,Q)P_c$. On the old carrier of P_c, the smooth refinement estimate gives

$$|K_{b_0^k}(c,Q) - \psi_{R,Q}(x)| \le C_{N_*} r/R\,.$$

The complete fine drift masks refine exactly to $\psi_{R,Q}$. Thus the difference from the smoothly localized complete centered current costs $C_{N_*} r/R$ against its summed absolute current. The following estimate controls that sum without counting the fine cells.

Apply Lemma 21.3 to each complete input prototype before testing the array or dividing by the single frame normalizer. It bounds the summed translated absolute current by $CN_*^2 A$ when $n \le N_*$. The positive coefficients $K_{b_0^k}(c,Q)$ and the unread assignment prefixes preserve the same envelope. Consequently the refinement error is at most $C_{N_*}(r/R)A$, and neither that error nor the bounded-real-mean reset constant contains a factor b_0^{2k}.

Here is the global mean estimate using that same companion. Choose buffered deterministic cuts

$$H_{\mathrm{old}} \Subset H_0 \Subset H_c\,,$$

where the old payload, observations and original retained menus are in H_{old}, and H_0, H_c have radii comparable to R. Use Proposition 4.1 to integrate unread original menus before making these cuts. The remaining reserve ranges lie on one side of each cut, so the residual covariance commutes with the corresponding Dirichlet projections.

Fix the new retained datum e_0 outside H_0, including the original nonselected coordinates and the new atlas geometry. Orthogonal Gaussian disintegration factors the mean likelihood into the H_0 innovation and its middle complement. The middle likelihood stays at this fixed e_0. Its score is a Gaussian linear functional on $E(H_c)$. The map from this coherent space into the physical Cameron–Martin space is Hilbert–Schmidt: its own-domain term controls H^{s_0}, and the embedding $H^{s_0} \to H^1$ on a fixed two-dimensional domain is Hilbert–Schmidt for $s_0 > 2$, since

$$\sum_{k\in\mathbb{Z}^2} (1+|k|^2)^{1-s_0} < \infty\,.$$

Bounded extension, a fixed cutoff, and Dirichlet projection give this assertion without imposing higher-order boundary compatibility. Composing with the middle projection and the residual inverse square root preserves the Hilbert–Schmidt bound, since the residual covariance is at least half the base covariance. The dual score thus has uniform covariance trace and operator norm and a small quadratic exponential moment. Include a fixed large cutoff on its norm among the actual local success predicates. This cutoff has arbitrarily high probability and is measurable from the fixed enlarged middle annulus.

On this event the real middle score is bounded by $CM\|x\|$. Its negative real quadratic may be discarded, and its imaginary modulus contributes at most $\beta_0\|y\|^2$.

For precision, perform the real translation on the complete finite Gaussian numerator before coefficient extraction or absolute values. After the unread original menus have been removed, the

reference residual covariance commutes with the H_0 Dirichlet projection. Write its inner variable as X_0, its fixed exterior as e_0, and $u = Q_{H_0}x$. For a finite covariance-source parameter $\mathcal{K}$, the Gaussian density identity is

$$L_{\mathcal{K},v}(X_0+u,e_0)\,d\gamma_0(X_0+u) = L_{\mathcal{K},v-u}(X_0,e_0)\,d\gamma_0(X_0).$$

Here the same exterior Gaussian density occurs on both sides and cancels. The formula is just translation of the numerator Gaussian density of mean v and covariance $C+\mathcal{K}$; it holds on the original complex covariance contours as well. Hence for the complete retained integrand J,

$$\int L_{\mathcal{K},v}(X_0,e_0)J(X_0,e_0)\,d\gamma_0 = \int L_{\mathcal{K},v-u}(X_0,e_0)J(X_0+u,e_0)\,d\gamma_0.$$

Every old selector, coefficient, and ordinary metric inside H_0 is translated inside J. None is held fixed merely because it was previously observed. The new atlas, its eligibility event, the selected outer arm variables, and the fresh companion are outside this translation and stay fixed. Coefficient extraction retains all contacts in this identity.

The old innovative means are $Q_{H_{\mathrm{old}}}u = Q_{H_{\mathrm{old}}}x$, with the strict coherent restriction already displayed. The remaining change is harmonic on each old payload. Compact pure scores annihilate that harmonic part; the normalized old flags retain their actual shifted values and are estimated by the spoiled-trial bound, and the ordinary frame is compared by its one linear Weyl factor. The one original imaginary likelihood remains in the same integral. These are precisely the factors used in the following real and imaginary quadratic estimates.

The support used for division is unchanged as well. With e_0 fixed, the old cap intervals and outer arm box contain the cost vectors almost surely under every fixed real Cameron–Martin translate of the reconstructed old law, by the finite-genealogy argument giving (19.1). This applies to X_0+u, even if some old selectors change value. The integrated numerator is therefore supported in the same old convolution support as before translation. Lemma 19.2 gives the lower bound for the same fresh density on that entire support, on each retained fresh fiber. Division consequently produces the same holomorphic full-mean conversion, with body maps, companion, physical directions, and labels independent of the mean.

The imaginary H_0 mean stays in its one original likelihood; its modulus gives a fixed quadratic term, included in β_0. The old converted child weights use $L\varphi$ and their joint Gaussian comparison variables, already bounded once in Proposition 21.2. Thus the integrated numerator has real and imaginary quadratic types at most

$$C\alpha_r\|L\|^2+\delta_0, \qquad C\beta_r\|L\|^2+\beta_0\,,$$

where δ_0 is arbitrarily small. The earlier choices give $\beta_c > 4\beta_0$, $\alpha_c > 0$, larger raw weights, and a shifted-trial loss within the reserved real margin. Increase the finite minimum accepted age k_0 so that the strict restriction on the actual old carrier makes both displayed quantities less than half the respective converted coefficients. This does not change b_0, the source contours, or any already chosen local cutoff. Absorb the remaining real linear term into a fixed part of the strict real quadratic margin.

Integrating the complete old numerator and dividing by this same fresh density gives

$$\|T_kP(\varphi)\| \le MA\exp\{W_c(\varphi)-\delta\|\varphi\|^2\}, \qquad \delta>0\,, \tag{21.5}$$

with constants independent of old order, history length and accepted age. The companion, body maps, physical direction basis and original labels are fixed independently of φ. The strict restriction is to the actual old carrier; projecting merely onto H_0 would not supply this margin.

On each bounded real mean ball, pure-reference forgetting and the output-label freezing error give

$$\sup_{\|x\|\le S} \|T_k P(x)\| \le \varepsilon_k(S) A, \qquad \varepsilon_k(S) \longrightarrow 0 \quad \text{for each fixed } S\,.$$

These are the same numerator, companion, and density division as in (21.5). Outside a large complex ball, the weighted bound is small by its strict quadratic margin. Inside, pair with a unit dual functional and apply the two-constants theorem to $T_k P(x+wy)$ on $\{|w|<3, \Im w>0\}$ at $w=i$. The diameter lies in a bounded real mean ball and the semicircle has the uniform bound (21.5). Its harmonic measure gives a fixed positive power of $\varepsilon_k(4S)$. Consequently the accepted weighted operator bounds satisfy

$$\sup_{k\ge k_0} a_k < \infty, \qquad a_k \longrightarrow 0\,. \tag{21.6}$$

The original output label is an element of the dual of the separable C^1 test space at the current scale. The common positive envelope of Lemma 3.3, followed by integration and the local conic estimate, dominates all these scalar pairings by one positive measure. Scalar Radon–Nikodym on a countable dense test subspace therefore gives weak-* measurable dual-valued densities with the stated norm bound. No Radon–Nikodym property of the whole dual space is required. Local bounded weak-* holomorphy gives norm-convergent Taylor coefficients on smaller disks by scalar Cauchy estimates on that dense subspace.

21.5 The atomic and spatial spaces

For each owner sector and stage, take the infimum of all total legal estimating costs. Lemma 21.1 makes this a seminorm on complete physical expressions. Quotient by zero cost and complete. Equivalently, use the weighted ℓ^1 space of legal atoms and quotient by the closure of its finite physical nullspace. Internal feature and phase integrals remain part of one legal atom.

Put $\Phi(n) = n\mathbf{1}_{\{n>N_*\}}$. The translation-covariant spatial space is the weighted ℓ^1 sum of the owner sectors, with norm

$$\|u\|_A = \sup_c \sum_{X\ni c} e^{A\Phi(|X|)} a_X(u_X)\,.$$

For a relative owner orbit this equals $|X| e^{A\Phi(|X|)} a_X(u_X)$. Finite complete families are dense. Every subsequent operation is first defined on complete finite expressions, estimated for arbitrary legal representations, and then extended after taking the separate input infima.

Theorem 7.1 defines the center by an unmarked scalar capture action and proves that its difference from any chosen original marked primitive tends to zero as the capture radius tends to infinity. The cheap-representation bound is applied to that scalar action before the limit, and the cost infimum is taken afterwards. Hence the center annihilates the physical nullspace, independently of all nonconstant-label choices. Together with Corollary 14.2, this gives a linear center satisfying

$$|\Pi_X(P)| \le C|X|^2 a_X(P), \qquad \Pi_X(P - \Pi_X(P) D_{\psi_c}) = 0\,.$$

Its complete-orbit construction makes it independent of the marked primitive and estimating representation. Its lift $\mathcal{C}$ preserves the owner and takes values in the after-mean space. For every $\delta>0$,

$$\|\mathcal{C}u\|_{A+\delta/2} \le C_{\delta,N_*} \|u\|_{A+\delta}\,, \tag{21.7}$$

since $\sup_n n^2 e^{-\delta\Phi(n)/2} < \infty$. This uses a genuine spatial reserve. Complete positive spline refinement intertwines the reference lift with zero-source unary transfer: the assembled original currents and complete fine drift orbit are unchanged, and $\sum_c \psi_{r,c} = 1$.

22 Uniform analytic normalization

For a finite set of physical slots and owners, write

$$Z(u) = \sum_{\mathcal{A} \text{ compatible}} \prod_{X \in \mathcal{A}} u_X \,,$$

including the empty family. The products are complete physical source expressions. On the plane this notation is coefficientwise: a connected coefficient with a fixed anchor, fixed positive response degree and a finite number of scales has only finitely many relative placements. Identities below mean identities of these locally finite complete coefficients, or their finite restrictions before translation assembly. They do not posit a convergent probability gas or a global metric on the entire plane. The complete mask D_1 is a formal physical constant shift at this stage. Global metric evaluation is used only after actual finite-depth torus periodization, where each physical slot is finite and the constant shift is an ordinary Weyl operation.

At a microstage whose input owner mesh is s, insert the original band/cell germ for the physical band $[s/b_0, s]$, namely $\exp(zD^2_{K_s}/2) - 1$, and then coarsen the whole owner to mesh $b_0 s$. Thus every live source birth in the before-mean output is at most its mesh divided by b_0. Expand before grouping support conflicts. Positive whole-owner assignments and the deterministic connected-component rule define the unnormalized block $S_L(z, u)$, with exact identity

$$Z(S_L(z, u)) = G_L(z) Z(u) \,. \tag{22.1}$$

Here G_L is the product of the original Gaussian factors in the L stages. A block has after-mean input and before-mean output. Atlas partitions change only its estimating representation.

For $v = S_L(z, u)$, write $\beta_X = \Pi_X(v_X)$ and $B_X = \beta_X D_{\psi_{c_X}}$. Expand

$$e^{-\sum_X B_X} Z(v) \,.$$

The counterterms have their original owner declarations and arbitrary multiplicities with exponential factorials. Only the original activity subfamily is subject to the original compatibility constraint. Group all nonempty connected families by their literal union owner; this adds no new halo. Denote the resulting family by $\mathcal{F}(z, u)$.

Proposition 22.1 (Uniform normalization at a fixed base)**.** *There exist a finite microblock length L, a spatial weight A, and fixed positive radii for which $\mathcal{F}$ and its assembled scalar $\mathfrak{b}$ are holomorphic, vanish at $(0, 0)$, and satisfy*

$$\|D_u \mathcal{F}(0, 0)\| \le \tfrac{1}{2}, \tag{22.2}$$

$$\|\mathcal{F}(z, u) - D_z \mathcal{F}(0, 0) z - D_u \mathcal{F}(0, 0) u\|_A \le C(|z| + \|u\|_A)^2, \tag{22.3}$$

$$|\mathfrak{b}(z, u)| \le C(|z| + \|u\|_A) \,. \tag{22.4}$$

The exact coefficientwise identity of complete physical expressions is

$$Z(\mathcal{F}(z, u)) = e^{-\mathfrak{b}(z,u) D_1} G_L(z) Z(u)$$

after complete formal translation assembly, and after actual finite torus assembly whenever globally evaluated. There are $\rho, \varepsilon > 0$ such that the iterates $u_0 = 0$, $u_{j+1}(z) = \mathcal{F}(z, u_j(z))$ are holomorphic on D_ρ and satisfy

$$\sup_j \|u_j(z)\|_A \le C|z|, \qquad \sup_j |b_j(z)| \le C|z|, \qquad b_j(z) = \mathfrak{b}(z, u_{j-1}(z))\,.$$

All maps commute with complex conjugation.

Proof. Each finite term in the product has a unique adjacency-component partition. Copies of a countertype have the same nonempty owner and hence lie in one component. Its factorial is therefore preserved; alternatively label the copies and cancel the allocation multinomial. Conversely a compatible family of connected components recovers the original term exactly once. This proves the finite identity. Its linear term is $(\iota - \mathcal{C}) D_u S_L(0,0)$, where ι is the norm-one pre-to-post inclusion. In particular the quadratic products $B_X^2/2$ and $-B_X v_X$ are retained.

We give the convergent majorant. If nonnegative owner weights $w(X)$ satisfy

$$\sup_c \sum_{X \ni c} w(X) e^{\mu|X|} \le \epsilon\,,$$

and the adjacency neighborhood of X has at most $D|X|$ cells, then the rooted-tree recurrence

$$T(X) = w(X) \exp\left(\sum_{Y \sim X} T(Y) \right)$$

gives $T(X) \le w(X) e^{\mu|X|}$ whenever $D\epsilon \le \mu$. This follows by induction on bounded tree depth. Rooting a connected family at a vertex containing a specified cell and summing its spanning trees gives the same bound, with the exponential multiplicities already included. A common vertex parameter and Cauchy's formula bound the degree h terms by C^h; degrees at least two are $O(\epsilon^2)$.

Fix L temporarily. Unfold the block into its original incoming owners and new active germs. Lemma 15.1 gives

$$\Phi(n_{\mathrm{out}}) \le \theta \sum_i \Phi(n_i) + C_L h, \qquad \theta < 1\,.$$

A fixed-block structural union of arbitrary legal input records costs at most

$$e^{c_L \sum_i n_i} C_L^h \prod_i a_{X_i}\,.$$

This can be obtained without any intermediate reset: carry the original source and observation lists through the fixed positive assignments. Their original reserves are evaluated once by Proposition 21.2. New slots and new opaque inputs contribute the displayed per-input constants. There is no independently chosen merger hierarchy. Each output cell has only finitely many possible input-cell witnesses, with a bound depending on L; the old owner size does not enter this witness count.

Choose separate reserves $\delta_\Pi, \delta_t > 0$, let $\delta = \delta_\Pi + \delta_t$, and take A so large that

$$(A + \delta)\theta + c_L \le A - 2\mu\,.$$

Using $n \le \Phi(n) + N_*$, one reserve in the remaining $e^{-2\mu \sum_i \Phi(n_i)}$ turns $\Phi(n_i)$ into n_i and the other pays the connected-tree sum. Hence

$$\|S_L(z,u)\|_{A+\delta,\mathrm{pre}} \le C_{A,L}(|z| + \|u\|_A)$$

on a fixed small ball, with a quadratic higher-degree estimate. The reference bound (21.7) uses only δ_Π and controls both its drift lift and its assembled scalar.

Identical physical mean slots must now be combined. At a distinguished cell c, set $\beta_c = \sum_{X:c_X=c} \beta_X$. Its factor is $e^{-\beta_c D_{\psi_c}}$. The fixed absent-axis estimate bounds $D^m_{\psi_c}/m!$ by C^m. Expanding β_c^m afterwards gives ordered provenance tuples, all containing c. If the weighted scalar per-site mass is ϵ, their per-site sum is at most

$$\sum_{m\geq 1} mC^m\epsilon^{m-1}\epsilon = O(\epsilon)\,,$$

and the terms $m \geq 2$ are $O(\epsilon^2)$. Distinct new mean axes are estimated jointly, keeping every old contour fixed. This avoids multiplying a Gaussian derivative factorial by a provenance-label factorial. The outgoing literal unions now use the connected-tree estimate with the remaining reserve δ_t. This proves the analytic bound and its quadratic remainder.

These bounds hold for arbitrary legal input representations. Lemma 21.1, followed by the separate cost infima, gives bounded multilinear coefficients on the completed stage spaces. Their series converge locally uniformly by the majorants, hence define the asserted holomorphic maps.

There is also a useful finite-degree statement. For each fixed depth j and response degree m, the coefficient $[z^m]u_j(z)$ is a finite complete physical family. Indeed every nonempty source or countervertex has positive response degree, so a degree-m connected component has at most m immediate positive-degree factors. At the finitely many scales in that prefix, fixed-range connected positions form a finite relative catalogue. Positive reassignment has finitely many anchor choices, and the reference changes only numerical coefficients. Induction on depth and degree proves the assertion; feature and private-phase integrals remain inside the corresponding physical expression. Thus the use of the completion to obtain a uniform analytic bound does not replace the actual finite physical coefficients needed at the terminal.

For the linear term, first take an input owner of at most N_* cells. The inequality $\theta N_* + C_0 \leq N_*$ preserves this class under unary coarsening, and $\Phi = 0$ throughout it. Its spatial weight is therefore one, so the following estimate is independent of A. Reference intertwining allows centering the input with a uniformly bounded cost in this small sector. Along a deterministic owner path, with $L \geq k_0$, use the delayed last-success identity

$$1 = \prod_{j=k_0}^{L} \mathbf{1}_{\Gamma_j^c} + \sum_{k=k_0}^{L} \mathbf{1}_{\Gamma_k} \prod_{j=k+1}^{L} \mathbf{1}_{\Gamma_j^c}\,.$$

No trial at an earlier age is introduced. There are no earlier indicators in the kth term. Sum its unread assignment prefix and delete the unread original menus only now. At its accepted age k, the complete centered expression has converted norm at most a_k by (21.6). All consumed source births are separated from an accepted query by the fixed delay, so this is a before-mean output representation.

A failed suffix of length s has cost at most $C\omega^s$, with $\omega = de^{-\kappa} < 1$, where d is the actual backward-address degree and κ is the stored failure gain after its read-address penalty. Choose the stored rate so that $\kappa > \log d$ after the fixed moment transfer. The one newly opaque input costs a fixed C_* once, not once per suffix level. At a final real mean x, only $O(1 + \log(1 + \|x\|))$ largest suffix annuli can exceed the fixed contrast tolerance, since their contrasts decay geometrically towards the smaller annuli. The loss is polynomial in $1 + \|x\|$ and fits an arbitrarily small real quadratic margin. The one imaginary likelihood is already included in the unintegrated weight.

Thus the small-sector norm is at most

$$C_{k_0}\omega^{L-k_0+1} + \sum_{k=k_0}^{L} a_k\omega^{L-k} \longrightarrow 0\,.$$

The fixed C_{k_0} pays any initial finite address prefix; no such factor is charged at every later age. The array estimate is applied to one relative input prototype at a time. The per-cell family norm then sums their costs; no sum over all fine translates is introduced. For a fixed $K \geq k_0$, its $k \leq K$ terms vanish as $L \to \infty$, and the remaining sum is at most $\sup_{k>K} a_k/(1-\omega)$. Choose one finite L making this part at most $1/8$. The delay changes only this estimating construction; the physical block is still the same fixed L-stage map at every absolute radius. There is no additional age variable in the activity state.

For an input owner of $n > N_*$ cells, use the output centering directly. Its first genuine large coarsening gains cn in the potential, while the fixed block costs at most $C_L e^{c_L n}$ and the output reference contributes at most $C_L(1+n)^2$. The centered weighted bound is consequently

$$C_L(1+n)^2 e^{-(cA-c_L)n}\,.$$

Increase A after L so its supremum is at most $1/8$, preserving the earlier reserve inequality. This also covers a large owner becoming small during the block. Summing the two input sectors proves (22.2).

Finally let $Q = D_u\mathcal{F}(0,0)$ and $H = D_z\mathcal{F}(0,0)$. Choose ε small enough for the quadratic remainder on $|z| + \|u\| \leq 2\varepsilon$, and then $\rho > 0$ so that

$$\tfrac{1}{2}\varepsilon + \|H\|\rho + C(\rho+\varepsilon)^2 \leq \varepsilon\,.$$

Induction gives the invariant ball on the same disk at every depth. Every iterate vanishes at $z = 0$. Cauchy's estimate on a smaller fixed disk gives $\|u_j(z)\| \leq C|z|$ uniformly in j, and (22.4) gives the scalar bound. Real physical coefficients, positive assignments, and the real reference imply conjugation symmetry. □

Lemma 22.2 (Exact coefficientwise representatives)**.** *Let $\mathcal{V}_A$ be the space of finite complete physical families with the estimating seminorm of Section 21, and let $\mathcal{B}_A$ be its separated completion. Suppose $F_j : D_{\rho_0} \to \mathcal{B}_A$ are holomorphic, $F_j(0) = 0$, and $\sup_j \|F_j(z)\|_A \leq M|z|$, where $M > 0$. Suppose each coefficient has a specified finite physical family $P_{j,m} \in \mathcal{V}_A$ whose image is $[z^m]F_j$. There are finite legal representatives $\mathcal{R}_{j,m}$ of these specified physical families, in their same complete owner sectors, such that*

$$\sup_j \sum_{m\geq 1} t^m \operatorname{cost}_A(\mathcal{R}_{j,m}) \leq 4Mt, \qquad 0 \leq t \leq \rho_0/4.$$

The resulting series is holomorphic in the weighted ℓ^1 space of legal atoms. The assertion includes physical coefficients of seminorm zero and requires no linear selection on the quotient.

Proof. Put $R = \rho_0/2$. Since the seminormed space embeds isometrically after separation, Cauchy's estimate gives $\|P_{j,m}\|_A \leq MR^{1-m}$. There are finitely many relative owner sectors in each $P_{j,m}$. In each choose an exact legal representative within a positive error of the defining infimum. Distribute the errors, including their actual per-cell weights, so that their total is at most $MR^{1-m}2^{-m}$. This is possible also when an infimum is zero: exactness is part of the class over which that infimum was taken. Thus

$$\operatorname{cost}_A(\mathcal{R}_{j,m}) \leq 2MR^{1-m}, \qquad \sum_{m\geq 1} t^m \operatorname{cost}_A(\mathcal{R}_{j,m}) \leq \frac{2Mt}{1-t/R} \leq 4Mt.$$

Absolute convergence proves holomorphy in the atomic ℓ^1 space. Coefficientwise exactness proves that its quotient is F_j. □

Lemma 22.3 (Uniform original-record representatives for the terminal)**.** *After a further choice of the spatial weight and a decrease of the fixed response disk, the normalized bulk iterates admit the original-record majorants in* (24.1)*. The same assertion holds for every unreset microstage of the fixed block and for its reference counterterms. These statements hold for actual legal representatives before any further cost infimum. They can be used in the guarded finite-depth terminal identity without changing the original primitive of an already marked record.*

Proof. First finish all geometric and Gaussian choices, including the single terminal reserve in C_*. Choose the finite block length L using the small-owner estimate. Lemma 21.3 makes its constant independent of L; also $\Phi(n) = 0$ throughout the small unary sector, so its choice does not depend on the spatial weight. Fix the tree, pin, and terminal spatial reserves next. Finally enlarge A as needed; this improves the large-owner contraction and preserves the normalization inequalities.

For each of the finitely many prefixes of this fixed block, the pre-infimum construction in Proposition 22.1 bounds its record cost by $e^{c_L \sum_i n_i} C_L^h \prod_i A_i$ and its owner mass by $C_L \sum_i n_i + C_L h$. Choose a target exponential weight A_b large enough for all the previously fixed terminal spatial reserves, and choose $A > C_L A_b + c_L$ with additional strict slack for the tree sum and the n^2 reference loss. The inequalities $n \le \Phi(n) + N_*$ and $\sup_n n^2 e^{-\delta\Phi(n)} < \infty$ show that these same finite construction maps send the weighted atomic input space at weight A to the desired spaces with weight $e^{A_b n}$. This is an estimate on the constructed records; it does not use a quotient norm for their outputs. The rooted-tree and identical-mean-slot majorants already proved in that proposition give locally uniformly convergent series for these maps. Their nonempty outputs vanish at the origin.

Apply Lemma 22.2 to u_j, using the finite-degree physical coefficients supplied by Proposition 22.1. The lifted input series have one $O(t)$ atomic bound, uniformly in j. From these inputs construct all the unreset prefixes and their reference counterterms by the just-described maps, keeping the original coordinate lists, contour orders, and positive assignments of each chosen record. Do not choose independent representatives for these prefixes. The fixed finite number of maps and their uniform analytic bounds give, on one smaller disk,

$$\sup_j \sup_c \sum_{X \ni c} a_j(X;t) e^{A_b |X|} \le Ct$$

for the actual coefficient majorants of each required stage. The n^2 estimate for the scalar center is applied to every chosen complete representative before the infimum, so it also gives the bound for the numerical counterterms and their original owner-preserving drift lifts. No new Gaussian or terminal moment is chosen in this step.

We verify compatibility with the exact terminal expansion at finite depth and degree. A passing complete owner is contained in a strict chart and satisfies the endpoint guards. Every contributing ancestor is contained in that complete owner, and therefore passes the same guards. Consequently restriction to passing outputs commutes with the complete connected-component identity, with numerical counterterms frozen at their planar values. Replace only an unmarked passing bulk coefficient by its chosen exact representative in that same complete owner sector. Lemma 21.1 preserves all subsequent unmarked physical source actions and full mean identities. Guard status and the available genuine exterior routes are unchanged. Generate its microstages from that representative. If an output first fails a guard, mark the resulting actual record at that moment

and choose its protected primitive by unfolding that record. Thereafter retain this primitive and all its labels; no later cheap representative is substituted for that marked component.

Induction over the finite stages proves the exact guarded identity. The replacements need not preserve a common genealogy across different unmarked stages. They preserve exact physical coefficients where unmarked congruence is sufficient, and preserve the chosen genealogy where a marked test requires it. Already bad or noninjective planar records are never newly periodized by this argument. The actual periodic source germs generate those branches before the first mark; only frozen numerical countertypes are periodized without a chart restriction.

Finally the bounded-test extension and the Mellin cap are proved for arbitrary finite legal representatives. The displayed majorants and the marked-forest bound therefore give normal convergence for these specific lifts. Passing to their series preserves the finite coefficient identities. This proves the pre-quotient terminal input; it is not an inference from a bound in the completed quotient alone. □

23 Gaussian identities for the torus observation

The first identity treats the global constant of a torus field. The second retains the contact terms when a Gaussian likelihood is evaluated at a shifted physical argument.

23.1 The exterior quotient

Let the flat torus have area one. In this subsection all pairings and averages use ordinary area measure. Put

$$\mathcal{E}(u,v)=\frac{1}{2\pi}\int_{\mathbb{T}^2}\nabla u\cdot\nabla v,\qquad G_U=2\pi(-\Delta_U)^{-1},\qquad \overline{f}=\int_{\mathbb{T}^2} f\,. \tag{23.1}$$

Here G_U is the zero-boundary Green operator on a proper smooth domain U. Let H_0 be the mean-zero torus GFF and let ζ be an independent centered Gaussian of variance $v_0>0$.

Proposition 23.1 (The scalar tilt)**.** *For proper smooth domains $U_1,\dots,U_k$ with disjoint closures and a nonempty open exterior, there is a decomposition*

$$aH_0=e+\sum_{i=1}^{k}(F_i-\overline{F}_i)\,, \tag{23.2}$$

where F_i are independent zero-boundary GFFs of covariance $a^2G_{U_i}$, extended by zero, and e is independent of them, harmonic in each U_i, and measurable from the exterior field modulo constants. Define

$$v_U=-a^2G_U1,\qquad Z_U=\exp\!\left(\frac{a^2}{2}\langle 1,G_U1\rangle\right)\,. \tag{23.3}$$

If T is a nonnegative measurable observation with $T(F+c)=e^cT(F)$, then, for every deterministic admissible shift u,

$$\mathbb{E}T(aH_0+a\zeta+u)=e^{a^2v_0/2}\left(\prod_i Z_{U_i}\right)\mathbb{E}T\left(e+\sum_i(F_i+v_{U_i})+u\right)\,. \tag{23.4}$$

The equality also holds for signed observations whenever the left side with $|T|$ is finite. If $U'\subset U$ and $Q_{U'}$ denotes Dirichlet-energy projection, then

$$Q_{U'}v_U=v_{U'},\qquad \|v_U\|^2_{\mathrm{CM},a,U}=a^2\langle 1,G_U1\rangle\,. \tag{23.5}$$

For dilations of a fixed smooth shape of radius r, the last quantity is $O(r^4)$ and $\|v_U\|_\infty = O(r^2)$.

Proof. The map $u \mapsto u - \overline{u}$ embeds $H_0^1(U_i)$ isometrically into the torus Dirichlet space modulo constants. The embedded subspaces for disjoint domains are orthogonal. Gaussian orthogonal decomposition gives (23.2). The complementary component pairs to zero with $-\Delta\phi$ for every smooth ϕ compactly supported in a domain, and is therefore harmonic there.

The kernel of restriction of a Cameron–Martin function to the exterior modulo constants consists exactly of those functions which, after subtracting that exterior constant, belong to $\bigoplus_i H_0^1(U_i)$. Gaussian conditioning therefore identifies the complementary component with the exterior quotient. Conditioning instead on the full mean-zero exterior would reveal an additional scalar and would not give these independent cores.

Homogeneity gives the pointwise identity

$$T\left(e + \sum_i (F_i - \overline{F}_i) + a\zeta + u\right) = e^{a\zeta - \sum_i \overline{F}_i} T\left(e + \sum_i F_i + u\right). \tag{23.6}$$

The variance of $\overline{F}_i$ is $a^2\langle 1, G_{U_i} 1\rangle$. Gaussian completion of the square yields

$$\mathbb{E}[e^{-\overline{F}_i} f(F_i)] = Z_{U_i} \mathbb{E}[f(F_i - a^2 G_{U_i} 1)]\,. \tag{23.7}$$

Integrating ζ and then the independent cores proves (23.4), first for nonnegative observations and then by positive and negative parts.

The function v_U solves $-\Delta v_U = -2\pi a^2$ with zero boundary data. Testing against $H_0^1(U')$ proves the projection identity by uniqueness of the weak Dirichlet problem. Testing against v_U proves the norm identity. Rescaling the equation to the fixed shape gives the stated powers of r. □

For a product $T = YK$, the proposition requires that the complete multiplier Y be invariant under constants and that K have degree one. An arbitrary local test does not satisfy this requirement. Differentiating (23.4) in physical directions is justified for smooth cylindrical observations with suitable domination; extension to metric sources requires the separate joint integrability estimates.

23.2 Contacts and unused Gaussian coordinates

For a symmetric complex matrix A with $\|A\| < 1$ and $m \in \mathbb{C}^d$, define

$$L_{A,m}(x) = \det(I + A)^{-1/2} \exp\left(\frac{1}{2} x^{\mathsf{T}} x - \frac{1}{2}(x - m)^{\mathsf{T}} (I + A)^{-1} (x - m)\right), \tag{23.8}$$

using the determinant branch equal to one at $A = 0$. For real A, m with $I + A$ positive definite, this is the density of $N(m, I + A)$ relative to $N(0, I)$.

Lemma 23.2 (A shift of the physical argument)**.** *For every $x, \theta \in \mathbb{C}^d$,*

$$L_{A,m}(x + \theta) = \exp\left(\theta^{\mathsf{T}} m + \frac{1}{2}\theta^{\mathsf{T}} A\theta\right) L_{A,m+A\theta}(x)\,. \tag{23.9}$$

Proof. Expand both quadratic exponents in (23.8). Their coefficients of x agree because $(I + A)^{-1}(m + A\theta) = (I + A)^{-1} m + (I - (I + A)^{-1})\theta$. The constant terms agree by the same identity and symmetry of A. The determinant factors are identical. □

Both the scalar factor and the changed mean $m + A\theta$ must be retained. The formula gives no justification for changing a fixed physical source direction to a projected direction.

There is also a precise condition for integrating out unused auxiliary coordinates. Let X be a centered Gaussian vector and let $A = (A_K, A_R)$ be a standard finite Gaussian vector with independent blocks, independent of X, and suppose

$$F = X + G_K A_K + G_R A_R, \qquad \operatorname{Cov} X = I - G_K G_K^* - G_R G_R^* \geq 0\,. \tag{23.10}$$

For any integrable $\Phi(F, A_K)$, put $X' = X + G_R A_R$. Then X' is independent of A_K, has covariance $I - G_K G_K^*$, and

$$\mathbb{E}\Phi(F, A_K) = \mathbb{E}\Phi(X' + G_K A_K, A_K)\,. \tag{23.11}$$

This is a direct Gaussian marginal identity. It applies only when the entire integrand, including selection indicators and retained multipliers, factors through (F, A_K). If a selection record reads A_R, that record must remain inside the integral until a valid partition has removed this dependence. Merely calling a coordinate unused does not establish (23.11).

Lemma 23.3 (The Mellin gauge with the original joint likelihood)**.** *Use the orthogonal coordinates of Proposition 23.1, including every retained original auxiliary coordinate, and whiten their joint law. Initially work in finite dimension. There are physical maps T, R and one real vector ϑ such that*

$$TX = e + \sum_i (F_i - \overline{F}_i) + a\zeta, \qquad RX = e + \sum_i F_i, \qquad TX = RX + \langle \vartheta, X\rangle \mathbf{1}.$$

Moreover,

$$\|\vartheta\|^2 = a^2 v_0 + \sum_i a^2 \langle 1, G_{U_i} 1\rangle, \qquad R\vartheta = \sum_i v_{U_i}.$$

Let $K > 0$ be a physical observation of degree one. If the indicated moments are finite, then, with $C = I + A$,

$$\begin{aligned}&\mathbb{E}[L_{A,m}(X) Y(X) K(TX)^w] \\ &\quad = \exp\Big(w\vartheta^{\mathsf{T}} m + \tfrac{1}{2} w^2 \vartheta^{\mathsf{T}} C \vartheta\Big)\, \mathbb{E}[L_{A,m+wC\vartheta}(X) Y(X) K(RX)^w]\,.\end{aligned} \tag{23.12}$$

Here $K^w = \exp(w \log K)$ uses the real logarithm of the positive observation. All complex products in the Gaussian formulas are bilinear. The multiplier Y is the complete original multiplier, including all flags and menus. The identity is valid on any common holomorphic germ on which both sides have the stated moment domination.

Proof. Homogeneity gives $K(TX)^w = e^{w\langle \vartheta, X\rangle} K(RX)^w$. Completion of the quadratic exponent gives the pointwise identity

$$e^{w\vartheta^{\mathsf{T}} X} L_{A,m}(X) = e^{w\vartheta^{\mathsf{T}} m + w^2 \vartheta^{\mathsf{T}} C\vartheta/2} L_{A,m+wC\vartheta}(X).$$

This proves (23.12) without moving a discontinuous selector to a complex field. The core coordinate of ϑ is the negative Gaussian score of $\overline{F}_i$, giving $R\vartheta = -\sum_i a^2 G_{U_i} 1$.

For real w, the equivalent change-of-variables formula is

$$e^{w^2\|\vartheta\|^2/2 + w\vartheta^{\mathsf{T}} m + w^2 \vartheta^{\mathsf{T}} A\vartheta/2} \mathbb{E}[L_{A,m+wA\vartheta}(X) Y(X + w\vartheta) K(RX + wR\vartheta)^w]\,.$$

Thus the physical Poisson means are wv_{U_i}, the source contact is $wA\vartheta$, and every retained menu receives its actual translated mean. For complex w this expression means the full physical mean

continuation obtained from the first formula, not pointwise evaluation of an indicator at a complex argument. Fixed Gaussian density moments and small positive and negative moments of K justify finite-rank limits and holomorphy. □

Only the Gaussian complement of the joint observation space generated by RX and every retained original menu may be marginalized. Conditional expectation of the whole likelihood onto this space compresses both its covariance perturbation and its shifted mean. It does not justify deleting the scalar coordinate alone when a retained menu is correlated with it. Original physical source directions are unchanged. For an opaque child, first use its original integrated numerator and then its full physical mean identity. The previously chosen fresh companion, body, and physical direction basis are reused; no new likelihood is put on fresh variables from an earlier conversion. These operations are performed for each deterministic type and maximal-hole family, summed over the complete family, and only then followed by the representation infimum.

Lemma 23.4 (Retained original menus in the torus gauge)**.** *Suppose the fixed-shape shell buffers of a finite original reserve genealogy have bounded overlap, radius at most $C_0 r_0$, and area comparable to the square of their declared radius. Here r_0 is the top covariance-band edge and C_0 is a fixed microscopic constant, allowed to include the owner offset b_0. Each reserve family is supported compactly in its buffer and has its fixed small Dirichlet synthesis norm. It admits an augmentation of the augmented torus law with residual covariance at least half the original covariance. Under the unit scalar tilt (the general real Mellin tilt multiplies all the following means by w), a reserve coordinate with actual physical bump g_v (including its small amplitude factor) receives the deterministic mean*

$$c_v = \overline{g}_v - \sum_i \overline{Q_{U_i} g_v}\,.$$

In particular, every menu wholly inside a maximal protected hole remains centered. Every menu wholly outside the maximal holes receives $c_v = \overline{g}_v$. These exterior means have a uniformly bounded total squared norm. They may be retained in the same joint Gaussian envelope, with one fixed density moment; no read coordinate is deleted or independently resampled.

Proof. The augmented torus Cameron–Martin norm is the sum of Dirichlet energy and $|\overline{g}|^2/(a^2 v_0)$. Bounded overlap bounds the energy synthesis by a fixed multiple of the largest local reserve norm. Poincaré on a shell buffer gives $\|\overline{G_v}\| \le C r_v^2 \|G_v\|$. Consequently the squared scalar-row norm is at most

$$C \sum_v r_v^4 \|G_v\|^2 \le C_{C_0} r_0^2 \sup_v \|G_v\|^2\,.$$

Here $\sum_v r_v^2 \le C$ follows from overlap of the full fixed-shape shell buffers after the fixed thinning, not merely from overlap of tiny private supports. The local synthesis norm already includes the complete finite menu, so there is no separate menu-size factor. Since $v_0 \asymp r_0^2$, a fixed further reduction of the reserve amplitude gives the claimed one-half bound. This requirement is included in the original menu choice, before its local density constants and C_* are fixed; no retained amplitude is rescaled at the terminal. There is no assertion of global orthogonality of nonzero-average torus bumps.

In the joint augmented law, $\mathrm{Cov}(F, A_v) = g_v$. Its Dirichlet core projection therefore satisfies $\mathrm{Cov}(F_i, A_v) = Q_{U_i} g_v$, whereas $\mathrm{Cov}(a\zeta, A_v) = \overline{g}_v$. Exponential tilting by $a\zeta - \sum_i \overline{F}_i$ shifts A_v by exactly c_v. The declared supports give $Q_{U_i} g_v = g_v$ for a bump inside U_i and $Q_{U_i} g_v = 0$ for one outside every U_i. The same scalar-row estimate bounds $\sum_v c_v^2$.

Define the exterior projection of a physical bump by

$$E_{\mathrm{ext}} g = g - \overline{g} - \sum_i \big(Q_{U_i} g - \overline{Q_{U_i} g}\big), \qquad G_{\mathrm{ext}} a = \sum_{v \text{ exterior}} a_v E_{\mathrm{ext}} g_v \,.$$

It is $g - \overline{g}$ for an exterior bump and zero for an inner bump. This is the synthesis in the mean-zero complementary field e; it is not the synthesis of the unprojected physical bumps. An inner core and its own original menu are independent of the exterior quotient and all exterior menus: each cross-covariance vanishes by the just computed Dirichlet projections. Its menu has zero tilted mean, and its residual core therefore has precisely the coherent Poisson mean v_{U_i}. The original shell orthogonality and commuting Dirichlet cuts are valid inside that core. Exterior menus retain their actual means c_v and their original independent covariance. The exterior field e itself has zero tilted mean. Consequently its residual $e_{\mathrm{res}} = e - G_{\mathrm{ext}} A$ has mean $-G_{\mathrm{ext}} c$, in addition to any original source and contact means; it is not silently recentered. The simultaneous shift of the residual and menu coordinates is $(-G_{\mathrm{ext}} c, c)$, one Gaussian shift of the full exterior block. Its squared Cameron–Martin norm is

$$\|c\|^2 + \|G_{\mathrm{ext}} c\|^2_{C_{\mathrm{res}}^{-1}} \le C \|c\|^2 \,.$$

The projected exterior residual differs from the original compact-bump residual only by the known scalar $-\sum_{v \text{ exterior}} \overline{g}_v A_v$. Normalized residual predicates may use either representative with its consistently normalized length unit. The degree-one body itself remains evaluated at the exact completed field $RX + wR\vartheta$ in the real-w form of (23.12); that scalar is never dropped from the distance. Every fixed density moment of this one Gaussian translation is at most $\exp(C_p \sum_v c_v^2)$, so one Hölder inequality transfers the conditional failed-observation bound with a fixed loss in its prescribed exponent. Every retained unintegrated exterior observation buffer avoids all maximal opaque holes. On such a buffer RX and TX differ by a scalar, so their observations modulo constants coincide exactly. The cutoff/Poincaré proof of the Gaussian angle bound applies in this quotient observation space. For the combined field/menu observations, the augmentation bound $\|G\| \le 1/\sqrt{2}$ makes the joint covariance uniformly comparable to the direct sum of its field covariance and the independent menu covariance: after whitening the field block its covariance has the form $\left(\begin{smallmatrix} I & B \\ B^* & I \end{smallmatrix}\right)$ with $\|B\| \le 1/\sqrt{2}$. Combining this comparison with the field angle bound and the independent menu axes gives a fixed block-angle bound for the centered annular observations and their retained menus. The same assertion follows after the exterior projection, whose residual inverse remains bounded. No claim is made that RX itself is a mean-zero torus field. Thus the complete exterior record is estimated once, without claiming an independent conditional law for its individual selected histories. □

Integrability convention. Proposition 23.1 is an identity of nonnegative extended integrals, or of absolutely integrable signed observations. It is not an assertion that an endpoint distance has a first moment. In the terminal argument below it is used only through the small Mellin formula (23.12). The norms of the total Poisson and menu shifts are uniformly bounded for disjoint maximal holes, and the coefficient families have fixed additional mean and Cauchy margins for this one joint evaluation.

24 A small Mellin terminal and its logarithmic derivative

Fix $a > 0$, distinct deterministic torus points p, q, and a fixed third pin o. Write $q(a) = aQ(a) > 0$. Choose $r_0 = 1/K$ small enough for every protecting buffer and connector enlargement, put $b = b_0^L$

and $r_j = r_0 b^{-j}$, and let the owner mesh at band edge r_j be $s_j = b_0 r_j$. The block from s_j to s_{j-1} changes exactly the covariance bands $[r_j, r_{j-1}]$. Thus N blocks change $[r_N, r_0]$.

The precise analytic input. We use Proposition 22.1 together with the original-record estimates of Section 21, before the representation infimum. Besides $|b_j(z)| \le C|z|$, the needed input is the following. On a smaller fixed response disk the complete coefficient families, their reference lifts, and every unreset microstage admit nonnegative original-record majorants $a_j(X;t)$ satisfying

$$\sup_c \sum_{X \ni c} a_j(X;t) e^{A_b|X|} \le Ct, \qquad 0 \le t < \rho, \tag{24.1}$$

with separate unused spatial reserves for trees, pin chains, and the terminal. Every finite union keeps the original source and read lists; its joint ordinary Gaussian envelope has the bound $C_p^{k+q} \prod_i A_i$ of Proposition 21.2. The atomic factor C_*^{k+q} has a separate reserve for one additional terminal evaluation of that complete union. It is not recharged at an ancestor or an empty age. Converted identities hold on the full retained physical mean family, with a fixed extra mean margin; their value mass is zero on every such fiber. Raw primary derivatives annihilate all retained selectors and the other maximal opaque bodies. These are properties of the stated legal records, not consequences of bounded-observable continuity alone.

For clarity, the quotient is tested initially by bounded smooth scalar functions of a positive local body glued to a finite positive directed exterior graph whose edges are genuine fixed routes. It includes the actual unintegrated whole-domain scalar distance. Equality is through every declared physical mean and its finite-dimensional holomorphic continuation. We will first prove all estimates for arbitrary legal representatives of finite complete families. The estimates then pass to the separate cost infima; no additional norm on arbitrary rough exterior data is introduced.

The representatives and all uniform pre-quotient bounds in this paragraph are those constructed in Lemma 22.3. Only unmarked passing bulk coefficients are replaced at a new normalized stage. All intervening microstages are generated from the chosen representatives, and a first marked record thereafter retains its actual original primitive.

At an unreset stage, the consumed physical slot list of its chosen representative is retained. The distinct terminal reserve in C_* is fixed after the ordinary moment orders and original coefficient contours, and before L and the spatial weights.

24.1 The fixed-depth transform

Let H_0 be the mean-zero torus GFF and

$$v_0 = \frac{1}{2}\left(\int \phi\right)^2 r_0^2, \qquad F_0 = a(H_0 + \zeta), \qquad \zeta \sim N(0, v_0), \quad \zeta \perp H_0.$$

Here ϕ is the compact covariance kernel specified in Section 25. Let D_0 be its canonical parameter-a metric, in the convention that a physical additive shift f has Weyl factor e^f. Let K_N be the full covariance of the white-noise bands $[r_N, r_0]$, including the spatial constant mode. Relative to the covariance C_a of F_0, set

$$A_N = C_a^{-1/2} K_N C_a^{-1/2}, \qquad \|A_N\| \le a^{-2}.$$

It is positive trace class for each finite N. In its eigenbasis define

$$L_{N,z}(X) = \prod_\ell (1 + z\lambda_\ell)^{-1/2} \exp\left(\frac{1}{2}\sum_\ell \frac{z\lambda_\ell}{1 + z\lambda_\ell} X_\ell^2\right), \qquad \Phi_N(z,w) = \mathbb{E}[L_{N,z} D_0(p,q)^w].$$

Write $\Phi_0(w) = \mathbb{E}[D_0(p,q)^w]$ and $B_N(z) = \sum_{j=1}^N b_j(z)$.

Lemma 24.1 (Fixed-depth joint holomorphy)**.** *There are positive ρ, η, independent of N, such that Φ_N is jointly holomorphic on $D_\rho \times D_\eta$. Its bound may depend on N. Moreover,*

$$\ell_N(z) := \partial_w \Phi_N(z,0) = \mathbb{E}[L_{N,z} \log D_0(p,q)].$$

For real $x \in (-\rho, \rho)$ and $t = \sqrt{a^2 + x}$,

$$\Phi_N(x,w) = \mathbb{E}[(D^{a,t}_{r_N}(p,q))^w], \qquad \ell_N(x) = \mathbb{E} \log D^{a,t}_{r_N}(p,q).$$

Proof. For a fixed a, the normalized deterministic-endpoint estimate of [4, Lemma 2.9] allows every positive order strictly below $2Q(a)/a$. On strict local charts, the local Gaussian comparison and a strict Hölder margin therefore give some $s_a > 0$ with $\mathbb{E}D_0(p,q)^{s_a} < \infty$. A fixed annulus separating p and q and the continuum-crossing estimate [4, Lemma 2.6] give $\mathbb{E}D_0(p,q)^{-s} < \infty$ for every fixed $s > 0$. These statements concern a deterministic pair; no supremum over points is taken.

Fix a finite $p > 1$ and then choose $\eta > 0$ such that $2p'\eta < s_a$. For $|z| < c_p a^2$, the one-dimensional Gaussian quadratic formula bounds the finite likelihood products locally in L^p; finite-rank products converge locally in L^p since A_N is trace class. The constants may depend on $\operatorname{Tr} A_N$, but the disk does not. Hölder's inequality, with the positive and negative moments just stated, gives joint holomorphy and differentiation at $w = 0$. The signed covariance identity in Section 25 identifies the real likelihood for both signs of x. □

24.2 Ordinary routes replace diameter estimates

Every maximal modified child has a complete protecting buffer V_i, and the terminal guards require

$$\operatorname{diam} V_i \le C \operatorname{dist}(p, V_i), \qquad \operatorname{diam} V_i \le C \operatorname{dist}(q, V_i). \tag{24.2}$$

The complete buffers are disjoint and have their unchanged ordinary surrounding circuits. These guards apply before a local reset. Constants in the strict-chart cutoff include all fixed buffer and connector enlargements.

Here V_i is the fixed complete protecting domain declared by the corresponding record, with its fixed surrounding-route enlargement included. It is not a field-selected minimal neighborhood of the modified paths. The domains H_0, H_c were fixed from the source-grid catalogue before the private atlas phase, and the random modification and its protecting region lie inside that declaration. Work first with one finite record type and one such maximal-domain family. All field-selected templates and eligibility flags remain in the original coefficient envelope; no ordinary metric law is conditioned on those flags when an annular moment is used. The annulus avoiding the V_i is therefore chosen from its deterministic geometric list. The uniform estimate is then summed over record types and integrated over the independent phases with their actual legal costs.

Lemma 24.2 (Small Mellin caps)**.** *Fix an ordinary exponent $h > 1$. There exist $\eta, \alpha, \kappa > 0$ and $C < \infty$, depending on a, h and the fixed geometries, with the following property. For a positive terminal scalar body K with m maximal guarded children, all evaluated in the original ordinary field or its independent-core gauge,*

$$\sup_{|w| \le 2\eta} \|K^w - 1\|_{L^h} \le C^{m+1} |w|. \tag{24.3}$$

Suppose a protected primary is contained in $B(p,r)$, with its entire declared body included in this ball. In the original physical expression it can be tested against a collapsed body K^0, or against the constant 1 *if the required collapse is macroscopic, so that the primary annihilates the subtracted test and*

$$\sup_{|w|\le 2\eta} \|K^w - (K^0)^w\|_{L^h} \le C^{m+1}|w|r^\kappa. \tag{24.4}$$

In the macroscopic case $(K^0)^w$ in this display means 1*. The assertion is symmetric in p,q. It requires only small positive endpoint moments and fixed annular moments.*

Proof. First fix the deterministic hole family. Each guard (24.2) implies that V_i intersects annuli about either endpoint only over a bounded interval of logarithmic radii: the ratio of its largest to smallest endpoint distance is at most $1+C$. Enlarging by the fixed protecting constants preserves this bound. Thus one can choose an ordinary annulus about q, disjoint from every complete buffer, with radius $u \ge cC_0^{-m}$, in a fixed ball not containing p. Every admissible route from p to q crosses it. If its ordinary crossing distance is A_u, then $K \ge A_u$. The same conclusion holds for the collapsed body below, as long as its collapsed ball stays outside this fixed ball about q.

For each fixed $v > 0$, the continuum annulus estimate and Gaussian circle-average moments give

$$\mathbb{E}A_u^{-v} \le C_v u^{-c_v} \le C_v^{m+1}.$$

Here and below one may increase a moment by a fixed Hölder factor when passing to a strict torus chart or to the independent-core gauge. This uses a bounded smooth Gaussian correction, not a bound for a metric diameter.

Choose an ordinary deterministic-endpoint route from p to a prescribed point at distance comparable to R, contained in a slightly larger disk, and an ordinary surrounding circuit. The endpoint estimate of [4, Lemma 2.9] and annular estimate of [4, Lemma 2.6], with exact metric scaling, imply that for every sufficiently small $s > 0$ the sth moment of their total cost is at most

$$C_s R^{sq(a)-cs^2a^2}.$$

The constant c accounts for a fixed number of Hölder comparisons and is independent of R and of the hole family. For example, the circle-average factor is Gaussian of variance $\log R^{-1} + O(1)$, and one strict Hölder inequality suffices to separate it from the normalized point-distance moment. Choose $\alpha > 0$ so small that every endpoint order used in this proof is strictly below $2Q(a)/a$ and $\alpha q(a) - c_h\alpha^2 a^2 > 0$. Shrinking it again if necessary gives a $\kappa > 0$ with the estimates below.

Starting above Cr, choose an enclosing annulus among a geometric sequence that avoids all complete buffers. Its radius satisfies

$$Cr \le R \le C_0^{m+1} r.$$

A fixed further enlargement ensures that each relevant buffer is wholly inside its inner side or wholly outside its outer side. Let Γ be an ordinary surrounding circuit in this annulus, disjoint from every complete buffer. Collapse the inner disk to a free initial vertex, producing $K^0 \le K$. The collapsed disk may meet child bodies. This causes no difficulty: splice the lifted path only at Γ, where every turn is an ordinary unrestricted turn. Take the last departure from the collapsed disk in a near-minimizing quotient path. Its tail meets Γ. Join p to Γ by the ordinary endpoint route. If that route visits a child, replace the portion between its first and last visit to the child's unchanged surrounding circuit by a subarc of that circuit. Disjoint complete buffers prevent a detour from entering another modified body. At most m detours are used. Join to the quotient tail by a subarc of Γ.

This proves

$$0 \le K - K^0 \le C_r, \qquad \|C_r^\alpha\|_{L^{2h}} \le C^{m+1} r^\kappa.$$

Indeed the endpoint route, Γ, and the at most m child circuits all have radii at most CR, by the endpoint guard. Their small moments have positive scaling powers; the factor $R/r \le C_0^{m+1}$ and the finite sum are absorbed into C^{m+1}. No arbitrary point on a circuit is used as an endpoint: the connections occur at actual intersections of the chosen routes.

For $x \ge y > 0$, $|w| \le 2\eta < \alpha$, elementary integration gives

$$|x^w - y^w| \le C_{\alpha,\eta} |w| (x-y)^\alpha (y^{-\alpha-2\eta} + y^{-\alpha+2\eta}).$$

To see this, write $x = y(1+v)$ and bound $|(1+v)^w - 1| \le |w| \log(1+v)(1+v)^{2\eta} \le C|w|v^\alpha$. Hölder with the preceding routing and inverse-crossing bounds proves (24.4) when the chosen collapse is microscopic.

A finite macroscopic route, with the same first-to-last child detours, gives $\mathbb{E}K^s \le C_s^{m+1}$ for some fixed $s > 0$, while the child-free annulus gives every fixed negative moment. Hence $|K^w - 1| \le |w||\log K| e^{2\eta|\log K|}$ proves (24.3), after reducing η. If the enclosing radius reaches the fixed chart threshold, $C_0^{m+1} r \ge c$; the global estimate therefore implies the required $C^{m+1} r^\kappa$ estimate without collapsing q.

For the gauge comparison, write $TX = RX + S\mathbf{1}$, where $S = \langle \vartheta, X \rangle$. Its variance is bounded uniformly in the disjoint family. For a nonnegative homogeneous route functional J and a real fixed v,

$$\mathbb{E}J(RX)^v = \mathbb{E}[e^{-vS} J(TX)^v] = e^{v^2 \operatorname{Var} S/2} \mathbb{E}J(TX - v \operatorname{Cov}(TX, S))^v.$$

The covariance shift has uniformly bounded supremum by the Green formula and disjointness. Weyl comparison therefore changes a moment only by a fixed factor. The coherent Poisson means have uniformly bounded total supremum as well. For a microscopic cap, both K and K^0 have degree one, so their Mellin difference acquires the common factor e^{wS}. For a global subtraction the exact formula is instead

$$K(TX)^w - 1 = e^{wS}(K(RX)^w - 1) + (e^{wS} - 1).$$

The last term is retained. Its fixed Gaussian moments are $O(|w|)$; equivalently apply the joint likelihood identity to its two terms and then Cauchy's estimate on the fixed larger mean disk. This uses the complete retained numerator and its original source contacts. It is not discarded as a constant. These observations justify the gauge comparisons for both kinds of cap.

Finally the annihilation statement is made before any Gaussian adjunction. A raw primary has physical range inside the collapsed region and annihilates every retained selector; thus the complete exterior functional of K^0 is independent of that physical source coordinate throughout all other source parameters. A converted primary has zero value mass on each full physical mean fiber and its whole body is ignored by K^0. Both statements persist through later physical derivatives and source contacts. They are not assertions that a separately extracted scalar contact kills the collapsed test.

The collapsed body is only a subtracted observation. It is never supplied as an old row to Proposition 19.1 or Lemma 19.2; in particular, no comparison between its value and the conditional essential infimum m_i is required. For a converted primary, cancellation is the fiber identity

$$\int h_\tau(v; y) g_{+,\omega}(y)\, dy = \int \mu_\tau(v; y)\, dy = 0, \qquad \mu_\tau(v) = -\operatorname{div} \sigma_\tau(v).$$

The collapse ignores the entire primary body, so its value is constant in the fresh variables integrated here. The exterior data and the other bodies are fixed on this fiber. The companion and denominator are fixed throughout the physical mean family, so subsequent mean derivatives preserve this zero identity. On the torus, first perform the gauge and source-contact calculation on the original integrated numerator, as in Lemma 23.3, and then use this same fiber identity; no new likelihood is applied to an earlier conversion's fresh variables. □

Extension from bounded tests. Choose smooth truncations χ_M of the real logarithm, equal to the identity on $[-M, M]$, with $|\chi_M(x)| \le |x|$ and $0 \le \chi_M' \le 1$. For fixed $\delta > 0$, the tests $\exp(w\chi_M(\log(t+\delta)))$ are bounded smooth scalar tests on $t \ge 0$ (real and imaginary parts are taken separately). Apply the finite physical identities first to these tests. The inequalities in Lemma 24.2, with a slightly larger fixed moment, dominate their values and differences. Taking $M \to \infty$ and then $\delta \downarrow 0$ proves the same identities for K^w and for its derivative at $w = 0$. This is done for an arbitrary legal complete representative before taking its cost infimum. Thus equality on the original quotient survives the extension. No expression $\mathbb{E}[|P|K]$, nor a first moment of K, is used.

24.3 The marked forest bound

Lemma 24.3 (Pinned finite forests)**.** *Under (24.1), the original-record union estimates, and whole-owner progress, the complete marked expansion has the following nonnegative majorant. Empty transport of an existing marked record has norm at most one. A complete nonempty attachment step, including new disconnected births at the three virtual pins, has norm at most Ct. A first endpoint contribution at block radius r_j has terminal mass at most Ctr_j^{κ}; a first macro contribution has mass at most $Ctr_j^{-2}e^{-c/r_j}$. All these bounds concern the original coefficient expressions, with their once-only terminal reserve.*

Proof. Work first at finite depth and coefficient degree. Before any reset, reject a whole owner that fails a guard at p or q, or that fails strict-chart injectivity. At its first failure pin it respectively to p, q, or o, with a deterministic cell chain. A pin chain has $O(n)$ cells for an n-cell endpoint owner; a macro owner has $n \ge c/s$ at mesh s, which pays its $O(s^{-1})$ chain. Frozen numerical countertypes retain their original planar mass as a birth decoration. Their periodic carrier and pin chain cost at most a constant times that mass. The decoration is not charged at later ages. Group components sharing a pin, keeping all old compatibility indicators inside the grouped expression. There are at most three marked components. A marked component is never reset and never made opaque again.

This preserves the finite identity: a passing output contains every contributing ancestor, so all its ancestors pass. With the numerical counterterms frozen, the unique connected-component partition commutes with restriction to passing outputs. Original exponential factorials are retained. Adding a pin chain only changes the grouping, not the physical support or an internal compatibility factor.

Use the total current cell mass n of the at most three pinned carriers, including their chains. Whole-owner progress, with constants enlarged once for the three pins, gives

$$n_{\mathrm{out}} \le \theta n + C_{\mathrm{att}} \sum_i n_i + Ck + C_0, \qquad \theta < 1.$$

Here n_i are newly attached owner masses and k their number. Choose $n_* \ge C_0/(1-\theta)$ and put

$\Upsilon(n) = (n - n_*)_+$. Then

$$\Upsilon(n_{\mathrm{out}}) \le \theta\Upsilon(n) + C_{\mathrm{att}} \sum_i n_i + Ck.$$

Give the marked carrier weight $e^{A_p\Upsilon(n)}$. At an empty stage the old record is literally unchanged: there are no new sources, reads, fresh variables, likelihoods, or conditioning events. The unread positive assignments sum to one, and Υ cannot increase. The empty transfer therefore has norm at most one. In particular the structural factor in the next paragraph is not charged to an empty stage.

For a nonempty step, the finite union estimate costs at most $e^{cn}C^k \prod_i e^{cn_i} a_j(X_i;t)$. Use a separate spatial reserve in (24.1) to pay $e^{A_pC_{\mathrm{att}}n_i + cn_i}C^k$. The rooted-tree recurrence from Section 22 now bounds the newly attached forests at an n-cell root by e^{Ctn}: if $T(X) = w(X)\exp(\sum_{Y\sim X} T(Y))$, induction on tree depth gives $T(X) \le w(X)e^{\mu|X|}$ as soon as the weighted per-cell mass is at most μ/D. The sum of trees incident to the root is at most Ctn, and exponential multiplicities give e^{Ctn}. This is a deterministic positive majorant and uses no independence.

Removing the empty forest and comparing the incoming and outgoing marked weights gives

$$e^{cn - A_p(1-\theta)\Upsilon(n)}(e^{Ctn} - 1) \le C't.$$

Choose $A_p(1-\theta) > c + 1$, then decrease t. For $n > n_*$, use $e^v - 1 \le ve^v$ and the remaining exponential decrease in n; for $n \le n_*$ use its bounded size. The same total-mass estimate includes mergers of marked carriers. An ordinary attachment which reaches two marks is simply included in the same rooted forest bound.

Keep an empty virtual root at each pin. A previously empty endpoint pin receives new owners of mass n from at most Cn^2 possible anchor positions; the chain and these positions are paid by unused exponential mass. Its total new value weight is at most

$$Ct \sum_{n\ge 1} n^2 e^{-cn} \le C't.$$

At a macro pin the corresponding sum is at most $Cts^{-2}e^{-c/s} \le C't$. These are direct per-pin sums. Three such sums and their products are still $O(t)$ for small t; no factorization of a Gaussian law is being asserted. Consequently the full propagation operator, including disconnected new marked births, has majorant norm at most $1 + Ct$ per fixed block.

Distinguish the unique earliest terminal event, breaking ties by a fixed order of pins and original vertices. Choose one primary in its finite original expression. Apply Lemma 24.2 to that primary before adjunction, and then the single joint ordinary Gaussian envelope to the flattened record. The resulting factor C^{k+q} is paid once by the reserved part of C_*^{k+q}; the original read and source lists prevent repeated charges. Later marks use their value identities, with no second cap. For an endpoint event the source radius is at most Cns, so its birth sum, including the cap, is bounded by

$$Cts^\kappa \sum_{n\ge 1} n^{2+\kappa} e^{-cn} \le C'ts^\kappa.$$

The macro birth uses the global cap and the mass lower bound $n \ge c/s$, giving $Cts^{-2}e^{-c/s}$. Grouping the fixed L microstages gives the stated bounds with r_j. This establishes the lemma by induction on finite construction depth and coefficient degree. □

Proposition 24.4 (Uniform Mellin remainder)**.** *After decreasing fixed positive ρ, η, there are jointly holomorphic R_N on $D_\rho \times D_{2\eta}$ such that*

$$e^{-wB_N(z)}\Phi_N(z,w) = \Phi_0(w) + R_N(z,w), \qquad |R_N(z,w)| \le C|z|, \qquad R_N(z,0) = 0, \tag{24.5}$$

uniformly in N. Consequently

$$\ell_N(z) = \ell_0 + B_N(z) + E_N(z), \qquad E_N(z) = \partial_w R_N(z,0), \qquad |E_N(z)| \le C|z|. \tag{24.6}$$

No zero-free assertion for Φ_N is required.

Proof. At finite depth and degree, form one germ for every actual periodic band/cell slot. Periodize only the frozen numerical countertypes, preserving their planar owner and mass decorations, and aggregate coincident physical mean directions before coefficient extraction. The complete mask orbit sums to 1, so its assembled counterterm is exactly $B_N D_1$. No unintegrated planar source list is periodized through a noninjective identification.

Apply the finite block identity, the guarded marking procedure, and the bounded-test extension above. A physical constant shift of K^w multiplies it by e^{wc}; hence the assembled counterterm contributes exactly e^{-wB_N}. All ordinary gauging uses (23.12) with the full original numerator and the retained-menu calculation of Lemma 23.4. The Poisson, source-contact, and menu means consume their fixed joint margin once. A new likelihood is never assigned to a previously fresh companion.

For completeness, the original torus slot synthesis has the same fixed bounds as the planar one. A compact radius-r slot has scalar row of squared norm $O(r^4)$; $O(r^{-2})$ slots at that scale contribute $O(r^2)$. Summing scales to r_0 gives $O(r_0^2)$, which is compensated by $v_0 \asymp r_0^2$. Disjoint maximal holes satisfy $\sum r_i^2 \le C$ and hence $\sum \|v_{U_i}\|_{E(U_i)} + \sum \|v_{U_i}\|^2_{E(U_i)} \le C$. Nested Poisson means are projections of the same original means. The additional factors therefore fit the single Gaussian terminal reserve, uniformly in depth.

A first endpoint event at radius r_j is followed by at most $j + O(1)$ propagation blocks. By Lemma 24.3, its absolute coefficient sum is bounded by

$$\sum_{j=0}^{N} Ctr_j^{\kappa}(1+Ct)^j \le \frac{Ctr_0^{\kappa}}{1 - b^{-\kappa}(1+Ct)} = O(t),$$

after choosing ρ so that $b^{-\kappa}(1 + C\rho) < 1$. First macro events contribute at most

$$\sum_{j=0}^{N} Ctr_j^{-2} e^{-c/r_j}(1+Ct)^j = O(t).$$

The remaining unmarked gas at the fixed top scale has finitely many cells. Its complete value majorant is at most $\exp(C_{r_0} t)$, so it multiplies each preceding bound by one uniform factor; when there is no earlier marked event, its nonempty part contributes $O(t)$ using the global Mellin cap. The cap is used once at the distinguished earliest event; later attachments were already included in the propagation bound. These estimates are uniform for $|w| \le 2\eta$ and for arbitrary legal coefficient representatives.

Thus the complete series converges normally on the common joint disks and defines R_N with the stated bound. Its finite coefficients are exactly those of the physical torus source product with $e^{-B_N D_1}$; Lemma 24.1 and the identity theorem identify the analytic series with the left side of (24.5). At $w = 0$ both transforms equal 1, so $R_N(z,0) = 0$. Differentiation and Cauchy's estimate in w give (24.6). □

25 The real pressure on both sides of the base parameter

Fix any two parameters $a, t > 0$. All constants in this section may depend on these two parameters. In particular, no assertion of uniformity as the base parameter varies is needed. Write

$$q(s) = sQ(s), \qquad \lambda(s) = 1 - q(s)\,.$$

We use the canonical metric associated with the field sh, with Weyl factor e^f for a physical additive shift f.

Choose a nonnegative radial function $\phi \in C_c^\infty(\mathbb{R}^2)$ with $\|\phi\|_{L^2(dx)} = 1$, and put $\phi_r(x) = r^{-1}\phi(x/r)$. In particular $\int \phi > 0$. Periodize this function on the area-one flat torus. On $L^2(\mathbb{T}^2 \times (0,\infty), dx\,dr/r)$ let W be real white noise. In the integral over $r \geq r_0$, remove the spatial constant mode. Here $r_0 > 0$ is fixed and small enough for the local charts used below. Set

$$U_\epsilon = \int_0^\epsilon \phi_r * W(dr), \qquad V_\epsilon = \int_\epsilon^{r_0} \phi_r * W(dr), \qquad I = \int_{r_0}^\infty (\phi_r * W(dr))_0\,. \tag{25.1}$$

The three fields are independent; the last two are smooth for fixed $\epsilon > 0$. The field $U_\epsilon + V_\epsilon + I$ has the law of a mean-zero torus GFF plus an independent centered Gaussian constant. Indeed, if Φ is the ordinary Fourier transform of ϕ, then, for $k \neq 0$,

$$\int_0^\infty r^2 |\Phi(rk)|^2 \,\frac{dr}{r} = \frac{1}{|k|^2}\int_0^\infty u|\Phi(ue_1)|^2\,du = \frac{2\pi}{|k|^2}\,.$$

The variance of the constant is $(\int \phi)^2 r_0^2/2$.

Define M_ϵ to be the canonical parameter-a metric of aU_ϵ, and set

$$F_\epsilon = tV_\epsilon + aI, \qquad D_\epsilon^{a,t} = e^{F_\epsilon}\cdot M_\epsilon\,. \tag{25.2}$$

The metric M_ϵ is a measurable function of U_ϵ alone. Here is a direct Gaussian justification of its definition and of this measurability. Put

$$g(u) = (2\pi)^{-1}\int_u^\infty v|\Phi(ve_1)|^2\,dv\,.$$

For each fixed $\epsilon > 0$, the nonconstant covariance eigenvalues of U_ϵ are $2\pi(1 - g(\epsilon|k|))/|k|^2$. They are strictly positive, since $\Phi(0) > 0$, and their ratios to the torus GFF eigenvalues tend to one faster than every power of $|k|$. The constant variance is $(\int \phi)^2\epsilon^2/2 > 0$. Thus its law is mutually absolutely continuous with that of a mean-zero torus GFF plus an independent Gaussian constant of any fixed positive variance. To verify the infinite-dimensional assertion, separate finitely many low Fourier modes. Their Gaussian densities are strictly positive. On the remaining modes the variance ratios lie in a fixed neighborhood of one and their squared deviations are summable. The finite Gaussian likelihood products, and their reciprocal products under the other law, then have bounded L^2 norms, by the elementary one-dimensional Gaussian integral and the bound $|\log(1-u) + u| \leq Cu^2$ near zero. Their martingale limits give the two absolute-continuity statements.

Apply the canonical measurable metric map to aU_ϵ on this equivalent Gaussian law. For every fixed smooth function the Weyl identity holds almost surely under that law. Conditional on the independent smooth field $V_\epsilon + I$, Fubini therefore gives

$$M_{a(U_\epsilon + V_\epsilon + I)} = e^{a(V_\epsilon + I)}\cdot M_{aU_\epsilon} \quad \text{almost surely}\,.$$

Consequently M_ϵ is independent of F_ϵ, and the metrics can equivalently be defined by removing the indicated smooth field from the canonical parameter-a torus metric. The absolute-continuity constants here need not be uniform in ϵ; the uniform fine estimates below use a different, exactly rescaled local comparison. Thus $D_\epsilon^{a,a}$ is independent of ϵ under this coupling. For a fixed deterministic pair $p \neq q$, the small endpoint moment and inverse annulus moments used in Lemma 24.1 imply

$$\mathbb{E}D_\epsilon^{a,a}(p,q)^{s_a} < \infty, \qquad \mathbb{E}D_\epsilon^{a,a}(p,q)^{-s} < \infty \quad (s>0), \qquad \mathbb{E}|\log D_\epsilon^{a,a}(p,q)|^k < \infty \quad (k<\infty). \tag{25.3}$$

The deterministic-endpoint estimate is [4, Lemma 2.9], with its strict range $p < 2Q(a)/a$; the continuum annular estimates are [4, Lemma 2.6]. Locality and Weyl scaling are those of the canonical metric [1, Definition 1.6 and Theorem 1.7] and the uniqueness theorem [4]. These statements, together with the previously established local restriction and gluing convention, specify the metric map used here.

Proposition 25.1 (Two-sided real logarithmic pressure)**.** *For every fixed $a, t > 0$ and deterministic $p \neq q$,*

$$\frac{\log D_\epsilon^{a,t}(p,q)}{\log \epsilon^{-1}} \longrightarrow q(t) - q(a) \quad \textit{in probability and in } L^1. \tag{25.4}$$

In fact these normalized logarithms are uniformly integrable to every fixed positive power. In particular,

$$\lim_{\epsilon \downarrow 0} \frac{\mathbb{E}\log D_\epsilon^{a,t}(p,q)}{\log \epsilon^{-1}} = q(t) - q(a).$$

The assertion holds for $t < a$, $t = a$, and $t > a$.

25.1 Fine routes and coarse crossings

Use a torus square grid of mesh between ϵ and 2ϵ, and let $Q^{(c)}$ denote the concentric fixed-factor enlargement of Q. Put

$$L_Q = M_\epsilon(Q, \partial Q^{(3)}; Q^{(4)}).$$

The internal-distance boundary convention is the increasing strict-domain limit. For a deterministic center z, let $A_{\epsilon,z}$ be the infimum of the M_ϵ-lengths of ordinary circuits in the annulus with radii $16\epsilon, 18\epsilon$, and let

$$R_{\epsilon,z} = M_\epsilon(z, z + 40\epsilon e_1; B(z, 60\epsilon)).$$

Changing these fixed numerical ratios, if needed for the chosen grid, only changes constants. Exact rescaling of the compact ultraviolet field gives

$$\sup_Q \mathbb{E}L_Q^{-v} \le C_v \epsilon^{-vq(a)}, \qquad \sup_z \mathbb{E}A_{\epsilon,z}^{v} \le C_v \epsilon^{vq(a)} \quad (v>0), \qquad \sup_z \mathbb{E}R_{\epsilon,z}^{s} \le C_s \epsilon^{sq(a)} \tag{25.5}$$

for sufficiently small $s > 0$. Indeed $U_\epsilon(\epsilon\,\cdot)$ is exactly the unit ultraviolet field on every fixed strict rescaled chart. Couple that field to a GFF with a fixed smooth normalization; the difference on the chart is smooth with Gaussian exponential supremum moments. Hölder with a strict margin transfers the continuum annular moments and the endpoint moment of order less than $2Q(a)/a$. The constants are uniform in ϵ and in deterministic translated centers. No metric diameter is used.

For a smooth real F, let d_F be the length metric $e^F|dx|$ and set

$$\omega_\epsilon(F) = \sup_{d_{\mathbb{T}}(x,y)\le 100\epsilon} |F(x) - F(y)|.$$

Lemma 25.2 (Fine and coarse comparisons)**.** *If p, q are separated by more than 100ϵ, then*

$$(e^F \cdot M_\epsilon)(p,q) \geq c\epsilon^{-1} \min_Q L_Q \, e^{-C\omega_\epsilon(F)} d_F(p,q). \tag{25.6}$$

There is an F-measurable finite sequence $z_0 = p, z_1, \dots, z_k$, with successive distances 8ϵ and $|z_k - q| \leq 8\epsilon$, such that

$$(e^F \cdot M_\epsilon)(p,q) \leq e^{C\omega_\epsilon(F)} \left(\sum_{i=0}^{k} e^{F(z_i)} A_{\epsilon,z_i} + e^{F(p)} R_{\epsilon,p} + e^{F(q)} R_{\epsilon,q} \right), \tag{25.7}$$

$$\sum_{i=0}^{k} e^{F(z_i)} \leq C\epsilon^{-1} e^{C\omega_\epsilon(F)} d_F(p,q).$$

Infimum-length circuits and endpoint routes are understood through arbitrarily close ordinary approximations. If F is independent of M_ϵ, then conditional expectation of the circuit sum, but not necessarily of the two endpoint terms, is at most $C\epsilon^{q(a)-1} e^{C\omega_\epsilon(F)} d_F(p,q)$.

Proof. For the lower bound, decompose a finite hybrid-length path by its successive first exits from radius-8ϵ Euclidean balls. Each full segment starts in some grid square Q and exits $Q^{(3)}$, so it has fine cost at least L_Q. There are finitely many full segments because the finite minimum $\min_Q L_Q$ is positive and F is bounded. Merge the final remainder with the preceding segment. A straight connector between the resulting endpoints has coarse cost at most $C\epsilon e^{F(z_i)+C\omega_\epsilon(F)}$, while the original full segment has hybrid cost at least $e^{F(z_i)-C\omega_\epsilon(F)} \min_Q L_Q$. Summing and taking infima proves (25.6).

For the upper bound choose from a countable polygonal family an F-measurable path of coarse length at most $2d_F(p,q)$ and use its radius-8ϵ exits. Neighboring annular circuits intersect: their inner disks overlap, and neither circuit can be contained inside the other's outer disk since $16 + 8 > 18$. The endpoint route from p to $p + 40\epsilon e_1$ meets the first circuit. Since q lies in the final inner disk, the analogous route from q meets the last circuit. Their connected union contains a route from p to q. All routes stay within the fixed 100ϵ comparison neighborhoods, proving (25.7).

Every full segment of the coarse path costs at least $8\epsilon e^{F(z_i)-C\omega_\epsilon(F)}$. The last weight is bounded by a fixed oscillation factor times the previous one. This proves the displayed sum bound. Conditional on F, the centers are deterministic, so (25.5) applies to each circuit. The circuits need not be independent. □

Lemma 25.3 (Coarse-field estimates at every parameter)**.** *For each fixed $c > 0$,*

$$\mathbb{E} e^{c\omega_\epsilon(F_\epsilon)} = \epsilon^{-o(1)}. \tag{25.8}$$

There is $s_* > 0$ such that, for every sufficiently small fixed $0 < s < s_*$,

$$\mathbb{E} d_{F_\epsilon}(p,q)^s \leq \epsilon^{s\lambda(t)-o(1)}. \tag{25.9}$$

For every $\delta > 0$,

$$\mathbb{P}\{d_{F_\epsilon}(p,q) \geq \epsilon^{\lambda(t)+\delta}\} \longrightarrow 1. \tag{25.10}$$

Proof. Put $\ell = \log \epsilon^{-1}$. Differentiating the covariance gives $\sup_x \mathrm{Var}(\nabla^j F_\epsilon(x)) \leq C_j \epsilon^{-2j}$ for $j \geq 1$. The canonical metric of $\epsilon \nabla F_\epsilon$ is at most $C\epsilon^{-1}|x-y|$, and its point variances are bounded. Covering the torus and integrating its Gaussian entropy bound gives expected supremum $O(\sqrt{\ell})$; Gaussian concentration then gives

$$\mathbb{E} \exp(c\epsilon \|\nabla F_\epsilon\|_\infty) \leq \exp(C_c \sqrt{\ell} + C_c).$$

This proves (25.8).

The nonconstant Fourier covariance of F_ϵ is

$$\frac{2\pi}{|k|^2}\{t^2 g(\epsilon|k|) + (a^2 - t^2) g(r_0|k|)\} = \frac{2\pi}{|k|^2}\{t^2[g(\epsilon|k|) - g(r_0|k|)] + a^2 g(r_0|k|)\}.$$

The latter expression is positive for all $a, t > 0$. Couple its modes to those of tH_ϵ, where H_ϵ is a heat-mollified mean-zero torus GFF. The inequality

$$\left|\sqrt{t^2u^2 + a^2v^2} - t\sqrt{u^2 + v^2}\right| \le |a - t|\,|v|$$

separates the smooth infrared difference. For the remaining difference use $\sqrt{g(u)} - e^{-cu^2} = O(u^2)$ at zero and rapid decay at infinity. Fourier summation shows that the difference S_ϵ has bounded point variances and derivative variances at most $C_j\epsilon^{-2j}$. The same entropy argument, including the bounded spatial constant, gives

$$\mathbb{E}e^{c\|S_\epsilon\|_\infty} = \epsilon^{-o(1)} \quad (c > 0). \tag{25.11}$$

The primary all-parameter LFPP input is [12, Propositions 1.1, 2.5, 3.2 and 4.6]: the median exponent is $\lambda(t) = 1 - tQ(t)$ with $Q(t) > 0$, and normalized fixed-shape crossings of the white-noise heat-scale approximation have every fixed positive and negative moment, also for nondyadic cutoffs. We now derive the endpoint estimate needed here from those crossing bounds.

In each dyadic annulus of radius $\delta \ge \epsilon$ around p and q, choose a circuit and a fixed finite radial crossing network using the field at scales between ϵ and δ. Consecutive networks intersect topologically. For every fixed $\zeta > 0$, exact rescaling and the median exponent bound the sth moment of a fine route by $C_{s,\zeta}\delta^s(\epsilon/\delta)^{s(\lambda(t)-\zeta)}$. The independent coarse field has center variance $\log\delta^{-1} + O(1)$ and centered oscillation with bounded exponential moments. With one fixed strict Hölder factor $h_1 > 1$, the resulting route cost C_δ satisfies

$$\mathbb{E}C_\delta^s \le C_{s,\zeta}\epsilon^{s(\lambda(t)-\zeta)}\delta^{s(q(t)+\zeta)-h_1 s^2 t^2/2}.$$

Choose $s < 1$ so small that $q(t) - h_1 st^2/2 > 0$. Subadditivity of x^s sums these dyadic costs. The final connector in the $O(\epsilon)$ ball has moment at most $C_s\epsilon^{s-h_1 s^2 t^2/2}$, which is bounded by the required upper scale because $q(t) - h_1 st^2/2 > 0$. A finite macroscopic chart network connects the two endpoint networks. This gives the point-distance bound for the planar heat-scale field.

To pass to the torus, periodize the heat-scale field between ϵ and fixed r_0. On a strict chart, couple it to planar white noise in a larger chart; the kernel and each derivative coming from outside the larger chart are exponentially small at fine scales. Their supremum has a uniformly bounded Gaussian exponential moment. The remaining infrared field and the removed constant also have bounded such moments. A finite chart cover and a strict reduction of s therefore transfer the bound to the heat-mollified torus GFF. Since $\zeta > 0$ is arbitrary, the result is $\epsilon^{s\lambda(t)-o(1)}$. A fixed annulus separating p from q, the negative crossing moments, and the median exponent give its lower bound in probability. Finally use

$$e^{-\|S_\epsilon\|_\infty} d_{tH_\epsilon} \le d_{F_\epsilon} \le e^{\|S_\epsilon\|_\infty} d_{tH_\epsilon}$$

and (25.11), retaining a strict moment margin. This proves both claims. □

Proof of Proposition 25.1. For fixed $\delta, M > 0$, choose v with $v\delta > M + 2$. The inverse fine-crossing estimate and a union bound over $O(\epsilon^{-2})$ squares give

$$\mathbb{P}\{\min_Q L_Q < \epsilon^{q(a)+\delta}\} \le C_{\delta,M}\epsilon^M.$$

Together with (25.6), the coarse lower bound, and $\omega_\epsilon(F_\epsilon) = o(\ell)$ in probability, this implies

$$\mathbb{P}\{D_\epsilon^{a,t}(p,q) \geq \epsilon^{q(a)-q(t)+C\delta}\} \longrightarrow 1.$$

For the upper bound, write the right side of (25.7) as a circuit term and two endpoint terms. Choose a sufficiently small $0 < s < 1$ with all needed strict moment margins. Conditional Jensen, independence of M_ϵ and F_ϵ, the circuit first moment in (25.5), and the small coarse moment give

$$\mathbb{E}(\text{circuit term})^s \leq \epsilon^{s(q(a)-q(t))-o(1)}.$$

For either endpoint $z = p, q$, independence and Gaussian moments give, for a fixed $h_2 > 1$,

$$\mathbb{E}\big[e^{C\omega_\epsilon(F_\epsilon)+F_\epsilon(z)}R_{\epsilon,z}\big]^s \leq \epsilon^{sq(a)-h_2s^2t^2/2-o(1)}.$$

Indeed $\operatorname{Var} F_\epsilon(z) = t^2\ell + O(1)$, the oscillation has the exponential moments above, and $R_{\epsilon,z}$ has the independent small fine moment. Choose s so small that $q(t) - h_2st^2/2 > 0$. Both endpoint terms are then smaller than the required scale. Subadditivity yields

$$\mathbb{E}D_\epsilon^{a,t}(p,q)^s \leq \epsilon^{s(q(a)-q(t))-o(1)}. \tag{25.12}$$

Markov's inequality gives the complementary probability bound. Hence $\log D_\epsilon^{a,t}/\ell \to q(t) - q(a)$ in probability.

It remains to justify expectation of the logarithm. The preceding display already bounds one positive moment by a fixed power of ϵ^{-1}. For a negative moment observe that every path from p to q contains at least $c\epsilon^{-1}$ radius-8ϵ exit segments. Thus

$$M_\epsilon(p,q) \geq c\epsilon^{-1}\min_Q L_Q, \qquad \mathbb{E}M_\epsilon(p,q)^{-v} \leq C_v\epsilon^v\sum_Q \mathbb{E}L_Q^{-v} \leq C_v\epsilon^{v-2-vq(a)}.$$

The Gaussian supremum estimate at variance $O(\ell)$ gives $\mathbb{E}e^{v\|F_\epsilon\|_\infty} \leq C_v\epsilon^{-C_v}$. Independence and $D_\epsilon^{a,t} \geq e^{-\|F_\epsilon\|_\infty}M_\epsilon$ therefore bound each fixed negative moment by another fixed power of ϵ^{-1}. Consequently, for some $s, A, C > 0$,

$$\mathbb{E}e^{s\log D_\epsilon^{a,t}} + \mathbb{E}e^{-s\log D_\epsilon^{a,t}} \leq Ce^{A\ell}.$$

For $\ell \geq 1$ this implies

$$\mathbb{P}\{|\log D_\epsilon^{a,t}|/\ell > u\} \leq 2Ce^{A-su} \quad \text{for } u > A/s.$$

These uniform exponential tails prove uniform integrability of every fixed power of the normalized logarithm. Probability convergence now implies the asserted L^1 limit. □

Signed covariance and the constant mode. Let C_{tor} denote the covariance of $H_0 + \zeta$ in the band decomposition, and let $K_\epsilon = \operatorname{Cov} V_\epsilon$ include its spatial constant. Then the exact identity is

$$\operatorname{Cov}(aU_\epsilon + tV_\epsilon + aI) = a^2C_{\text{tor}} + (t^2 - a^2)K_\epsilon.$$

It changes the amplitude of an existing independent band; it is valid for $t < a$ without a Gaussian field of negative variance. The relative eigenvalues of K_ϵ are

$$a^{-2}\{g(\epsilon|k|) - g(r_0|k|)\} \quad (k \neq 0), \qquad a^{-2}(1 - \epsilon^2/r_0^2) \quad (k = 0).$$

They lie in $[0, a^{-2}]$ and are summable for fixed $\epsilon > 0$. These are precisely the eigenvalues used in Lemma 24.1. In particular no constant-mode correction is left over in the real pressure identification.

26 Passage to the analytic pressure

Proof of Theorem 1.1. Fix any $a > 0$. Proposition 22.1 and (24.6) give

$$|B_N(z)| \le CN|z|, \qquad \ell_N(z) = \ell_0 + B_N(z) + E_N(z), \qquad |E_N(z)| \le C|z|$$

on one disk independent of N. Hence the holomorphic functions $H_N(z) = \ell_N(z)/(N \log b)$ form a locally bounded family. For real x in a two-sided interval about zero, Lemma 24.1 and Proposition 25.1 identify their limits as

$$H_N(x) \longrightarrow q(\sqrt{a^2 + x}) - q(a),$$

since $\log r_N^{-1} = N \log b + O(1)$. Montel's theorem gives subsequential holomorphic limits, and the identity theorem makes them equal because their real values agree on this interval. A counter-subsequence argument gives local uniform convergence of the entire family to a holomorphic function H_a.

The branch $t \mapsto t^2 - a^2$ and the local square root then give

$$q(t) = q(a) + H_a(t^2 - a^2)$$

for real t near a, hence a holomorphic continuation there. As $a > 0$ was arbitrary, $q(t) = tQ(t)$, $Q(t) = q(t)/t$, and $\lambda(t) = 1 - q(t)$ are real analytic on $(0, \infty)$. Cauchy's estimate on a smaller disk gives uniform factorial derivative bounds on each compact subinterval after a finite cover. □

A Finite target exploration and typical geodesics

In this section, $D = M_{ah}^a$ for a fixed $a > 0$. A filled ball aimed at z is the complement of the component containing z of the complement of the ordinary closed metric ball. We write it as $K_r^z(p)$, where p is the root and $r < D(p, z)$. We allow $D(p, z) = \infty$. The set $K_r^z(p)$ can be unbounded in the plane. Its boundary is nevertheless a Jordan curve, is compact in D, and consists of points at distance r from p; see [3, Theorem 1.4 and Lemma 3.4]. All references to left and right below use the orientation seen from the component containing the target.

Lemma A.1 (Metric approximation through an accessible side)**.** *Suppose x is nonsingular, V is Euclidean open, and there is a simple Euclidean arc $\kappa : [0, 1] \to \mathbb{C}$ with $\kappa(0) = x$ and $\kappa((0, 1]) \subset V$. Then there are rational points $q_n \in V$ with $D(x, q_n) \to 0$. The points can also be required to lie in any prescribed Euclidean neighborhood of x.*

Proof. Take surrounding rectifiable loops J_n about x whose Euclidean diameters, D-lengths, and access distances from x tend to zero, as in [3, Lemma 3.2]. For large n, $\kappa(1)$ is outside J_n, so J_n meets $\kappa((0, 1])$ at a point $y_n \in V$. Choose the access path to J_n within an error tending to zero and then follow J_n to y_n. Thus

$$D(x, y_n) \le D(x, J_n) + \mathrm{len}_D(J_n) + o(1) \longrightarrow 0.$$

The continuity of the identity from the D topology to the Euclidean topology implies that some D-ball about y_n lies in V and in the prescribed Euclidean neighborhood of x. Rational D-density supplies a rational q_n in this ball with $D(y_n, q_n) < 1/n$. □

For a Jordan frontier the required arc in its complementary domain exists at every frontier point by the Jordan–Schoenflies theorem. For a specified local side of a simple geodesic, apply the same theorem in a small disk straightening the geodesic. In particular, metric density alone is not the justification for a sided approximation. For the extremal construction, first approach x along the designated frontier arc by points x_n and then approximate each x_n from the target component. The agreement of the two topologies on the compact frontier gives $D(x_n, x) \to 0$. Taking the approximation error smaller than one fourth of the distance from x_n to the opposite frontier arc ensures that the root geodesic to q_n exits the filled ball on the designated arc. Indeed, if the frontier level is s, then

$$D(P_{q_n}(s), q_n) = D(p, q_n) - s \le D(q_n, x_n),$$

so such an exit cannot lie on the opposite arc. This gives the sided rational approximations used in the extremal-envelope argument.

Lemma A.2 (Finitely many extremal prefixes)**.** *Fix a deterministic point p. Almost surely, for every $z \neq p$ and every $0 < t < s < D(p, z)$, the leftmost geodesics from p to $\partial K_s^z(p)$ have only finitely many distinct restrictions to $[0, t]$. The same assertion holds for the rightmost geodesics.*

Proof. We give the finite target modification of the exploration argument in [3, Section 5.2]. The two changes are a Euclidean separation estimate for the component containing a finite target and a stopping rule that does not condition on later regularity events.

First fix a deterministic target z and deterministic levels $0 < t < s$. Write $K_u = K_u^z(p)$, and set $K_u = \mathbb{C}$ after $T = D(p, z)$. The exploration sigma field at time u is generated by K_u and the field revealed on K_u, with the usual completion and right continuity. Stopped filled balls are local sets for this filtration by [3, Lemma 5.1].

The extremal geodesics and their order are obtained just as in [3, Lemmas 5.4 and 5.6]. Here are the points needed for a finite target. Since ∂K_u is D-compact and the identity to the Euclidean plane is continuous, the two topologies on ∂K_u agree. Lemma A.1, using rational D-density from [1, Proposition 1.13], supplies rational approximations in D from the designated side through the component containing z. Their distances from p tend to u. The compactness result [3, Lemma 5.5] gives geodesic limits, and uniqueness of geodesics to rational points gives the two extremal limits and their sandwich property. In particular, this use of compactness includes convergence of endpoint distances. The region between two filled frontiers is a topological annulus on the sphere. Noncrossing in this annulus proves that an extremal geodesic restricts to the corresponding extremal geodesic at each earlier level. It also proves that, for disjoint half-open arcs on an earlier frontier, the sets of points on a later frontier whose leftmost geodesic passes through those arcs are disjoint connected arcs or empty. These are precisely the ordering arguments of [3, Lemma 5.6]; their only topological input is the annulus between the frontiers.

We next record the Euclidean separation estimate. Suppose a Jordan filled set K satisfies

$$B_d(p) \subset K, \qquad \partial K \subset B_M(p), \tag{A.1}$$

where $0 < d < M$, and let V be its complementary component containing z. For all sufficiently large n, at least half of any n disjoint arcs of ∂K can be separated from z in V by crosscuts of diameter at most $C(d, M)n^{-1/2}$. If V is bounded, this follows from [5, Lemma 2.14], applied to V with distinguished point z; its area is at most πM^2. If V is unbounded, translate p to zero and apply that lemma to

$$U = \{1/w : w \in V\} \cup \{0\},$$

whose area is at most πd^{-2}, with distinguished point $1/z$. The crosscuts have endpoints on the image frontier, where $|w| \geq M^{-1}$. For large n the entire crosscut lies in $\{|w| \geq (2M)^{-1}\}$. Inversion back to V increases its diameter by at most $4M^2$. This is a deterministic use of inversion; neither the field nor the metric is transformed.

For clarity, we state the local exploration estimate used in the iteration. Let τ be an exploration stopping time, let $x \in \partial K_\tau$ and a dyadic $\epsilon > 0$ be measurable at time τ, and assume that the pinning circle has already been revealed. Use the good annulus radii $\rho_\epsilon^j(v)$ of [3, (5.13)], and put

$$R_\epsilon(\partial K_\tau) = 6 \max\left\{\rho_\epsilon^{\lfloor \eta \log \epsilon^{-1} \rfloor}(v) : v \in (\epsilon/4)\mathbb{Z}^2 \cap B_\epsilon(\partial K_\tau)\right\} + \epsilon. \tag{A.2}$$

Only the compact frontier occurs in this maximum. Let σ be the first level whose filled set contains the open R_ϵ-neighborhood of K_τ, with absorption at T. There are constants $C, \alpha > 0$, depending only on a, and an event G_x measurable at time σ such that

$$\mathbb{P}(G_x^c \mid \mathcal{F}_\tau) \leq C\epsilon^\alpha. \tag{A.3}$$

If

$$R_\epsilon < \tfrac{1}{4} \min\{\mathrm{dist}(p, \partial K_\tau), \mathrm{dist}(z, \partial K_\tau)\}, \tag{A.4}$$

then, on G_x, no geodesic from p to a point outside that open neighborhood enters $B_\epsilon(x) \setminus K_\tau$. Consequently, if $B_\epsilon(x)$ separates an arc I from z in the complement of K_τ, no such geodesic to the later frontier passes through I.

We justify the conditional estimate without treating an inside-field event as outside-measurable. Write $K = K_\tau$ and choose, measurably from $\mathcal{F}_\tau$, a grid point z_0 with $B_\epsilon(x) \subset B_{2\epsilon}(z_0)$. Use the annulus events E_r^U, their intersection E_r over the finite square-deletion catalogue, and the internal event H_r^U of [3, (5.3)–(5.12)]. In particular, H_r^U bounds internal distances between the finitely many square centers in each component of U; it does not bound that component's metric diameter. For each candidate dyadic radius r, put

$$U_r = A_{3r,4r}(z_0) \setminus \bigcup\{S : S \text{ is a catalogue square and } S \cap K \neq \varnothing\}.$$

This is one member of the finite catalogue and is disjoint from K. Define a second sequence of radii, distinct from the universal radii used in (A.2), by

$$\widetilde{\rho}_0 = \epsilon, \qquad \widetilde{\rho}_n = \min\{r \geq 6\widetilde{\rho}_{n-1} : r \text{ is an admissible dyadic radius and } E_r^{U_r} \text{ occurs}\}.$$

The inclusion $E_r \subset E_r^{U_r}$ implies, by induction, $\widetilde{\rho}_n \leq \rho_\epsilon^n(z_0)$. For $N = \lfloor \eta \log \epsilon^{-1} \rfloor$, set

$$G_x = \bigcup_{n=1}^{N} H_{\widetilde{\rho}_n}^{U_{\widetilde{\rho}_n}}.$$

Here is the conditional-law argument at a selected radius. Fix possible values of ϵ, z_0, r, U with $\widetilde{\rho}_n = r$ and $U_r = U$. The local-set property of K_τ is available for finite-target balls by [3, Lemma 5.1]. On $K \cap U = \varnothing$, the explored data and the square-deletion domain are determined outside U. All tests at earlier dyadic radii are also outside U: a preceding radius is at most $r/2$, its tested annulus ends at $5r/2$, whereas $U \subset A_{3r,4r}(z_0)$. Thus the selected-radius event is measurable with respect to the outside field together with $1_{E_r^U}$. The pinning circle is in the revealed set, so there is no additional normalization constraint on the zero-boundary innovation in U.

The event E_r^U is not itself outside-measurable. Its short-loop condition uses the field in U, and this is exactly why the conditional FKG estimate of [3, Lemma 5.10] is needed. That estimate gives

$$\mathbb{P}\Big(H_{\widetilde{\rho}_n}^{U_{\widetilde{\rho}_n}} \mid \epsilon, z_0, \widetilde{\rho}_n, U_{\widetilde{\rho}_n}, h|_{\mathbb{C}\setminus U_{\widetilde{\rho}_n}}\Big) \geq p \tag{A.5}$$

for the selected domain and radius, with $p > 0$ independent of the explored ball. Conditioning on each of the countably many selected values proves (A.5). The previous internal tests and $\mathcal{F}_\tau$ are determined outside this selected domain, because successive radii have ratio at least six. Iterating conditional expectations therefore gives

$$\mathbb{P}(G_x^c \mid \mathcal{F}_\tau) \leq (1-p)^N \leq C\epsilon^\alpha.$$

No independence between different boundary arcs is asserted or used.

All selected annuli, including the field needed to test their radii, lie strictly inside the neighborhood in (A.2). The level σ is a stopping time: at any proposed level v, the event that this neighborhood is contained in K_v is verified by the finite list of dyadic radius tests inside K_v and the corresponding set inclusion. Take the countable union over the possible lists. Conversely, if the inclusion holds, all the tests defining that list are already revealed. Right continuity handles the infimum, and absorption at T handles the case in which the target is reached first. The same observation, metric locality, and the strict factor-six buffer show $G_x \in \mathcal{F}_\sigma$. Only grid points near the compact frontier are needed, even when K is unbounded.

We give the local geometric implication as well. Suppose the clearance condition holds and select a successful tested radius r. The connected set K, containing p and meeting the small central ball near x, crosses both circles of $A_{3r,4r}(z_0)$. The good annulus event supplies a cyclic chain of squares separating the two annulus boundaries. Mark the squares which meet K. Between two successive marked squares, the unmarked squares belong to a single component of U_r. The event $H_r^{U_r}$ joins their first and last centers inside that component. The short surrounding loops supplied by $E_r^{U_r}$ attach these joins to ∂K; consecutive marked squares require just one such loop. These are the paths constructed in [3, proof of Lemma 5.12, Step 2(a)]. Each has length strictly smaller than $D(\text{across } A_{2r,3r}(z_0))$, and their union separates the inner from the outer boundary in $\mathbb{C} \setminus K$. This assertion concerns paths and finitely many square centers, so it remains valid in the presence of singular points.

A root geodesic to a point outside the exposed neighborhood meets ∂K only once. Read it backwards. If it enters $B_\epsilon(x) \setminus K$, it must first meet one of the shortcut paths and then cross $A_{2r,3r}(z_0)$ before reaching K. Its remaining length at that shortcut contact is therefore at least the radial crossing cost plus τ. Following the shortcut to ∂K and then a root geodesic costs strictly less, because every frontier point has root distance τ. This contradicts geodesicity. The argument uses no boundedness of K and no filling toward infinity.

Finally suppose $B_\epsilon(x)$ separates a frontier arc I from the finite target in $\mathbb{C} \setminus K$. The target component after the neighborhood has been swallowed is contained in the target component of its complement. Any path from that component to I, before its first contact with K, must enter the separating ball. The same is true starting at its frontier, since the separating ball lies strictly inside the exposed neighborhood. The point-shield conclusion therefore excludes geodesics through I to the later filled frontier as well as to points beyond it.

There is a minor normalization issue in this conditional argument. For each deterministic dyadic $\rho > 0$, first use the representative $h - h_\rho(p)$. Begin the exploration after the stopping time when $\overline{B}_{2\rho}(p) \subset K_u$. The pinning circle and a neighborhood of it have then been revealed, so the preceding conditional innovations are zero-boundary GFFs without a remaining circle-average constraint.

We prove the result for rational levels in each of these countably many representatives. Almost surely, every positive filled ball contains a Euclidean neighborhood of its root [3, Lemma 3.3]. A sufficiently small ρ therefore gives the conclusion at any prescribed positive levels, after their common Weyl rescaling. Rational interpolation of the levels removes the rational restriction. In the rest of the proof we work in one of these representatives and absorb the iteration immediately if its pin has not been revealed by time t.

On $\{s < T\}$, there are almost surely random $d > 0$ and $M < \infty$ such that, for every $u \in [t, s]$,

$$B_d(p) \subset K_u, \qquad B_d(z) \subset \mathbb{C} \setminus K_u, \qquad \partial K_u \subset B_M(p). \tag{A.6}$$

For the last assertion use precompactness of the metric net segment [3, Proposition 3.12] and continuity of the identity to the Euclidean plane. The first two follow from nesting and the openness of the appropriate complementary components. Lemma 5.11 of that paper, a union bound on a deterministic compact grid, and Borel–Cantelli give

$$R_\epsilon(\partial K_u) \le C\epsilon^{1/2} \tag{A.7}$$

simultaneously in $u \in [t, s]$ for all sufficiently small dyadic ϵ. Likewise, the simultaneous estimate [3, Lemma 3.11] and Borel–Cantelli give a deterministic $\beta > 0$ and a random positive cutoff such that

$$\operatorname{dist}(\partial K_u, \partial K_v) \ge c(v - u)^{1/\beta} \quad (t \le u < v \le s) \tag{A.8}$$

when $v - u$ is below the cutoff. To pass from dyadic relative increments to arbitrary increments, choose a dyadic δ between $(v - u)/(2v)$ and $(v - u)/v$ and use $K_u \subset K_{(1-\delta)v}$. The root and target exclusions in the published estimate are absent under (A.6) at sufficiently small scales. For nested filled sets, the frontier separation in (A.8) implies inclusion of the corresponding Euclidean neighborhood. Thus

$$\sigma_\epsilon(u) - u \le C\epsilon^{\beta/2} \tag{A.9}$$

whenever the right endpoint is at most s and ϵ is below the cutoff.

Exhaust $\{s < T\}$ by countably many events E on which all constants and cutoffs in (A.6)–(A.9) have prescribed deterministic bounds, including a revealed pin at time t. Fix one such E. We do not condition the iteration on E.

Take any finite $\mathcal{F}_t$-measurable partition of ∂K_t into half-open arcs. At a current stopping level τ, let n be the number of its nonempty image arcs under leftmost propagation. For sufficiently large n, the separation estimate gives at least $n/2$ arcs separated from z by balls of a common dyadic radius

$$\epsilon_n \asymp C(E)n^{-1/2}.$$

Apply (A.3) to these arcs and expose to their common level $\sigma_{\epsilon_n}(\tau)$. Independence between arcs is unnecessary: their conditional expected number of unsuccessful events is at most $Cn\epsilon_n^\alpha$. Hence, for all $n \ge N_0(E)$, the conditional probability that at least $n/4$ arcs disappear is at least $3/4$. On E the time increment is at most

$$b(n) = C(E)n^{-\kappa}, \qquad \kappa = \beta/4 > 0.$$

For the unconditional iteration expose only until $\sigma_{\epsilon_n}(\tau) \wedge (\tau + b(n))$. Absorb with next count zero if the neighborhood has not been swallowed by the latter time, or if a localization or radius cutoff visible at that time fails. If the neighborhood is swallowed, the shield events are measurable there and define the next count. All step times are thus finite. Any new absorption only improves

the reduction inequality. On the localized event E, no absorption occurs while the sum of the allowed time increments remains less than $s-t$, by the deterministic radius and time bounds on E. This keeps the iteration unconditional; E is used only to identify its consequence.

Since absorption only decreases the next count, we retain

$$\mathbb{P}\{n_{\text{next}} \le 3n/4 \mid \mathcal{F}_\tau\} \ge 3/4, \qquad n \ge N_0(E). \tag{A.10}$$

Stop when $n \le N$, where $N \ge N_0(E)$. In each interval $[2^j N, 2^{j+1}N)$, a bounded number of reductions by $3/4$ suffices to leave the interval. By (A.10), the expected number of steps in that interval is bounded by an absolute constant. Therefore

$$\mathbb{E} \sum_{\text{steps before stopping}} b(n) \le C(E) \sum_{j\ge 0} (2^j N)^{-\kappa} \le C'(E) N^{-\kappa}.$$

If N_s is the number of initial arcs met by leftmost geodesics to ∂K_s, Markov's inequality gives

$$\mathbb{P}(E \cap \{N_s > N\}) \le \frac{C'(E)}{s-t} N^{-\kappa}. \tag{A.11}$$

This bound is uniform in the number and mesh of the initial arcs.

Use nested measurable partitions whose maximal Euclidean arc diameter tends to zero; for example, use a boundary parametrization from the normalized Riemann map of the target component. If more than N distinct points of ∂K_t are met, every sufficiently fine partition has more than N met arcs. Fatou's lemma and (A.11), followed by $N \to \infty$, show that only finitely many points are met. An extremal prefix to each such point is unique, so there are only finitely many prefixes. Repeat the argument for rightmost geodesics.

Finally take the countable intersection over rational targets, rational levels, and the pinning representatives. For an arbitrary target w and $s < D(p,w)$, choose a rational z in the component of $\mathbb{C} \setminus \overline{B}_s^D(p)$ containing w. Then $K_u^z(p) = K_u^w(p)$ for every $u \le s$. Inserting rational levels strictly between t and s, and restricting extremal geodesics, proves the simultaneous assertion. □

Lemma A.3 (Protected subsegments of a typical geodesic)**.** *Fix deterministic distinct points p, q, and let $P : [0, D(p,q)] \to \mathbb{C}$ be their almost surely unique geodesic. Almost surely, for every $0 < u < v < D(p,q)$ there are Euclidean neighborhoods U of p and V of q such that every geodesic with one endpoint in U and the other in V contains $P|_{[u,v]}$. The assertion concerns endpoints joined by a finite geodesic. It does not assert a finite metric diameter for U or V.*

Proof. Include p, q in a fixed countable Euclidean dense set $\mathcal{R}$ of deterministic points. Intersect the full-probability events for Lemma A.2 with roots in $\mathcal{R}$ and for uniqueness of all geodesics between pairs in $\mathcal{R}$.

We first prove a one-endpoint assertion. Let $x \in \mathcal{R}$, and suppose that the geodesic $P_{x,z}$ from x to a nonsingular point z is unique. For every Euclidean neighborhood W of z, there is a neighborhood W' of z such that every geodesic from x to a point of W' agrees with $P_{x,z}$ outside W.

To prove this, let $z_n \to z$ and choose geodesics P_n from x to z_n. By [3, Lemma 3.2], choose surrounding rectifiable loops J_n with shrinking Euclidean diameter, enclosing z_n, such that

$$D(z, J_n) + \operatorname{len}(J_n; D) \longrightarrow 0.$$

The loops can be chosen from a sequence about z by taking its index sufficiently slowly. Let w_n be the first contact of P_n with J_n. Comparing with a path through z and then along J_n gives

$$D(x, w_n) \le D(x, z) + D(z, J_n) + \operatorname{len}(J_n; D).$$

Lower semicontinuity gives the reverse limiting inequality, so $D(x, w_n) \to D(x, z)$. Thus the prefixes to w_n converge uniformly in the Euclidean metric to $P_{x,z}$ by [3, Lemma 5.5] and uniqueness. After the first contact with J_n, the path cannot leave a fixed neighborhood of z for large n: otherwise it pays the positive crossing cost of a fixed annulus before returning to J_n, whereas a replacement along J_n has length tending to zero. Positivity of that crossing cost follows from lower semicontinuity on two disjoint compact boundary curves and positivity off the diagonal.

Choose $0 < t < s < D(x, z)$ such that the tail of $P_{x,z}$ after t lies in W. For large n, z_n is in the component containing z outside the level-s ordinary ball, and $P_n(s) \in \partial K_s^z(x)$. The extremal geodesics to $P_n(s)$ converge to $P_{x,z}|_{[0,s]}$: their lengths are exactly s, the limiting endpoint is $P_{x,z}(s)$, and that prefix is unique. Lemma A.2 places their level-t points in a fixed finite set. Consequently both level-t points eventually equal $P_{x,z}(t)$. Their sandwich property forces the same equality for $P_n(t)$, and uniqueness of the prefix gives agreement on $[0, t]$. The preceding tail estimate gives the one-endpoint assertion.

We use this assertion to construct two barriers. Choose $u' < u < v < v'$ strictly inside the parameter interval of P. A Schoenflies straightening of the simple arc supplies small disjoint Jordan neighborhoods of $P(u')$ and $P(v')$ in which the reference is a proper crosscut. By applying the one-endpoint assertion first with root p near $P(v')$, and then with a rational root on one chosen side near $P(u')$, we obtain a unique geodesic between two points of $\mathcal{R}$ on that side which agrees with P outside the two neighborhoods. Do the same on the other side. Their off-P portions stay on the chosen sides: they are confined to the two neighborhoods, and two separated intersections with P followed by a departure would contradict uniqueness of the corresponding subsegment. The neighborhoods may be chosen small enough that the two barriers contain $P|_{[u,v]}$.

The planar argument of [6, Lemma 3.11] now applies. In a sufficiently thin neighborhood of P, a path passing from its entrance to its exit can avoid the common middle only by entering a side region and then leaving it. A geodesic cannot do this: it would meet one of the unique barrier geodesics twice and differ from it between the contacts. Thus every geodesic sufficiently close to P and passing from its entrance to its exit contains $P|_{[u,v]}$.

Finally let $p_n \to p$ and $q_n \to q$, with geodesics Q_n between them. Surround p and q by disjoint shrinking rectifiable loops enclosing the respective endpoints and having lengths and access distances tending to zero. Truncate Q_n between its first contact with the first loop and its last contact with the second. The omitted portions are in the shrinking components. The remaining length is at most $D(p, q)$ plus the two loop lengths and access distances. Its lower limit is at least $D(p, q)$ by lower semicontinuity. Compactness and uniqueness therefore give Euclidean convergence of the truncated paths to P. The two barriers force the required subsegment for all large n. A sequence contradiction supplies the neighborhoods U, V. □

B Geometric estimates at every positive parameter

Fix $a > 0$, let h be a whole-plane GFF, and write $D = M_{ah}^a$ and $q = aQ(a) > 0$. A shift $h - f$ changes lengths by the factor e^{-af}. Every assertion below is for this fixed parameter. Null sets may be intersected over the countable deterministic domains, points, scales, and smooth shifts used in the proof. We never assume that an ambient Euclidean compact set is compact for D.

B.1 Compactness and short loops

Lemma B.1 (Ordinary local compactness)**.** *Let K be a Euclidean compact set and let d_K be the intrinsic length metric obtained by restricting D-rectifiable curves to K. Suppose $P_n \subset K$ joins x_n*

to y_n, its lengths are bounded, and

$$\mathrm{len}_D(P_n) - d_K(x_n, y_n) \longrightarrow 0, \qquad d_K(x_n, y_n) \longrightarrow d_K(x, y) < \infty, \qquad x_n \to x, \quad y_n \to y.$$

The endpoint convergence in the last display is Euclidean. After arclength parametrization and passage to a subsequence, P_n converges uniformly in the Euclidean metric to a d_K-geodesic from x to y. The same conclusion holds in an open comparison domain if the paths and their limits remain in a fixed compact subset of that domain.

Proof. Extend the length parametrizations constantly to a common bounded interval. Positivity off the diagonal and lower semicontinuity imply that, for every $\epsilon > 0$,

$$\inf\{D(u, v) : u, v \in K,\ |u - v| \geq \epsilon\} > 0.$$

Consequently these parametrizations are Euclidean equicontinuous. Arzelà–Ascoli gives a uniform limit in K. Lower semicontinuity, applied to each finite partition, bounds its length by the limiting lengths. This gives a minimizing curve because its endpoint distance is the assumed limit. No convergence of distances follows merely from Euclidean endpoint convergence. This is the local version of the argument in [3, Lemma 5.5]. □

Lemma B.2 (Metric convergence along a rectifiable limit)**.** *Let $Q_n : [0, 1] \to \mathbb{C}$ have a common finite D-Lipschitz constant and converge uniformly in the Euclidean metric to a nonconstant constant-speed D-geodesic Q. Then $Q_n \to Q$ uniformly in D. The same conclusion holds for a rectifiable limit which is injective on a neighborhood of each interior parameter and has the stated accessible loops there, provided the endpoint convergence is in D.*

Proof. First fix $t \in (0, 1)$ and a small $\delta > 0$ with $t \pm \delta \in (0, 1)$. A sufficiently small surrounding loop J about $Q(t)$ has $Q(t - \delta), Q(t + \delta)$ outside it and has arbitrarily small sum of access distance and perimeter. For all large n, $Q_n(t)$ is inside J, whereas $Q_n(t \pm \delta)$ are outside it. Consequently Q_n meets J within parameter distance δ of t. If C is the common Lipschitz constant, then

$$\limsup_{n\to\infty} D(Q_n(t), Q(t)) \leq C\delta + D(Q(t), J) + \mathrm{len}_D(J).$$

First make the last two terms arbitrarily small and then let $\delta \downarrow 0$. This proves pointwise convergence at interior parameters. The common Lipschitz bound, together with the D-continuity of Q, promotes convergence on a finite parameter net to uniform convergence on compact subintervals of $(0, 1)$. The same bound controls the two remaining endpoint intervals by choosing one nearby interior net point. Thus for a nonconstant geodesic limit no separate endpoint hypothesis is necessary. In the second assertion, use the assumed D convergence at the endpoints for these intervals. □

Lemma B.3 (Short loops and access)**.** *Fix a bounded ordinary comparison domain and a smaller compact observation set. For every $\chi \in (0, 1)$ there is $\theta > 0$ such that, almost surely at all sufficiently small scales r, the following holds simultaneously for ordinary internal geodesics of length at most M: if a geodesic visits $B_r(z)$ and both its endpoints and the domain boundary are outside $B_{Cr^\chi}(z)$, there is a separating rectifiable loop in $A_{Cr, Cr^\chi}(z)$ of length at most CMr^θ. Constants and the scale cutoff may depend on a fixed compact localization. The assertion permits random geodesics and random visited points.*

If both endpoints of a geodesic are outside a separating loop J, its portion between its first and last contacts with J has length at most $\mathrm{len}_D(J)$. In particular, two nearby points on two uniformly protected geodesics can be joined through a common short loop. On compact families of such protected interiors, Euclidean convergence therefore implies D-convergence, uniformly.

Proof. The simultaneous loop estimate is [4, Lemma 2.13] with zero perturbation. Its comparison domain U is fixed and deterministic. For each domain in our countable catalogue, apply the estimate on a smaller compact observation set, take a dyadic union, and use Borel–Cantelli. Intersect these countably many probability-one events. On the resulting event, the visited point, the endpoints, and the internal geodesic may all be random. Every use for an internal geodesic below is localized inside one of these deterministic buffered domains; no simultaneous statement over arbitrary random domains is being inferred from the cited lemma. The passage bound follows by replacing the geodesic portion with an arc of J. Each of the two points then has access to J at cost at most its length, and a loop arc joins the two access points. Constants absorb the finite changes in annular ratios. Arbitrarily short accessible loops at a single nonsingular point are supplied by [3, Lemma 3.2]. □

The same access argument applies to a filled frontier. If $x \in \partial K_s^z(p)$ and a loop surrounds x but neither p nor z, take the first contact of a p-to-x geodesic with the loop. Its remaining length is at most the loop length. Otherwise every loop point has root distance strictly less than s, so its interior is swallowed by $K_s^z(p)$, contradicting $x \in \partial K_s^z(p)$. Combined with [3, Corollary 3.7], this gives the same uniform access estimates on frontier portions separated from their root and target.

B.2 Protected references and local forced patterns

A *protected X* consists of two unique geodesics whose common nonconstant segment is a segment of a reference geodesic. Their two initial arms lie on opposite local sides of the reference, and their two final arms do likewise. A sufficiently thin topological neighborhood of the reference has the following property: a minimizing path traversing that neighborhood from entrance to exit contains the common segment. Indeed, avoiding it requires entering and leaving a side region, hence meeting one of the unique barriers twice and differing from it between the contacts. This is the planar separation argument of [6, Lemma 3.11].

Lemma B.4 (Protected reference with a flat lowering)**.** *Let P be a unique D-geodesic from u to v, of length $d > 0$, with a protected X on a strict interior subsegment. There is $\delta > 0$ such that the following holds for every continuous $0 \le f \le A$ which equals A on P and both barriers. If a D_{h-f}-geodesic Q has points x, y satisfying*

$$D_{h-f}(u,x), D_{h-f}(v,y) \le \delta e^{-aA} d,$$

then the portion between those points contains the protected segment, with its specified sidedness. The conclusion is uniform in A and in the values of f off the indicated plateau.

If the reference is a strict subsegment of a geodesic between one of the countably many fixed deterministic endpoint pairs, the barriers and four short endpoint loops with connectors can be chosen so that, upon also requiring $f = A$ on these loops and connectors, the entire query avoids fixed Euclidean neighborhoods of all four barrier endpoints whenever its full endpoints lie outside the local carrier.

Proof. The Weyl floor is $D_{h-f} \ge e^{-aA} D$. The reference and each barrier attain this floor and remain unique geodesics. For the oriented segment $Q[x,y]$,

$$D(u,x), D(v,y) \le \delta d, \quad \operatorname{len}_D Q[x,y] \le (1+2\delta)d, \quad D(x,y) \ge (1-2\delta)d.$$

If no suitable δ existed, let it decrease to zero. The paths have bounded length and lie in a common Euclidean bounded metric ball about u. Compactness and uniqueness identify every limit with P.

The two barriers remain unique in every changed metric, so the planar separation just described forces the protected segment, a contradiction.

For the last assertion, use the stated deterministic-endpoint origin of the reference and first choose the protected segment in a very short metric-time interval of the strict middle of P. Lemma A.3, applied successively from rational roots on either side, gives the barriers. They may have total lengths at most ϵd, while their common segment has distance at least cd along P from u, v, where $c > 0$ and $3\epsilon < c/2$. Here the rational endpoints can be chosen close in D, not just in Euclidean distance: a small accessible loop about either selected reference point has points strictly on both local sides, and rational D-density followed by the metric-to-Euclidean modulus gives rational approximations on each chosen side. The resulting barrier length is bounded by the two endpoint access costs plus the short reference interval.

For every barrier endpoint p_i, choose a loop J_i surrounding it and a connector from p_i to J_i, with combined length at most ϵd. They lie in a small open neighborhood of p_i disjoint from P. Choose a closed Euclidean ball B_i about p_i inside J_i. All these curves can be included in the plateau. Choose the capture neighborhood of P disjoint from the B_i and then decrease $\delta < c/2$.

If Q visits $w \in B_i$, it meets J_i on both sides of w. Loop replacement, followed by the connector and the short barrier, gives, for a point z of the captured stem,

$$D_{h-f}(w, z) \le 3\epsilon e^{-aA} d.$$

The captured interval extends at least $(c - \delta)e^{-aA}d$ before and after z. Since Q is a geodesic and $3\epsilon < c - \delta$, its visit to w must belong to that interval, which lies in the neighborhood disjoint from B_i. This contradiction proves whole-query avoidance. □

We need this configuration as a local event, with a deterministic tolerance. Start with a unique geodesic between fixed deterministic points. Select a crossing of a deterministic horizontal or vertical half-annulus and construct the barriers strictly inside it. A countable exhaustion fixes the rational barrier endpoint labels, strict Euclidean buffers, a bound for the ratio of the longer reference length to the crossing length, and the short endpoint loops, while retaining positive probability.

Before fixing δ, choose $S > 3$ and require that there exists $t \in [3r, Sr]$ such that

$$D(\text{around } A_{t,2t}) \le \tfrac{1}{8} D(\text{across } A_{2t,3t}).$$

Annular recurrence gives a deterministic S for which this has probability sufficiently close to one to meet the preceding positive-probability event; see [4, Lemma 5.2 and the proof of Lemma 5.7]. A path between the inner reference endpoints which exits B_{3t} pays a crossing cost and can be shortened along the enclosing loop by a fixed positive amount. Thus all reference paths and sufficiently accurate near minimizers are determined inside B_{3Sr}, regardless of the exterior completion. The compactness proof above is now local. A dyadic cutoff retains $\delta \ge \delta_0 > 0$ on an event of positive probability.

Choose $\lambda > 0$ with $C\lambda < \delta_0$. Only then choose the endpoint mesh. The proof of [4, Lemmas 5.4–5.6], using Lemma B.3 along the longer reference, gives deterministic endpoint disks and the estimates

$$\left|\{x \in \partial B_s(u_0) : D(x, u; \overline{B}_s(u_0)) > \lambda d\}\right| \le (\lambda/2)s,$$

and the corresponding estimate at v. The mesh event has probability tending to one as the mesh tends to zero; hence it preserves positive probability. In that calculation the error is bounded by a constant times $(2\sqrt{t})^\theta$, which tends to zero. This is the error needed here. No smaller power is required. All tests are invariant under adding a scalar to h.

Lemma B.5 (A local forced-X comparison)**.** *For every $p_0 < 1$ sufficiently close to one there are fixed radii $0 < b_- < b_+ < \lambda_3 < \lambda_4$, local events $E_{z,r}$ of probability at least p_0, and a target-independent local path predicate $C_{z,r}$ with the following properties. The events depend on a fixed-ratio annular observation modulo constants. The predicate is determined by the same local field and the query stopped at its last exit from a fixed enlarged carrier. It asserts that the query contains a protected X of diameter $O(r)$ and avoids four endpoint balls of radii at least cr throughout that carrier.*

For every deterministic pair of endpoints outside the carrier, with their almost surely unique geodesic P, one has

$$\mathbb{P}(C_{z,r}(P),\ P \cap B_{b_+ r}(z) \neq \varnothing \mid h|_{\mathbb{C}\setminus B_{\lambda_3 r}(z)}) \geq \Lambda^{-1}\mathbb{P}(E_{z,r},\ P \cap B_{b_- r}(z) \neq \varnothing \mid h|_{\mathbb{C}\setminus B_{\lambda_3 r}(z)}). \quad \text{(B.1)}$$

The pinning support is outside the fresh hole. The constants are deterministic, and the assertion is uniform in center and scale.

Proof. We specify the change in the deterministic forcing argument of [4, Sections 5.4–5.10]; its strict-shortcut conclusion is not used. Put positive-probability blocks just constructed into the finite chain of half-annuli and endpoint disks from Section 5.4 of that paper. In each good half-annulus, include the reference, barriers, and endpoint loops in the maximum plateau. Their containment is available before choosing the fine endpoint mesh. The disk-observation recurrence [4, Lemma 2.3] makes the required supply of good blocks on each of two long circular arcs arbitrarily probable by taking their relative scale sufficiently small. All remaining annular, lower-modulus, and geodesic-boundary regularity tests in that construction have probability tending to one. The finite smooth profiles are deterministic after their catalogue labels have been specified.

Here are the deterministic estimates, including the hypotheses whose use changes. Write $N_r = r^q e^{a h_r(z)}$. Let $U \Subset V \Subset A_{r,3r}(z)$ be the separating plateau and its support, and let $0 \leq f \leq A$ be A on U and zero off V. Retain deterministic bounds

$$D(V, \partial A_{r,3r}(z)) \geq c_0 N_r, \quad D(\text{around } A_{3r,4r}(z)) \leq C_0 N_r, \quad D(\text{around } U) \leq L_0 N_r.$$

Put $J = 2 + \lceil C_0/c_0 \rceil$ and fix $b = c_0/(4J)$. This substantial-excursion threshold is chosen before the pre-support contact-separation scale. Set $\kappa = a(Q+3)$ and retain the lower-modulus test

$$D(x, y; A_{r,4r}(z)) \geq \left(\frac{|x-y|}{r}\right)^{\kappa} N_r \quad \text{(B.2)}$$

for all x, y in the compact annular region used below with $|x-y| \leq \lambda^{-1}\eta r$, where η is the small comparison scale. This is condition 6 of [4, Section 5.5]. For each fixed $a > 0$, $\kappa > a(Q+2)$, and the reverse-Hölder estimate in [4, Lemma 5.14] makes this test arbitrarily probable as $\eta \downarrow 0$. We use this exponent in both the confinement and lower crossing estimates.

Retain also the strengthened plateau-circuit event from condition 7 of [4, Section 5.5]: there is a rectifiable circuit $\Pi \subset U$ separating the two boundary components of U, with

$$\operatorname{len}_D \Pi \leq L_0 N_r, \qquad \sup_{x \in \partial_{\text{out}} U} \operatorname{dist}_{\text{Eucl}}(x, \Pi) \leq a_6 r. \quad \text{(B.3)}$$

Here $a_6 > 0$ is chosen after the outside-excursion confinement scale, and the transition width defining V is at most $a_6 r$. For each of the finitely many deterministic annular templates U, choose a thin fixed boundary neighborhood inside U whose inner boundary is within $a_6 r$ of every point of $\partial_{\text{out}} U$. A separating circuit in that neighborhood has finite length almost surely, by the annular

path estimates and a finite rectifiable network. Taking a deterministic length cutoff, and a finite union over templates, makes (B.3) arbitrarily probable. This is the construction of [4, Lemma 5.15], which already applies at every positive parameter.

A first or last contact with V is within $2a_6r$ of Π. The pre-support short-loop estimate, used with an inner radius $4a_6r$, gives a loop of length at most $Ca_6^\theta N_r$ about each contact. It meets both the query and Π: the latter has a point inside its inner boundary and, being a macroscopic separating circuit, a point outside its outer boundary. Replacing the query between these two contact loops by their arcs and a plateau arc of Π therefore gives

$$\sigma - \tau \le Ca_6^\theta N_r + e^{-aA} L_0 N_r. \tag{B.4}$$

Choosing a_6 small and then A large makes this less than the unchanged crossing cost of every outside-support excursion of the prescribed confinement diameter. This supplies the precise access hypothesis and length bound used in the confinement argument.

The choices have the following noncircular order. First fix the protected-block tolerance and endpoint mesh, then the macroscopic radial crossing and surrounding-loop cutoffs, the threshold b, and the uniform pre-support short-loop cutoff. These determine a positive separation scale for the first and last support contacts. Next choose the block density ρ, the outside-excursion confinement scale, and the smaller contact-loop scale a_6. Choose the boundary-following circuit cap L_0 only after these choices. Then choose the lowering amplitude A to satisfy (B.4) and the original-hit saving inequality. Only after A is fixed choose the support transition width and impose the thin-boundary occupation test for shifts in $[0, A]$, on each endpoint circle as well as the plateau boundary. Retain also the zero boundary-length contact conclusion used when taking closures of the accessible sets; these are the tests used in [4, Lemmas 5.11 and 5.32]. Finally choose the Gaussian score cutoff for the resulting finite smooth catalogue. The boundary-occupation exponent is allowed to depend on A; no estimate uniform over an amplitude tending to infinity is used.

A D_{h-f}-geodesic excursion with support-contact intrinsic distance at least bN_r cannot enter the inner disk: the lowered plateau circuit is cheaper than the unchanged inner annular crossing. The pre-support loop bound gives a positive Euclidean separation between its two support contacts. Two contact loops and the lowered plateau circuit bound its total length, and the unchanged lower modulus then confines every outside-support excursion to a sufficiently small Euclidean diameter. These are the proofs of Lemmas 5.19–5.22 of the cited paper and involve only D.

The circular chain and the good blocks on its two long arcs now force the path to cross the enlarged endpoint disks of one good block, by the four-component separation in Lemmas 5.23–5.25. Let u, v be the reference endpoints of that block and $d = D(u, v)$. In an endpoint disk, call a boundary point accessible if it can be reached from the exterior side without meeting the query. If accessible points have total length greater than $3\lambda s$, the most-boundary estimate and the thin-boundary occupation test give a path from u to the query of changed length at most $2\lambda e^{-aA} d$. The radial segment used to pass between the two endpoint circles is chosen away from the query; it is a topological device and is not assigned a finite metric length.

Otherwise, most of the endpoint-circle arc is covered by disjoint blocked arcs I. Each is spanned by an outside-support query excursion. Discard the arcs smaller than the support transition width, using the thin-boundary occupation test. For every remaining arc with an accessible metric endpoint, a short outer loop joins the query to the endpoint access path and gives times $s_I < t_I$ with

$$D_{h-f}(P(t_I), u; A_{r,4r}(z)) \le 2\lambda e^{-aA} d + C(|I|/r)^{-1/4}(t_I - s_I), \qquad t_I - s_I \ge c(|I|/r)^{a(Q+3)} N_r. \tag{B.5}$$

The annular test used here is explicitly

$$D(\text{around } A_{2\rho,3\rho}(y)) \le (\rho/r)^{-1/4} D(\text{across } A_{\rho/4,\rho/2}(y)).$$

Both orientations of the fixed-ratio annular comparison can be required using [4, Lemma 2.8]. Let R be the endpoint-circle radius. The parent boundary arc has endpoint chord at least cR for a fixed $c > 0$, so its angular length is at most $2\pi - 2\arcsin(c/2)$. Hence a deterministic K satisfies $|I| \le K|y_I - y'_I|$ for every subarc I. Outside-support confinement bounds this chord. Choose its scale smaller by K, and then small enough relative to $t\rho$ that each blocked arc has $|I| \le R$. Thus

$$|y_I - y'_I| = 2R\sin\big(|I|/(2R)\big) \ge |I|/2.$$

The first exit from $B_{|I|/2}(y_I)$ therefore precedes the second endpoint of the blocked excursion. The four-component separation puts the crossing segment outside both the endpoint disk and the remaining support. Its changed and ordinary lengths agree. Applying (B.2) gives directly

$$\Delta_I := t_I - s_I \ge D(\text{across } A_{|I|/4,|I|/2}(y_I)) \ge \left(\frac{|I|}{4r}\right)^{\kappa} N_r.$$

The outer endpoint comparison disk is enlarged by the same fixed factor K, so all loops and access paths remain admissible.

For $|I| \in [2^{-k-1}r, 2^{-k}r)$, this lower bound and the contact-time upper bound give

$$2^{-\kappa(k+3)} N_r \le \Delta_I \le N_r.$$

There are at most $\lceil \kappa(k+3) \rceil + 2$ dyadic time bins. The arc-length pigeonhole argument supplies at least $c2^{k/2}t\rho$ disjoint comparison disks at some such scale, and hence a subcollection of

$$N \ge c_a (k+1)^{-1} 2^{k/2} t\rho$$

ordered disjoint time intervals whose increments lie in $[v, 2v]$. The first and last endpoint-access estimates and geodesicity give

$$(N-1)v \le C\lambda e^{-aA} d + C2^{k/4} v.$$

Choose the confinement scale after t, ρ so that the lower bound for N is at least $2C2^{k/4} + 2$ for every allowed k. Then

$$v \le C_a \lambda \frac{(k+1)2^{-k/2}}{t\rho} e^{-aA} d.$$

Substitution in (B.5) yields

$$D_{h-f}(P(t_I), u; A_{r,4r}(z)) \le C_a \lambda \left(1 + \frac{(k+1)2^{-k/4}}{t\rho}\right) e^{-aA} d.$$

A further decrease of the same confinement scale bounds the parentheses by a fixed constant. The argument at v is identical. This proves the endpoint-closeness conclusion of [4, Lemma 5.26], with its auxiliary positive threshold replaced by $D(u, v)$. Lemma B.4 then gives the protected X and whole-query endpoint avoidance.

For completeness, an original hit produces the substantial changed excursion needed above. An original geodesic hitting $B_r(z)$ crosses a plateau loop twice. Replacing the intervening part with its lowered loop arc saves at least

$$(c_0 - 2e^{-aA} L_0) N_r.$$

Choose A so this saving exceeds $c_0 N_r/2$. The unchanged outer loop bounds the changed length between the first and last support contacts by $C_0 N_r$. Each intervening departure to the carrier boundary and return pays two unchanged radial crossings, so at most J carrier excursions meet the support. If each support-contact pair had ordinary endpoint distance below bN_r, replacing all such portions by ordinary almost-minimizers would increase length by at most $JbN_r + o(N_r) < c_0 N_r/2$, contradicting the saving. At least one pair therefore has ordinary, and hence intrinsic, endpoint distance at least bN_r. The changed geodesic hits $V \subset B_{3r}(z)$, which explains the enlarged hit radius in (B.1).

Finally fix the field outside the slightly larger fresh hole, and write μ for its conditional Gaussian law. For each deterministic catalogue profile put $T_i h = h - f_i$. Its Dirichlet pairing with the harmonic conditional mean vanishes. Write $\langle h, f\rangle_\nabla = \langle h, -\Delta f\rangle_\mu$ and $\|f\|_\nabla^2 = \langle \nabla f, \nabla f\rangle_\mu$. For every measurable set A_i one has

$$\mu(T_i A_i) = \int_{A_i} \exp\left(\langle h, f_i\rangle_\nabla - \tfrac{1}{2}\|f_i\|_\nabla^2\right)\, d\mu(h).$$

Retain $|\langle h, f_i\rangle_\nabla| \le M$ and $\|f_i\|_\nabla^2 \le E$ on the selected part, so this is at least $e^{-M-E/2}\mu(A_i)$. The finite family of score cutoffs has probability as close to one as required. Partition the original-hit event by its m catalogue labels; every translated part belongs to the changed-hit event B. Finite overlap gives

$$\mu(B) \ge m^{-1} e^{-M-E/2} \sum_{i=1}^{m} \mu(A_i),$$

which is (B.1). Dirichlet norms are unchanged by planar scaling.

To make the success predicate local, retain an inverse-shift event: the local field $g + f_i$ has the base reference configuration, and the observed g-query has its stated X incidence and endpoint avoidance. The enclosing local loop makes the base barriers globally minimizing. The common floor preserves global minimality and uniqueness after the shift. This is one predicate independent of the terminal label. The likelihood comparison is valid for each exterior; it is not a lower success bound uniform over arbitrary exterior harmonic functions. □

B.3 Root prefixes and the exploration filtration

Lemma B.6 (Unique strict prefixes from countably many roots)**.** *For every fixed countable deterministic set of roots, almost surely every geodesic from one of these roots has a unique strict prefix.*

Proof. Fix deterministic rational levels $0 < t < s$ and a rational filling target with $s < D(p, z)$. Write X for the finite set of level-t prefix endpoints for leftmost geodesics to the level-s frontier. For every such leftmost geodesic, use the sided rational approximations from Lemma A.1. The unique rational-target geodesics in the approximating sequence are also leftmost to their level-s contacts. Their level-t contacts therefore lie in X. They converge to the chosen prefix endpoint and hence eventually equal it, since X is finite. A rational-target geodesic through that prefix endpoint has a unique prefix to it: any alternative prefix could be concatenated with its remaining tail. Consequently each prefix endpoint in X has a unique root geodesic.

The prefix classes on the later frontier are connected cyclic arcs. A root geodesic reaching a point in the relative interior of a prefix endpoint class has the same prefix through t. To verify this, bracket its frontier point from both sides by frontier points in that class, and approximate the two corresponding extremal paths by unique rational-target paths. For sufficiently close approximations

their level-t contacts are the class's prefix endpoint, by the preceding finite-set argument. The unique rational-target paths cannot be crossed by a root geodesic with a different prefix, since splicing at the first crossing would contradict their uniqueness. The bracket therefore forces the stated prefix.

It remains to exclude the finitely many class endpoints Y for all geodesics which continue outside the level-s filled ball. The set Y is measurable from the filled ball and its revealed field. Fix a positive Euclidean neighborhood width d and truncate to $|Y| \le m$, a compact frontier, positive root and filling-target clearances, and a revealed pin. Apply the point shield at every $y \in Y$ with dyadic radius ϵ. Their total conditional failure probability is at most $Cm\epsilon^{\alpha}$. The common exposed radius tends to zero almost surely, by the good-annulus grid estimate. If a geodesic to a point outside the d-neighborhood passes through y, its continuation immediately enters $B_\epsilon(y) \setminus K_s$ and violates the shield for every sufficiently small ϵ. Letting $\epsilon \downarrow 0$ shows that the probability of any such geodesic is zero. Then exhaust m, the localization cutoffs, and the countable neighborhood widths d. Every point outside the closed filled ball has some positive Euclidean neighborhood, so the exclusion is simultaneous over all exterior targets.

Now let $P : p \to w$ be an arbitrary finite geodesic, and let r be a strict parameter time. Choose rational $r < t < s < D(p, w)$ and a rational filling target in the ordinary level-s complementary component containing w. Filling toward that rational point gives the same sets through level s. The preceding argument gives a unique prefix to $P(t)$. Any second prefix to $P(r)$ would concatenate with $P[r, t]$ to contradict this uniqueness. All exceptional events used above were indexed by countably many deterministic roots, rational levels, rational filling targets, and localization cutoffs. □

We now use only balls filled toward infinity, writing $K_t(p)$ and $\tau_r(p) = D(p, \partial B_r(p))$. The quantitative prefix-count estimate needed below is the published [3, Theorem 5.15]; finite-target confluence is not being used as a quantitative estimate.

An *eligible-root localization* fixes a compact observation buffer, a Euclidean scale $\ell > 0$, and deterministic upper and lower cutoffs for the following quantities: exploration times up to $\tau_{16\ell}$; the Euclidean neighborhood of the root contained in K_{τ_ℓ}; the positive separations between the required macroscopic exit times; fixed-scale circle-average comparisons; the lower metric-to-Euclidean modulus; and the uniform loop, good-annulus, and estimates for filled-frontier neighborhoods at all sufficiently small scales. It also requires every center of one absolute-origin candidate grid to have an E-good radius in each sufficiently small dyadic window $[\epsilon^{1+\nu}, \epsilon]$. For this last condition choose p_0 in Lemma B.5 close enough to one that annular recurrence and the polynomial grid union have summable failure probability. Include the resulting eventual scale cutoff in the localization. None of these future conditions is imposed in a conditional trial law.

For each deterministic root choose a deterministic small pinning circle inside its retained root neighborhood. Start the exploration after the circle and a buffer have been swallowed, and use deterministic time increments in the representative pinned on that circle. The delay is absent on eligibility. A cutoff for the difference from the macroscopic circle average compares these time units with the published estimates. The fresh holes then have ordinary zero-boundary innovations without an unrevealed pinning constraint.

Lemma B.7 (Two-loop frontier visit bound)**.** *Fix $R, d, c > 0$ and $\eta > 0$. Except on an event smaller than every power of ϵ, simultaneously for all s with $B_d(p) \subset K_s(p) \subset B_R(p)$ and all root geodesics P destined outside $K_s(p)$,*

$$\operatorname{area}\big(B_{c\epsilon}(P[s, \operatorname{len} P]) \cap B_{c\epsilon}(\partial K_s(p))\big) \le \epsilon^{2-\eta}.$$

The probability bound is uniform for deterministic roots in a fixed compact set. It does not require an upper bound on the total length of P.

Proof. Fix $0 < \zeta < \eta/4$. At each point of an ϵ-grid near the frontier, examine separated scales $10(c+1)\epsilon \le r \le 10^{-2}\epsilon^{1-\zeta}$. For any required probability power, annular recurrence [4, Lemma 2.1] gives, simultaneously at all grid points, a scale with a separating loop $J_1 \subset A_{r,2r}$ satisfying

$$\mathrm{len}_D(J_1) \le Cc_r, \qquad c_r = D(\text{across } A_{2r,4r}).$$

The constant C is chosen after the required probability power. An outer loop $J_2 \subset A_{4r,5r}$ and the crossing obey $\mathrm{len}_D(J_2) \le \epsilon^{-A}S$ and $c_r \ge \epsilon^A S$, simultaneously for the polynomially many center-scale pairs, where $S = e^{ah_1(p)}$. This follows from all positive loop moments, all negative crossing moments, Gaussian circle-average tails, and a union bound; see [4, Section 2.4]. No point-to-point positive moment is used.

Record the first time at or after s at which the path is within $2c\epsilon$ of the frontier, and thereafter record a new such visit only after exiting the disk of radius $\epsilon^{1-\zeta}$ about the previous recorded visit. Call a visit completed if this exit occurs before the path ends. Choose the good loops around a grid point within $O(\epsilon)$ of that visit. A root-to-nearby-frontier geodesic crosses both loops by time s, so every point of J_i has root distance at most $s + \mathrm{len}_D J_i$.

For a completed visit v_j, the query subsequently crosses J_1 and then J_2 before the specified disk exit. Let w_j be its first subsequent J_2 contact and put $\Delta_j = w_j - v_j$. Geodesicity gives

$$v_j - s \le \mathrm{len}_D J_1 \le Cc_j, \qquad c_j \le \Delta_j \le \mathrm{len}_D J_2.$$

The first inequality uses a subsequent J_1 contact and is asserted only for completed visits. Since $v_{j+1} \ge w_j$,

$$\sum_{i<j} \Delta_i \le v_j - s \le C\Delta_j.$$

The partial sums therefore grow geometrically. The simultaneous bounds $\epsilon^A S \le \Delta_j \le \epsilon^{-A}S$ limit the number of completed visits to $O(\log \epsilon^{-1})$. There is at most one remaining recorded visit, and after it the path stays in its recording disk. Thus every frontier-near visit is covered by these $O(\log \epsilon^{-1})$ disks. Enlarging by $c\epsilon$ gives area $O(\epsilon^{2-2\zeta}\log\epsilon^{-1}) \le \epsilon^{2-\eta}$ for small ϵ. The auxiliary constants may depend on the requested probability power; since the claimed failure event is fixed, this proves its superpolynomial bound. No total path-length bound is used. □

Lemma B.8 (Stable levels for all outgoing geodesics)**.** *There are positive exponents β, ϑ such that the following holds on any eligible-root localization. In a fixed macroscopic exploration interval take*

$$s_k = s_0 + k\epsilon^\beta, \qquad t_k = s_k + \epsilon^{2\beta}, \qquad 0 \le k < K \asymp \epsilon^{-\beta},$$

using the pinned time units described above. Except on a superpolynomially small event intersected with eligibility, at least $(1-\epsilon^\vartheta)K$ stages have stable candidate cells for every geodesic whose target lies beyond the prescribed later filled frontier. Stability means that all ordinary exterior geodesics from the root to first contact with the cell's larger closed obstacle have the same prefix through time s_k.

Proof. Let $\mathcal{G}_t = \sigma(K_t, h|_{K_t})$. At time t_k, partition its frontier into the arcs corresponding to the level-s_k leftmost prefix endpoints, and let Y_k be their endpoints. By [3, Theorem 5.15], for some $\alpha_g > 0$,

$$\mathbb{P}(\text{eligibility},\ \max_k |Y_k| > \epsilon^{-\omega}) = o(\epsilon^m) \quad \text{for every } m, \qquad 2\beta < \omega\alpha_g.$$

For a frontier endpoint y, consider the region cut off together with y from infinity by connected crosscuts of diameter at most δ. The deterministic enclosure [13, Lemma 4.14] puts this entire region behind one crosscut of diameter at most 16δ. Choose a frontier point adjacent to that crosscut as in the proof of Lemma 4.15 there. A point shield at that point blocks every path from the region to outside the exposed neighborhood. A shield centered at y itself would not suffice for a long fjord.

Set $\delta = \epsilon^{\kappa}$. The conditional point-shield failure is at most $C\epsilon^{\alpha_s \kappa}$ by [3, Lemma 5.12]. Union over Y_k and take $2\omega < \alpha_s \kappa$. The good-annulus grid bound gives a neighborhood of radius $O(\epsilon^{\kappa/2})$, and [3, Lemma 3.11] swallows it within time $O(\epsilon^{\kappa b_f/2})$ for a positive b_f. Choose $\beta < \kappa b_f/2$.

For an adapted trial, declare automatic success if the prefix-count bound fails or the neighborhood has not been swallowed by s_{k+1}; otherwise require all the shields. Its failure probability conditional on $\mathcal{G}_{t_k}$ is at most ϵ^{ω}, and it is measurable by $\mathcal{G}_{s_{k+1}}$. The bounded-increment martingale estimate gives failure of more than $\epsilon^{\vartheta} K$ stages with probability at most $\exp(-c\epsilon^{2\vartheta-\beta})$, where $0 < \vartheta < \min\{\omega, \beta/2\}$. On eligibility and the prefix-count event the automatic successes are absent. The geometric shield conclusion already quantifies over all outgoing geodesics.

We give the two replacements for the upper Euclidean-to-metric modulus in [13, Lemmas 4.13 and 4.16]. On the retained localization, write

$$|x - y| \le C D(x, y)^{\chi_-}, \qquad \chi_- > 0,$$

for the lower metric modulus on the relevant compact set. Fix a short-loop exponent $\chi \in (0, 1)$ and its positive length exponent θ. The frontier levels are bounded, their roots have a fixed Euclidean clearance, and the short-loop estimates hold simultaneously on this localization. Constants below absorb these fixed bounds.

First let P be a root geodesic to a point beyond the later filled frontier, and let I be the level-t_k arc through which it leaves. Suppose a connected witness of diameter ρ in $\mathbb{C} \setminus K_{t_k}$ joins $P[t_k, \operatorname{len} P]$ to $\partial K_{t_k} \setminus I$. With an arbitrarily small enlargement, the witness may be taken to be an arc whose only frontier contact is its latter endpoint x; the frontier is Jordan. A short loop about x, with inner radius larger than 2ρ, surrounds the witness and has length at most $C\rho^{\theta}$. A root-to-x geodesic reaches this loop by time t_k. Hence every loop point has root distance at most $t_k + C\rho^{\theta}$. Choose ρ so that this is strictly before the later frontier. The loop's interior is then swallowed, and the query target is outside it. A contact of P with the witness must therefore be followed by a loop contact; its time is at most $t_k + C\rho^{\theta}$. The portion of P from its frontier departure to that witness contact has Euclidean diameter at most $C\rho^{\theta\chi_-}$.

Adjoin this query portion to the witness. The planar crosscut argument in [13, Lemma 4.13] now puts the departure point together with an endpoint of I behind a crosscut of diameter at most $C\rho^{c_1}$, where

$$c_1 = \min\{1, \theta\chi_-\} > 0.$$

This is the only use of that lemma here: the time and diameter bounds which it originally obtained from an upper metric modulus have just been proved using actual loop contacts. A shielded stage excludes such a departure whenever $C\rho^{c_1} \le \epsilon^{\kappa}$.

Second let O be a candidate larger closed obstacle, disjoint from K_{t_k} and within $C\epsilon$ of its frontier. Choose a frontier point x within that distance and use a loop based at x whose inner radius contains the entire obstacle and its full carrier. Its outer radius is $C\epsilon^{\chi}$, and its length is at most $C\epsilon^{\theta}$. In particular the loop itself avoids the obstacle. The initial part of a root-to-x geodesic lies in K_{t_k}, so its access to the loop also avoids the obstacle. Every loop point therefore has an obstacle-avoiding access of length at most $t_k + C\epsilon^{\theta}$.

Let P' be any existing exterior minimizer to its first contact with O. Its last loop contact has exterior arclength at most $t_k + C\epsilon^\theta$, by comparison with the preceding access. The subsequent tail stays inside the loop. Furthermore P' reaches ∂K_{t_k} exactly at exterior time t_k and never returns: any later frontier contact could be shortened by a root geodesic of length t_k, which lies in K_{t_k} and avoids O. Its prefix through that frontier is an ordinary root geodesic, since any shorter root path of length at most t_k also avoids O. The lower modulus before the last loop contact and Euclidean confinement after it give

$$\operatorname{diam}_{\mathrm{Eucl}} P'[t_k, \operatorname{len} P'] \le C\epsilon^{c_2}, \qquad c_2 = \min\{\chi, \theta\chi_-\} > 0.$$

If the query hits the fresh hole inside O and P' leaves the frontier through a different prefix arc, join its final contact to the query inside the obstacle. Together with the exterior tail of P' this is a connected witness in $\mathbb{C} \setminus K_{t_k}$ of diameter $C\epsilon^{c_2}$ joining the query to the complement of its prefix arc. The first estimate contradicts the stage's shields when $\kappa < c_1 c_2$. All exterior minimizers must therefore leave through the interior of the query's prefix arc. Bracketing by its extremal paths and Lemma B.6 identifies their prefixes through s_k. Both exponents c_1, c_2 were fixed before κ; there is no circular choice of exponents. After choosing κ, choose ω and then β sufficiently small for the earlier prefix-count, shield, and neighborhood inequalities, and also so that every loop used above is swallowed before s_{k+1}.

Use a fresh hole strictly smaller than the closed obstacle. Every exterior comparison curve then lies in the open revealed complement of the fresh hole, including at its obstacle-boundary contacts. This is the locality buffer required in this assertion. □

Proposition B.9 (Abundance from eligible roots)**.** *Fix an eligible-root localization and an admissibly small window exponent $\nu > 0$. Uniformly in deterministic p, w in the observation buffer,*

$$\mathbb{P}\left(p \text{ eligible},\ w \notin K_{\tau_{4\ell}}(p),\ \text{no protected-X carrier along } P_{p,w} \text{ in } K_{\tau_{2\ell}}(p) \setminus K_{\tau_\ell}(p)\right) = o(\epsilon^m) \quad \text{(B.6)}$$

for every $m > 0$. Carrier radii lie in $[\epsilon^{1+\nu}, \epsilon]$, up to fixed constants. After a polynomial union over deterministic candidate roots, the corresponding conclusion holds simultaneously for every geodesic from each eligible root to every target outside $K_{\tau_{16\ell}}(p)$.

Proof. We give the filtration and counting argument. Run early stages between τ_ℓ and $\tau_{2\ell}$ and set

$$\mathcal{F}_k = \sigma(K_{t_k}, h|_{K_{t_k}}, P_{p,w}|_{[0,s_k]}).$$

The conditional shields of Lemma B.8 were proved in $\mathcal{G}$, not in this enlarged filtration. Their simultaneous pathwise consequence may now be used in $\mathcal{F}$.

Use a fixed grid of mesh $c\epsilon^{1+\nu}$ and the separated admissible radii. Candidate enlarged obstacles lie outside K_{t_k}, within $O(\epsilon)$ of its frontier. Remove any candidate whose full carrier contains the fixed terminal point w. This is an adapted, outside-measurable restriction. The final event $w \notin K_{\tau_{4\ell}}$ and the retained macroscopic separations give a positive Euclidean neighborhood separating w from every early explored set, so the restriction removes no candidate there for sufficiently small ϵ.

Count stable candidates satisfying respectively E and the smaller hit, or C and the larger hit, and denote these sets by Z_k^E, Z_k^C. Both hit balls lie inside the fresh hole, whereas stability is defined with the strictly larger closed obstacle. On a stable fresh-hole hit, the old prefix is determined by the outside field: it is the common prefix of every exterior minimizer to the obstacle. Thus the elementary Bayes lemma [13, Lemma 4.9] transfers (B.1) to $\mathcal{F}_k$. The event used in that lemma is the common event of stability and a fresh-hole hit. Equality of the two smaller hit radii is unnecessary.

The same obstacle-surrounding loop gives the requisite observability without assuming attainment of an exterior infimum. Its complete carrier and its interior are contained in a strictly earlier filled ball than $K_{s_{k+1}}$. Consider any finite exterior competitor to first obstacle contact, and split it at its last loop contact. If its prefix is more expensive than the known avoiding route to that contact, replace the prefix by the latter. Otherwise retain the prefix. The chosen prefix has length at most $t_k + C\epsilon^\theta$, so every point on it has root distance at most that value. The remaining tail is inside the loop. The resulting competitor is wholly in the revealed filled interior and has no greater cost. Thus the full exterior infimum equals the infimum restricted to this revealed region. For an attained exterior minimizer its original prefix already has this bound, so its entire curve is in the revealed region as well. This proves observability of both infima and the existing minimizing family. It does not claim that an infimum in an arbitrary open exterior domain is attained.

Let F_k be the event that these loop and access bounds hold, all full carriers are contained in the revealed ball before s_{k+1}, and

$$w \notin K_{s_{k+1}}. \tag{B.7}$$

It is measurable from the next revealed ball. Each event holds on the final localized event. Condition (B.7) is essential for adaptation: a geodesic to a point outside a filled ball cannot return to it after its frontier time, so its entire path through every contained carrier is determined by the prefix. A target inside a filled cavity would not give that implication. Thus

$$A_k = \{Z_k^E \neq \varnothing\} \cap F_k, \qquad B_k = \{Z_k^C = \varnothing\} \cap F_k \quad \text{belong to } \mathcal{F}_{k+1}.$$

The events can be expressed by strict rational length tests for paths and loops compactly inside the revealed domain. The definition of the event uses only the revealed data.

Lemma B.7 bounds the number of candidates met by the outgoing query by $\epsilon^{-2\nu-\eta}$ outside a superpolynomial exception, for every fixed $\eta > 0$. On a current measurable compact cutoff the total candidate count is polynomial. Conditional Markov therefore bounds its expected excess tail given $\mathcal{F}_k$ by a superpolynomial error, outside a superpolynomial exception. Summing the transferred comparison and truncating this count yields

$$\mathbb{P}(Z_k^C \neq \varnothing \mid \mathcal{F}_k) \geq c\epsilon^{2\nu+\eta}\mathbb{P}(Z_k^E \neq \varnothing \mid \mathcal{F}_k) - o(\epsilon^m) \tag{B.8}$$

for arbitrary m, on current cutoffs. If the terminal point is in the current filled ball, the candidate count can be set to zero; on the final event it is outside every early ball.

On the eventual good-radius cover, every outgoing query has an E-candidate. Indeed, choose its first point at Euclidean distance $3L\epsilon/2$ from the current filled set, with L larger than all obstacle ratios. The nearest candidate-grid point has distance in $[L\epsilon, 2L\epsilon]$ from the frontier. Its obstacle is disjoint from the ball for every admissible radius, and the query hits its smaller hit disk. The cover supplies an E-good radius. At a shielded stage this candidate is stable. Thus, on eligibility and apart from superpolynomial failure, at least $(1-\epsilon^\vartheta)K$ of the events A_k occur.

Apply the adapted concentration lemma [13, Lemma 4.18] first to A_k, with conditional threshold $1/2$. Except on superpolynomial failure, all but $4\epsilon^\vartheta K$ stages have $\mathbb{P}(A_k \mid \mathcal{F}_k) \geq 1/2$. Formula (B.8) then gives a success probability at least $c\epsilon^u$, where $u = 2\nu + \eta$. Apply the same concentration lemma to B_k, whose conditional probabilities are at most $1 - c\epsilon^u$ at these stages. Its error is at most $\exp(-c'\epsilon^{2u}K)$. We use B_k, rather than subtracting an unknown conditional probability of event failure. On eligibility all events hold, so the resulting deficit of B_k is a positive number of actual C-successes. Choose ν, η after β, ϑ with $4\nu + 2\eta < \beta$. This proves (B.6).

For the all-target conclusion run a separate late shield exploration between $\tau_{4\ell}$ and $\tau_{8\ell}$, with carriers swallowed before $\tau_{12\ell}$. Every geodesic to outside $K_{\tau_{16\ell}}$ meets a stable candidate. Its

deterministic center w lies outside the current late ball and hence outside $K_{\tau_{4\ell}}$. The query and the unique geodesic to w have the same prefix through a late prefix time at least $\tau_{4\ell}$. Union (B.6) over all polynomially many deterministic centers. The resulting early X carrier lies in $K_{\tau_{2\ell}} \setminus K_{\tau_\ell}$, and neither geodesic can revisit this set after $\tau_{2\ell}$. Their entire stopped paths in that carrier agree. Since C is a single target-independent predicate, its X configuration and whole-query endpoint avoidance transfer to the arbitrary target. There is no union over arbitrary targets and no conditioning on their later separation event. □

B.4 Deterministic clouds and finite multiplicity

Lemma B.10 (Polynomial clouds on geodesic interiors)**.** *For each fixed compact frame of bounded-length geodesics with a positive Euclidean clearance from their endpoints, there are deterministic finite sets H_r with $|H_r| = O(r^{-2})$ and a deterministic $\theta > 0$ such that, almost surely at every sufficiently small dyadic r, every frame point x has $u \in H_r$ with*

$$|x - u| \le Cr, \qquad D(x, u) \le CMr^\theta.$$

The failure at one scale has a positive polynomial bound. The exponent may be fixed over countably many compact frames. A corresponding assertion holds on filled-frontier portions with bounded level and polynomial Euclidean clearance from root and target.

Proof. We first obtain arbitrarily high moments for the minimum of a finite deterministic cloud's access costs. Choose m rational points $u_i \in B_{1/4}$ with disjoint private disks. In each private disk write $h = g_i + H_i$, where the zero-boundary fields g_i are independent, and the harmonic parts are independent of all these innovations. The normalization circle is outside the disks. Choose concentric disks $B_{\rho_i}(u_i) \Subset B_{2\rho_i}(u_i) \Subset B_{3\rho_i}(u_i) \Subset B_{4\rho_i}(u_i)$ inside the ith private disk, and choose a fixed rational point v_i in $B_{4\rho_i}(u_i) \setminus \overline{B}_{3\rho_i}(u_i)$. Let X_i be the g_i-internal distance from u_i to v_i in a fixed smaller buffered disk containing $\overline{B}_{4\rho_i}(u_i)$ and compactly contained in the private disk. By [4, Lemma 2.9], each X_i has a positive moment of every order $s < 2Q/a$. The same statement for a zero-boundary field follows from the Markov decomposition on a smaller disk: on an independent positive-probability event the harmonic part is bounded, and Weyl comparison transfers the full-field moment.

Independence gives

$$\mathbb{P}(\min_i X_i > t) \le C_s^m t^{-ms}, \qquad t \ge 1.$$

Thus the minimum has any prescribed finite moment if m is chosen sufficiently large. The exponential of the maximum harmonic fluctuation on the smaller disks has all moments. Include in the connecting network a separating loop in $A_{\rho_i,2\rho_i}(u_i)$. Every access path from u_i to v_i must meet this loop; this is why the second endpoint is chosen outside it. Choose the outer loop J in $A_{8,10}$, which is a subannulus of the stated $A_{8,16}$. From each private annulus make a finite chain of annuli of the same radius, with successive centers at distance $3\rho_i/2$, along a deterministic route to a point of modulus 12. Use slightly shorter final steps, still strictly between ρ_i and $2\rho_i$, if necessary. All these annuli lie inside B_{16}. For two successive annuli their inner disks intersect, and each inner disk contains a point outside the other outer disk. The interiors of any two separating loops therefore intersect, but neither can contain the other; their boundaries must intersect. Thus each loop chain is connected. It has points in B_8 and points outside B_{10}, so it meets J. This explicitly joins every private access path to J. The total network length is a finite sum of annular loop lengths and therefore has all positive moments. Only the minimum private access cost uses

independence; the network itself need not be independent of those costs. Consequently, for any prescribed $p < \infty$, a sufficiently large fixed cloud satisfies

$$\mathbb{E}\left[e^{-ah_1(0)}\left(\min_i D(u_i, J; B_{16}) + \mathrm{len}_D(J)\right)\right]^p < \infty.$$

One does not claim all moments for a single cloud independent of p.

Scale this construction at every point z of an r-grid, and put $N_r(z) = r^q e^{ah_r(z)}$. For every $\eta > 0$, inverse crossing moments give, simultaneously on this grid,

$$D(\text{across } A_{2r,3r}(z)) \geq r^\eta N_r(z)$$

outside an arbitrarily high polynomial failure. Choose the cloud moment order large enough that its minimum access plus loop length is at most $r^{-\eta}N_r(z)$ at every grid point with the same failure bound.

If a protected geodesic of length at most M visits $B_r(z)$, Lemma B.3 gives a larger surrounding loop of length at most CMr^{θ_0}. The geodesic reaches that loop from its visit and crosses $A_{2r,3r}(z)$, so the crossing cost is at most CMr^{θ_0}. It also meets the cloud loop on both sides of the visit. Loop comparison and the cloud access bound give

$$D(x, H_r) \leq Cr^{-\eta}N_r(z) \leq CMr^{\theta_0 - 2\eta}.$$

Take $2\eta < \theta_0$. The uniform loop event has polynomial failure, and the remaining failures can have any requested polynomial order. Borel–Cantelli proves the assertion.

For a filled-frontier point, use [3, Corollary 3.7] and the frontier access observation following Lemma B.3. A root-to-frontier geodesic crosses the inner annulus before reaching the frontier; its last part from the surrounding loop has length at most the loop length. This supplies the same crossing bound and cloud access. If the root-target clearances are powers of a larger scale, choose r to be a sufficiently small power of that scale. Replace the larger random scale by a comparable dyadic scale before using the deterministic cloud. □

Lemma B.11 (Finite multiplicity on extendable points)**.** *Let $\mathcal{E}$ be the set of strict interior points of finite global geodesics. There is a deterministic $N = N(a) < \infty$ such that, almost surely, every pair in $\mathcal{E} \times \mathcal{E}$ is joined by at most N geodesics.*

Proof. Lemma B.6 gives unique strict prefixes from every root in the fixed countable deterministic set.

We spell out the three-geodesic reduction. For a fixed finite-target prefix configuration at levels $t < s$, choose three distinct prefix endpoints whose classes are I_1, I_2, I_3. At most one exterior target can be reached by root geodesics through all three prefix endpoints. If w is such a target, the three tails after level s are disjoint until w, by strict-prefix uniqueness, and their contacts lie in the interiors of the three classes. In the Jordan target component these three tails partition the domain into three regions. A second target lies in one of them, whose boundary uses only two tails and the intervening frontier arc. A route from the remaining class to this target must meet one of those tails or pass through w. The former would contradict strict-prefix uniqueness at the intersection; the latter would make the nonunique prefix to w a strict prefix. Thus both alternatives are impossible.

If four root-to-target geodesics exist, take rational levels after all four tails have separated, and choose a rational filling target in the appropriate ordinary-ball component. Select a rational inner ball visited by the fourth tail and a larger rational ball whose closure misses the other three tails

and the explored filled set. Fix a smooth nonnegative bump ϕ supported in the larger ball and equal to one on an open neighborhood of the inner ball. Prefix labels, both balls, and the bump can be chosen from countable catalogues.

For this fixed data, let E_α be the fourth-route event for $h + \alpha\phi$. Consider two coefficients $\alpha < \beta$ for which the root-prefix and prefix conclusions hold. On E_α, increasing the coefficient preserves the ordinary ball through level s, hence its finite-target fill and the prefix data: distances cannot decrease, and every original root geodesic of length at most s avoids the support. The three avoiding geodesics remain minimizing to the same target and preserve its root distance. If E_β also holds, the preceding uniqueness for a three-prefix target forces its target to be the same point. A fourth route in the β-metric which enters the inner ball has strictly smaller length in the α-metric, since it traverses a positive-length portion of the plateau. This contradicts its preserved root distance. Therefore E_α and E_β cannot both hold.

For almost every base field, the root and prefix conclusions hold for Lebesgue-almost every coefficient, by Cameron–Martin equivalence and Fubini. On that full-measure coefficient set, at most one coefficient can satisfy the fourth-route event. Sampling a coefficient with a density gives probability zero for that event. Equivalence under each fixed Cameron–Martin shift then gives probability zero in the original field. Countable prefix endpoint and ball choices exhaust every four-geodesic configuration. This proves at most three root geodesics to every nonsingular target, simultaneously for the fixed countable family of roots.

Here is the quantitative ingredient. Fix n disjoint private bump carriers, with smaller plateau sets. Let $E_\epsilon(z, w)$ be the event that $n + 1$ geodesics connect possibly different points in the two D-balls of radius ϵ about deterministic z, w; path i meets its private plateau and avoids every other carrier, and path zero avoids all carriers. Require the endpoint balls to avoid each other and every carrier. On a cutoff that all transverse plateau costs are at least C^{-1}, the private-bump comparison of [7, Lemma 4.6] gives

$$\mathbb{P}(E_\epsilon(z, w),\ \text{cutoff}) \le C_{t,C}\epsilon^{n/t}, \qquad t > 1. \tag{B.9}$$

Its deterministic proof only uses the following facts. The endpoint balls are unchanged by all the bumps, since the first visit to a carrier already costs more than ϵ. Their points can be joined through their centers at cost at most 2ϵ. If two coefficient vectors differ by more than $C'\epsilon$ in one coordinate, the corresponding private arm changes cost by more than the possible endpoint adjustments, while the zero arm is unchanged; one of the two fields then has a strictly shorter competitor. The set of coefficient vectors producing the event has volume $O(\epsilon^n)$. The finite-dimensional translation density has every moment, yielding (B.9). Positive transverse costs follow from lower semicontinuity on disjoint compact sets. No upper ambient modulus is used.

Cover a compact geodesic frame by Lemma B.10. With $\epsilon_r = CMr^\theta$, its endpoint pairs are assigned to at most $O(r^{-4})$ deterministic cloud pairs. If an $n + 1$-arm configuration exists, the corresponding endpoint balls miss the fixed private carriers at all sufficiently small scales. Formula (B.9) and a union bound give a bound $Cr^{\theta n/t-4}$. Choose n with $\theta n > 4$ and $t > 1$ close enough to one. Along dyadic scales this and the cloud failures are summable. Countable compact frames, carrier configurations, and cost cutoffs show that no pair in $\mathcal{E}^2$ has $n + 1$ geodesics with a private point on each path. Every point of $\mathcal{E}$ belongs to one such frame by trimming its witnessing geodesic.

The remaining reduction is deterministic. The proof of [7, Lemma 4.11] uses only strict rational-root prefixes and the three-geodesic bound: if two prefixes of a path meet again, choose a rational point in a bubble component opposite its continuation and splice a geodesic from this rational point through the bubble boundary. It follows that every later strict segment is unique and there

are at most three suffixes. Thus two same-endpoint geodesics make at most two excursions away from one another. The finite-family argument of Lemma 4.13 of the same paper then gives, on each path, an interior point hit by at most three family members, including that path. All these arguments concern Euclidean continuous simple curves; their endpoint pairs are finite-distance pairs.

For a family of M paths, join two vertices of a conflict graph if either path hits the other's selected point. It has at most $2M$ edges. A random-ordering argument and Cauchy–Schwarz give an independent set of size at least

$$\sum_v \frac{1}{\deg(v)+1} \geq \frac{M^2}{2|E|+M} \geq M/5.$$

An independent family has a private point on every path, so its size is at most n. Hence $M \leq 5n$. Apply this to every finite subfamily. A larger normalization circle disjoint from the finitely many bump supports handles their additive normalization; the conclusions are invariant under an overall metric scalar. □

B.5 Nearby disjoint geodesics cannot persist

For two simple curves, their *one-sided distance* is the minimum over a fixed choice of a side of each curve of the following symmetric access bound: every point of either curve can be reached from the selected side of the other by a path of length at most that bound which otherwise avoids the other curve. This is the prime-end definition in [10, Section 5.1]. It is at least their ordinary D-Hausdorff distance.

Lemma B.12 (Global protected ladders)**.** *Fix a nonconstant geodesic η_* of length T_*. Suppose disjoint geodesics η approach it in one-sided distance $\delta \to 0$. Form the endcap rectangle from the two curves and short endpoint connectors. Every internally shortest ladder from the reference side to a protected central point of η is, for sufficiently small δ, a global geodesic and globally minimizes distance to the reference side.*

If L_* is the supremum of these protected ladder lengths, the central rectangle contains a ball of radius cL_* with center on a globally geodesic ladder or on a finite-target filled frontier. The center has polynomial Euclidean clearance from the endpoints or root and target used in its short-loop estimate. A polynomial deterministic cloud therefore contains a point inside this ball with a fixed metric margin.

Proof. The endpoint orientation estimate [10, Lemma 38] is a metric triangle-inequality argument: after orientation the lengths differ by at most 2δ, the corresponding endpoints are at distance at most 5δ, and corresponding arclength points are within 7δ. Endpoint connectors can be clipped so their long-side contacts are intervals. Their contact lengths are at most 5δ, and the connectors cannot meet each other for small δ.

Work in the closed endcap rectangle, allowing its planar interior to be unbounded. Bounded minimizing sequences of ladders have a common Euclidean compact range and endpoints on the metric-compact reference; Lemma B.1 gives a shortest ladder. Rightmost ladders are obtained by right-envelope splicing and compact limits. The deterministic separation arguments of [10, Lemmas 39–41] give length at most δ, ordered ladders, and the protected central rectangle. They use only the geodesicity of the four sides and this minimizing property. In particular they require no upper Euclidean-to-metric modulus.

Let ξ be an internally shortest ladder, from $a \in \eta_*$ to $b \in \eta$, of length L. It has no reference-side contact other than a, since otherwise its initial part could be deleted. At its first contact c with η, replace the remaining tail by the η-interval from c to b. The replacement has no greater length; internal minimality forces equality. We keep the name ξ for the replacement and write $v = \xi(L/2)$. Its initial part $\xi[a, c]$ is a proper crosscut of the endcap rectangle: its interior avoids both long sides and both caps. The caps are avoided because $L \le \delta$, the reference endpoint a belongs to a fixed strict middle of η_*, and the caps shrink to the two original reference endpoints.

The same observations, together with the lower metric-to-Euclidean modulus on a common compact set, imply that

$$\mathrm{Hull}_{\mathrm{Eucl}}(\overline{B}^D_{2L}(v)) \quad \text{misses both original endcaps} \tag{B.10}$$

for all sufficiently small δ. To see this directly, each such metric ball lies in a Euclidean disk of radius tending uniformly to zero about a. Every bounded component of the complement of a subset of that disk is also in the disk. The disks have positive clearance from the two original endcaps. This reasoning allows the chosen rectangle itself to be unbounded in the plane.

The reference is D-compact, so choose a closest reference point to b and a global geodesic G from b to that point. Write $d = \mathrm{len}_D\, G \le L$. Let x be the first reference-side contact of G, and let y be its last η-contact before x. The open bridge $G(y, x)$ avoids the long sides. Moreover $G \subset \overline{B}^D_{3L/2}(v)$, so it avoids the caps by (B.10). It therefore lies entirely inside the rectangle or entirely outside it. In the first case, the concatenation of the η-interval from b to y with $G[y, x]$ is an internal competitor of length at most d. Internal minimality gives $L \le d$, hence $d = L$.

Suppose instead that this bridge is external. The four arcs

$$\xi[a, c], \qquad \eta[c, y], \qquad G[y, x], \qquad \eta_*[x, a] \tag{B.11}$$

form a Jordan curve, with zero-length boundary intervals omitted. Indeed, the two bridge interiors lie in opposite complementary domains of the rectangle's Jordan boundary, and neither meets a long side; the long-side intervals meet the bridges only at their listed endpoints. This Jordan curve separates the two caps. For a precise separation argument, the proper internal crosscut splits the rectangle into two components, one adjacent to each cap. Join a point of one cap to a point of the other by a curve in the rectangle interior which crosses this crosscut transversely once and avoids its endpoints. This curve is disjoint from the external bridge and from the long-side intervals, so it has odd intersection number with (B.11). Its two endpoints are on opposite sides of that Jordan curve. Each cap is connected and disjoint from the curve, so one complete cap is on its bounded side.

We verify that (B.11) lies in $\overline{B}^D_{2L}(v)$. The ladder part lies in $\overline{B}^D_{L/2}(v)$, and the comparison bridge lies in $\overline{B}^D_{L/2+d}(v)$. Put $q = \mathrm{len}_D\, \xi[a, c]$ and $d_y = \mathrm{len}_D\, G[b, y]$. If $q \le L/2$, then v is on the η-interval from c to b, and the distance to v on the η-interval $[c, y]$ is bounded by the larger of its endpoint values, hence by $L/2 + d_y$. If $q > L/2$, put $B = D(c, y) \le L - q + d_y$. A point of this interval at arclength t from c satisfies

$$D(v, z) \le \min\{q - L/2 + t,\ L/2 + d_y + B - t\} \le \tfrac{1}{2}(q + d_y + B) \le L/2 + d_y.$$

Finally, if $A = D(a, x) \le L + d$, a point of the reference interval at arclength t from a satisfies

$$D(v, z) \le \min\{L/2 + t,\ L/2 + d + A - t\} \le L + d \le 2L.$$

The bounded side of (B.11) is therefore contained in the Euclidean filled hull in (B.10), contradicting the enclosed cap. This rules out the external alternative. We have proved $D(b, \eta_*) = L$. In

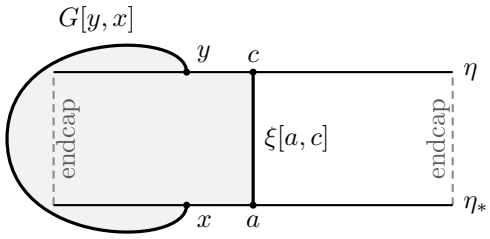


Figure 1: The internal ladder and an exterior crossing, joined by intervals of the two long sides, separate the endcaps. In the proposed shortcut, all four pieces would lie in a small metric ball. Its Euclidean filled hull misses both endcaps, giving the contradiction in Lemma B.12.

particular $D(a,b) = L$, so the original ladder, as well as its replaced version, is a global geodesic. Its midpoint satisfies $D(v,\eta_*) \geq L/2$.

The ball $\overline{B}^D_{L/8}(v)$ misses the reference side and the two central end-ladders. The latter assertion follows from the protected parameter margins: an intersection would join b to the corresponding end-ladder endpoint on η at cost at most $L+\delta$, smaller than the prescribed parameter separation. It cannot meet both η-portions on opposite sides of b. Indeed, two such contacts would have mutual distance at most $L/4$, while each has distance at least $D(v,b) - L/8 = 3L/8$ from b; geodesicity of η contradicts these inequalities.

If the ball misses both portions, its geodesic radial paths cannot cross any rectangle boundary, and it is already an interior ball. Otherwise choose the portion it misses, and call its corresponding original endcap the avoided cap. Let C be the component on that side of the proper initial crosscut $\xi[a,c]$. More explicitly, C is the Jordan region bounded by this crosscut and the rectangle boundary arc through the avoided cap. Any part of ∂C which is not on ξ misses $\overline{B}^D_{L/8}(v)$. This remains true if $\xi[c,b]$ follows η: if c is on the avoided portion, then $D(v,c) \geq L/8$, while if c is on the other portion, the additional boundary interval from c to b is itself part of ξ.

Put $r = L/64$ and $v_i = \xi(L/2 + (-1)^{1-i}r)$ for $i = 0, 1$. Fill the radius-r ball about v toward the original endpoint of η at the avoided cap. The two tails of the global ladder, followed by the reference side, the avoided cap, and the avoided η-portion as appropriate, connect the ladder points outside $[v_0, v_1]$ to this target without entering the closed ordinary ball. They thus belong to the target component. The strict ladder interval between v_0 and v_1 is in the ordinary open ball. Consequently the filled frontier intersects ξ exactly at v_0, v_1.

This frontier is a D-compact Jordan curve, all of whose points have D-distance r from v, by [3, Theorem 1.4 and Lemma 3.4]. One of its two v_0-to-v_1 arcs lies in C. Indeed, it cannot meet $\partial C \setminus \xi$, which misses the larger metric ball; its only possible contacts with ξ are v_0, v_1. The strict intervening ladder interval is in the interior of the filled set, so the two frontier arcs lie on opposite local sides of that interval. This is the Jordan separation theorem applied to ∂C and the filled frontier.

On the arc in C, continuity in D gives a point v_2 with

$$D(v_2, v) = r, \qquad D(v_2, v_0) = D(v_2, v_1) \geq r.$$

If $|t - L/2| \geq 3r/2$, global geodesicity of ξ gives $D(v_2, \xi(t)) \geq r/2$. Otherwise $\xi(t)$ is within $r/2$ along ξ of one of v_0, v, v_1, giving the same bound. Hence $D(v_2, \xi) \geq r/2$. Every geodesic from v_2 to a point of $B^D_{r/4}(v_2)$ stays in $B^D_{5r/4}(v) \subset B^D_{L/8}(v)$. It cannot hit ξ or any other boundary piece of C, and therefore stays in C. This proves $B^D_{L/256}(v_2) \subset C$. The chosen frontier arc and these radial

geodesics also remain between the central end-ladders: they start there and lie in the larger ball which misses both end-ladders. Thus the smaller ball is in the central rectangle.

Here is the polynomial clearance needed for the cloud application. A short loop about either ladder endpoint, obtained using its protected macrogeodesic, has length at most $CM\rho^\theta$ and surrounds a Euclidean ball of radius comparable to ρ. Choose ρ to be a sufficiently small fixed power of L so this length is less than $L/10$. If both ladder endpoints were inside the loop, their protected macrogeodesic access paths and a loop arc would give $D(a,b) < L$. If only one endpoint is inside, every ladder contact with the loop is at distance at most twice its length from that endpoint. Thus the midpoint cannot be inside the loop. It follows that v has a positive power of L as Euclidean clearance from a, b. Now apply the short-loop estimate to the global ladder at v, on a sufficiently small further power of L. The resulting loop has total perimeter and access cost less than $r/2$, so its trace lies in the ordinary radius-r ball. The distant filling target is outside it; hence its interior is swallowed by the filled ball. The frontier point v_2 is outside this interior, giving a positive-power Euclidean clearance from its root v. Its target has macroscopic clearance. These are precisely the hypotheses for the two cloud alternatives.

Choose $L \geq L_*/2$ and put $R = L_*/2048$. The resulting ball has radius at least $4R$ before a harmless change of constants, and a cloud point within $R/4$ of its center retains a ball of radius at least $2R$ inside the rectangle. Every protected ladder has length at most $2048R$. Replace R by a comparable dyadic scale before choosing the cloud. Its cardinality is polynomial in R^{-1}. □

Lemma B.13 (One-sided intersection)**.** *Almost surely, for every nonconstant global geodesic η_* there is $\delta_* > 0$ such that no geodesic disjoint from η_* has one-sided distance less than δ_* from it.*

Proof. We give the replacements for the two upper-modulus uses in the proof of [10, Proposition 37]. Use the notation of Lemma B.12, and choose the cloud point p in its interior ball. It has $D(p,a) \leq C_0 R$ for a point a on the strict reference middle. The purely planar splicing argument of [10, Lemma 42] supplies, after exchanging the two long sides if necessary, a global geodesic σ from p to an original endpoint of η_* contained in the endcap rectangle and the reference. Its length is uniformly bounded by $T_* + C_0 R$.

The selected roots satisfy common eligible-root cutoffs. Indeed, short accessible loops about their reference anchors are reachable at cost $C_0 R + 2\,\mathrm{len}(J)$. Choosing a fixed sufficiently small loop scale puts a Euclidean neighborhood of each root inside K_{τ_ℓ}. A bounded reference continuation gives the upper exit-time bound. Positive macroscopic annular crossing costs give the separation bounds, and the published simultaneous estimates give the remaining cutoffs. These assertions use a countable localization. Choose ℓ small enough that the early X interval lies between fixed times $0 < \phi < \psi < T_*/200$, with the lower time bound supplied by a fixed annular crossing and the upper time bound by an accessible reference loop. The target is Euclidean-separated from the reference middle. All cloud-root and center unions are performed before selecting p.

Fix a lower-modulus exponent $b > 0$, so $|x - y| \leq CD(x,y)^{1/b}$ on the compact localization. Ordered ladders subdivide the protected rectangle into strips. If their two endpoint parameters differ by at most d and the ladder lengths are $O(d)$, the reference-side length is also $O(d)$ by the triangle inequality along the two global sides. The four boundary pieces lie in a Euclidean disk of radius $Cd^{1/b}$; their bounded region lies in its convex hull. In particular the protected rectangle shrinks toward a strict reference middle and stays away from the target.

We need a bound for a portion $\sigma[\alpha,\beta]$ trapped in a disk of radius $r \geq cR^{1/b}$ meeting this strip. If the starting point p is outside the larger short-loop neighborhood, ordinary two-sided loop comparison gives $\beta - \alpha \leq CMr^\theta$. If it is inside, then its reference anchor a is within Cr^χ of this disk. A second accessible loop about a, with inner radius $C'r^\chi$ and outer radius $C''r^{\chi^2}$, surrounds

both the disk and p. Every point on this loop is reachable from p at cost at most $C_0R + Cr^{\chi\theta}$. The target is outside it, so a contact after the trapped portion bounds that entire portion. Thus

$$\beta - \alpha \le C(R + r^{\vartheta}), \qquad \vartheta > 0. \tag{B.12}$$

This handles a possible initial portion near the root; a two-sided loop estimate alone would not.

If a segment meets no two ordered ladders whose parameters differ by $4d$, it lies in a limiting substrip with boundary lengths $O(d)$. Formula (B.12) gives duration at most $C'd^{\rho}$ for some $\rho > 0$, first with $d = \delta$ and then with $d = 2048R$. Before the first use, the middle-ladder separation in the deterministic setup, namely distance greater than $T_*/10$ from the opposite middle ladder, keeps the initial segment of length $C\delta^{\rho}$ inside the protected rectangle. Rightmost-ladder ordering then keeps the later path on the inward side of a ladder reached in that initial segment. This is the containment argument in Step 2 of the cited proof; it precedes, rather than follows from, the passage estimate. It yields

$$\sigma[\phi/2, 2\psi] \subset \text{a more strictly protected central rectangle.}$$

The same anchored-loop comparison shows that this interval stays a fixed positive Euclidean distance from p: otherwise a small reference-based loop surrounding p and one of its points would give time at most $C_0R + C\,\mathrm{len}(J) < \phi/2$. Lemma B.3 now gives a uniform positive-power modulus

$$D(x, y) \le C|x - y|^{\theta_f}$$

for nearby points on this interval and on the protected reference frame.

Choose $\beta_2 > 0$ so small that $\beta_1 = (1+\nu)\beta_2 < 1/b$. Apply Proposition B.9 with small parameter comparable to R^{β_2}. The polynomial cloud union is harmless to its superpolynomial estimate. It supplies an X stem in $\sigma[\phi, \psi]$ of carrier radius $r_X \in [R^{\beta_1}, R^{\beta_2}]$, up to constants. Its barriers lie within Cr_X of the stem and the whole query avoids their endpoint balls of radius cr_X.

Choose

$$0 < \kappa < \min\{\rho, \beta_2\theta_f, 1\}, \qquad h_R = R^{\kappa}.$$

The passage estimate supplies a ladder contact in each of the two padding intervals at distance between h_R and $2h_R$ before and after the stem. Their ladders have length at most $2048R$. If a barrier met one of them at v, its reference endpoint a_v would satisfy $D(v, a_v) \le 2048R$. Euclidean proximity to the stem and the frame modulus would give

$$D(v, \text{stem}) \le 2048R + CR^{\beta_2\theta_f}.$$

But the selected query contact on the same ladder is at time distance at least h_R from the stem, giving $D(v, \text{stem}) \ge h_R - 2048R$, a contradiction.

The strip between these two ladders lies in the $CR^{1/b}$-neighborhood of the padded query segment. To see this, subdivide it by ladders at parameter increments at most R; each small strip has four sides of length $O(R)$, and the query crosses its intervening ladders. The lower modulus and convex-hull containment give the assertion. Since $R^{\beta_1} \gg R^{1/b}$, whole-query endpoint avoidance puts every barrier endpoint outside the strip. Each barrier must therefore exit through its appropriate long side, as it cannot meet the two chosen ladders. The X sidedness and uniqueness of the barriers force both long sides to contain the common stem, by the final planar separation in Step 5 of the cited proof. This contradicts their disjointness.

All choices use strict exponent inequalities in the order b, χ, θ; ρ; ϕ, ψ; θ_f; β_2, β_1; κ. The stochastic estimates are simultaneous on countable compact localizations and dyadic scales; hence the conclusion holds for every reference geodesic on one probability-one event. □

B.6 No bubbles and strong confluence

Theorem B.14 (Strong confluence at every positive parameter)**.** *Almost surely every strict interior subsegment of every finite D-geodesic is unique between its endpoints. If geodesics P_n converge to a nonconstant geodesic P in D-Hausdorff distance, then every fixed strict middle segment of P is contained in P_n for all sufficiently large n, after consistent orientation. The same conclusions hold for ordinary internal geodesics whose images lie in a fixed compact subset of their open comparison domain.*

Proof. First, convergence in D-Hausdorff distance implies convergence of lengths and corresponding endpoints after orientation, and uniform D-convergence with the length parameter. This is the triangle-inequality argument of [10, Lemma 38], already used above.

Suppose a sequence has no intersection with a fixed strict interval of P. Trim to a smaller interval and take a topological rectangle around the reference arc. Uniform Euclidean convergence gives proper subcrosscuts of the approximating curves on one side of the reference. They bound a strip. Cover a further strict middle by finitely many surrounding loops of arbitrarily small length. Both crosscuts meet each loop, and consecutive contacts with different crosscuts give a loop arc inside the strip joining them. The geodesic passage bound controls the arclength displacement of these contacts. Following such a loop arc and a short part of the other geodesic proves one-sided convergence on this smaller interval; endpoint time strips can be made arbitrarily short. Lemma B.13 is contradictory. Thus every fixed nonconstant strict interval is eventually intersected. This proves the needed nearby-intersection statement without an ambient upper modulus.

Now suppose a reference geodesic has an alternative segment between two strict interior points. Trim the full reference endpoints so that they belong to $\mathcal{E}$ and the alternative remains strictly inside; write the trimmed path as $P : [0, T]$ and the alternative as $R : P(u) \to P(v)$ with $0 < u < v < T$. For every s in a countable infinite subset of $(0, u)$, choose rational q_n with $D(q_n, P(s)) \to 0$. Geodesics $G_n : q_n \to P(T)$ have lengths tending to $T - s$. Compactness yields a geodesic limit $Q^s : P(s) \to P(T)$, and Lemma B.2 upgrades it to uniform D-convergence.

This limit cannot begin with a nonconstant segment of P before u. Otherwise the nearby-intersection statement forces G_n to meet that reference interval for large n. At this contact, replace its remaining tail first by the reference tail and then by the tail containing R. Both are geodesics from the rational point q_n to $P(T)$, but their strict prefixes to $P(v)$ differ, contradicting Lemma B.6. Prepending $P[0, s]$ to Q^s now gives distinct geodesics from $P(0)$ to $P(T)$ for distinct s: equality for $s < t$ would force Q^s to begin with $P[s, t]$. This contradicts Lemma B.11, since the trimmed endpoints belong to $\mathcal{E}$. Strict interior uniqueness follows.

For the convergent sequence P_n, nearby intersection supplies contacts in small strict intervals near both ends of P. Interior uniqueness identifies the segment between them. Shrinking these intervals gives the asserted containment of every prescribed strict middle.

Finally fix $K \Subset V \Subset W$ with $\overline{W}$ contained in the original open comparison domain, and internal geodesics with images in K and lengths bounded by M. Choose a deterministic smooth φ equal to one on K and supported in V. A competitor leaving W pays a positive unchanged crossing cost c between ∂V and ∂W. For a large enough integer H, $e^{-aH} M < c$. Each original internal path is a global geodesic for $h - H\varphi$: inside W every competitor pays at least the same Weyl floor times its old cost, while an exiting competitor pays at least c. The countable family of these fixed shifts is Cameron–Martin equivalent to the original law. Both conclusions therefore transfer to the internal metric, including convergent sequences. A countable compact exhaustion removes the auxiliary bounds. □

B.7 Local crossing counts and their moments

Proposition B.15 (All moments of the ordinary crossing and trace counts)**.** *Let*

$$K_{z,r} = \overline{A_{r/2,4r}(z)}, \qquad O_{z,r} = A_{r/3,5r}(z).$$

Let $N_{z,r}$ be the supremum of the sizes of finite families of pairwise closed-disjoint proper crossings of $A_{r,2r}(z)$, each minimizing between its endpoints among curves in $K_{z,r}$. For the parameter-a pure metric,

$$\mathbb{E}N_{z,r}^p < \infty \qquad (p < \infty),$$

with bounds independent of z, r. The count is measurable from $h|_{O_{z,r}}$ modulo constants. For geodesics in one common open ordinary comparison domain, with full endpoints clearing a larger observation disk, the number of complete traces in a smaller closed disk has moments of every order as well.

Proof. First we prove almost-sure finiteness at one shape. Choose two rectifiable circuits in disjoint strict subannuli of $A_{r,2r}(z)$, one near each boundary, and a finite rectifiable connection between them in $\operatorname{int} K_{z,r}$. Such circuits and connections follow from the annular estimates and a finite circuit network; see [4, Lemmas 2.6 and 2.8]. Trim a crossing at its last contact with the inner circuit before its first contact with the outer one. Its image lies in the compact annular region between these circuits. Its endpoints belong to fixed rectifiable curves, and its length is bounded by the sum of their lengths and the connecting length. This is a geodesic for the internal metric of $\operatorname{int} K_{z,r}$: every shorter competitor in that open set would also be a competitor in $K_{z,r}$.

The endpoint sets are compact for this internal metric, since their finite-length parametrizations are continuous for it. A sequence therefore has a subsequence with metrically convergent endpoints. The triangle inequality gives convergence of the endpoint distances. Lemma B.1 gives Euclidean convergence to an internal geodesic. The limit is nonconstant, since the two circuits have positive Euclidean separation. Lemma B.2, applied in the same fixed ordinary comparison domain, upgrades this convergence to uniform convergence in the internal metric. Thus the trimmed path class is compact in the metric Hausdorff topology.

For each path in this compact class, Theorem B.14 supplies a neighborhood whose members contain one fixed nonconstant strict subsegment of that path. A finite subcover bounds every finite disjoint subfamily: at most one member can be assigned to each neighborhood. This proves a uniform finite packing bound on each realization, and hence $\mathbb{P}(N_{0,1} > k) \to 0$. Merely excluding a single infinite disjoint family would not give this conclusion.

We amplify this qualitative tail. Use separated concentric observation annuli, with a fixed ratio large enough that their strict buffers have bounded overlap. Lemma 4.3 and Proposition 4.4, at the pure base law, give a constant $K < \infty$ such that the probability that m specified annuli all have count greater than k is at most

$$\mathbb{P}(N_{0,1} > k)^{m/K}.$$

Additive constants in the field have no effect on these counts. Exact scaling identifies their one-annulus laws. Since there are $c\log(1/\delta) + O(1)$ separated scales between δ and a fixed small radius, choosing $k = k(H)$ large enough makes persistence through all these scales have probability at most $C_H\delta^H$, for any prescribed H.

Suppose a macroscopic family contains more than $Ck\delta^{-2}$ disjoint minimizing crossings. A deterministic δ-grid covering a middle circle has a ball visited by at least k of them. Their endpoints lie outside a fixed larger disk around its center. For each intermediate radius, trim at

the last inner-circle contact before the first outer-circle contact. This produces k closed-disjoint proper crossings. Require each local $K_{z,r}$ to lie in the macroscopic comparison domain. A shorter competitor in this closed local annulus would shorten the original crossing, so these local crossings really are counted by $N_{z,r}$. A union over the $O(\delta^{-2})$ possible centers gives

$$\mathbb{P}\{N_{\mathrm{macro}} > Ck(H)\delta^{-2}\} \le C_H \delta^{H-2}. \tag{B.13}$$

Taking $H > 2p + 2$ proves the required pth moment. Translation and scaling give the same constants at every physical radius.

For complete traces use Lemma 6.1, whose deterministic proof needs only simple Euclidean-continuous finite geodesics and strict-interior uniqueness. Its incoming/outgoing-arm argument gives at most $1 + 2N^2$ traces. Their arm counts are bounded by the local closed-annulus counts just proved. A finite cover handles any fixed buffered observation shape. All representatives are compared in the same open ordinary domain; the local closed-annulus counts bound their arms, and do not replace that common comparison domain. □

Measurable versions. Use the jointly measurable lower-semicontinuous extended kernel of the canonical metric. It can be coded by its closed epigraph in the Euclidean product space; continuity of the ambient kernel is not assumed. For a fixed Euclidean compact set K, the intrinsic kernel d_K is lower semicontinuous: bounded almost-minimizing paths with convergent endpoints have a Euclidean-uniform subsequential limit in K, whose length is bounded by partition lower semicontinuity. The same argument, together with measurable projection, supplies its jointly measurable completed-law version. For an open domain O, use a countable compact exhaustion: every finite curve with compact image in O belongs to some member, so d_O is the infimum of these compact intrinsic kernels on the relevant endpoint domains.

The space of Euclidean-continuous curves on a compact time interval is Polish. The relation

$$P([0,1]) \subset K, \qquad D(P(s), P(t)) \le L|s - t| \quad (s, t \in \mathbb{Q} \cap [0,1])$$

has compact sections for fixed D, K, L, by the positive off-diagonal modulus and Arzelà–Ascoli. Lower semicontinuity extends the inequalities to all times. Length is the supremum of the rational partition sums. Consequently finite length, minimizing equality with d_K or d_O, crossing conditions, closed disjointness of a finite tuple, and equality of complete closed-core traces are measurable relations. Their path projections are analytic, and are therefore measurable after completion. Universal assertions are treated by the complement of the corresponding bad-path projection. This establishes the measurability used in the crossing argument and permits completed-measurable selection of representatives. It does not assert that an arbitrary analytic projection is Borel, or that the varying-endpoint geodesic relation itself has compact sections.

The occupation of an arclength-parametrized path is the pushforward of Lebesgue measure. Equal trace images have equal restricted occupation measures, since a simple geodesic carries metric one-dimensional Hausdorff measure. Enumerate distinct traces by selecting a representative and then excluding its trace class. The finite trace bound terminates this procedure. Rational compact localizations, length cutoffs, and a fixed enumeration of the observation charts give versions measurable from the unchanged ordinary region alone.

B.8 Quantitative ports and exact closed rows

Theorem B.16 (A universal finite port representation)**.** *Let W_0 be a fixed compact ordinary annular buffer, let $O = \operatorname{int} W_0$, and let $A \Subset O$ be a middle annular band. The inner filling*

and all modified supports lie strictly inside its inner boundary; exterior attachments and query endpoints clear its outer boundary. Suppose the relevant ordinary query passages have a common cap furnished by a surrounding ordinary rectifiable loop.

There is an exterior-measurable Jordan separator $\Sigma \subset A$ and a finite set of transverse ordinary ports on Σ with the following properties. Their number, and the number of ordered port rows, have moments of every order under the pure law, uniformly under exact scaling. Ordinary rows in the bounded closed interior I of Σ are finite and attain their infima. Rows with a compact inner set collapsed to a free vertex have attained ordinary arms. For every compatible positive inner filling, the full query infimum is exactly the finite port-graph infimum using these local closed rows and the genuine exterior-restricted kernels. No attainment in the inner filling or in an exterior open domain is required.

Proof. Cover A by finitely many closed observation cores with strict ordinary buffers in O. In each larger buffer retain a rectifiable circuit surrounding the core, with its closed interior compactly contained in O. Trim a relevant closed-W_0 minimizing path at the component of its visit to this circuit's interior that contains the core visit. The resulting whole finite segment is an O-internal geodesic, since every shorter O competitor is also allowed in W_0. Its length is at most the circuit length, and its endpoint set is compact in d_O. The compactness argument in Proposition B.15, including Lemma B.2, applies. Interior strong confluence makes its complete core trace locally constant. Endpoint clearance ensures that a short terminal tail cannot create an additional core visit, even for a tangent visit to the closed core. Select one whole representative for each distinct complete core trace in each member of the finite observation cover. Their union contains every relevant trace on A.

Let J be the total number of these representatives. It is bounded by the sum of the complete-trace counts in the fixed finite observation cover, so Proposition B.15 gives all moments of J. Write $s_R = R^q e^{a h_R(z_0)}$ for the local unit. The representative lengths are bounded by a random variable U with $\mathbb{E}(U/s_R)^p < \infty$ for every p, using the finite set of retained annular circuits. Strict-interior uniqueness makes the intersection of two representative interiors connected. The finitely many full endpoints can add only finitely many further contacts. Their whole union is thus a finite embedded graph with at most CJ^2 edges and total distinct edge length at most JU. We retain whole representatives before forming this graph; clipping at a circle could create infinitely many pieces.

The edge count alone is insufficient for a quantitative port bound, since one edge could oscillate across a band many times. Fix a deterministic $\rho > 0$ small compared with the width of A and with the clearance of all representative buffers. Take a finite $\rho/16$-net of their compact union region, and let c_ρ be the minimum of the intrinsic crossing costs of $A_{\rho/8,\rho/4}(z)$ at these net centers. Each such annulus lies in the observed ordinary region. A curve whose Euclidean diameter is at least ρ crosses one of them and therefore costs at least c_ρ. The fixed-shape negative annular moments give

$$\mathbb{E}(s_R/c_\rho)^p < \infty \qquad (p < \infty).$$

Subdivide each graph edge at arclength increments at most $c_\rho/2$. Every resulting edge has Euclidean diameter less than ρ, and their number E' satisfies

$$E' \le CJ^2 + \frac{2JU}{c_\rho}. \tag{B.14}$$

Hölder's inequality gives all moments of E'; no independence is used.

Choose a_0 so that the radial band from a_0 to $a_0 + 2\rho$ lies in A and surrounds the hidden region. Select every whole small edge meeting $\overline{B}_{a_0}$, and let K_0 be their union with that disk. Then K_0

is connected and lies in $B_{a_0+\rho}$. An unselected edge meets K_0 only at some of its at most two endpoints. Around these finitely many vertices, choose disjoint vertex-star charts contained in $B_{a_0+2\rho}$ and trim short outgoing arm pieces. The remaining unselected edge cores are compact and disjoint from K_0.

A finite planar graph is tame: straighten each vertex star in a topological disk, then each remaining edge in a disjoint topological rectangle; the planar Schoenflies theorem extends these charts. In these charts take a connected polygonal neighborhood of K_0, avoiding the remaining cores and crossing each outgoing arm once. Away from the finitely many crossing windows, positive compact separation permits a sufficiently fine polygonal neighborhood. Fill its bounded complementary components. Its outer boundary is a Jordan separator Σ in $B_{a_0+2\rho} \setminus \overline{B}_{a_0}$; filling components can delete contacts but creates none. Selected edges lie in its interior, and every unselected edge contributes at most its two endpoint crossings. Thus, if P is the number of ports,

$$P \le 2E', \qquad \#\{\text{ordered rows}\} \le P^2. \tag{B.15}$$

This proves the claimed moments. The separator need not have finite metric length; it is a topological boundary only.

Move its finitely many ports to rational arclength parameters in their transverse edge charts. Choose the first feasible finite tuple in a fixed enumeration and select a Jordan curve realizing it. Injectivity, avoidance away from the selected contacts, local transverse charts and the winding condition are measurable curve relations. Analytic projection and completed-measurable selection therefore make both ports and separator measurable from the ordinary neighborhood. This is the same version convention as above.

Transversality supplies a finite rectifiable edge access to either side of every port. Any two nonsingular points in a connected open ordinary set have a finite internal connection. Indeed, an arbitrarily short surrounding loop and its short access connect each to a rational point in that set; a finite chain of rational disks connects the rational points. Fixed rational point connections in these disks are finite simultaneously by [4, Lemma 2.9]. Thus ordinary port rows in I are finite. A bounded minimizing sequence in the compact set I has a Euclidean-uniform limit and attains the infimum by lower semicontinuity. With an inner compact set collapsed, delete repeat visits to the free vertex. The remaining candidate has at most two ordinary arms, whose endpoint-to-compact-set infima are attained by the same compactness argument. No assertion about completion or attainment in an arbitrary open domain is made.

It remains to prove exactness for every filling. Let

$$c = \inf\{D(x,y) : x \in \partial W_0,\ y \in A\} > 0.$$

The infimum is positive by compactness, lower semicontinuity and the positive Euclidean separation. Take a vanishing-defect sequence of full queries. Trimming between first and last surrounding-loop contacts bounds its relevant cost by a fixed cap plus the defect. Every component of its visit to $\operatorname{int} W_0$ that reaches A costs at least $2c$, except at most two endpoint components if that variant is allowed. There are consequently only finitely many relevant episodes. Raw contacts with ∂W_0 may still be infinite.

Replace each relevant episode by a shortest ordinary path with the same endpoints in the compact closed set W_0. Its finite infimum is attained as above, and the replacement does not increase cost. Every strict fragment near A is a geodesic in the same open domain O: a shorter O competitor would also be a W_0 competitor. Hence its entire trace on A lies in the selected graph. Each tightened episode is simple and meets Σ only at the finite port set. Split at these

contacts. Inside pieces cost at least the corresponding closed local row; outside pieces cost at least the genuine exterior-restricted kernel. Tangencies can be omitted or represented by zero diagonal edges, and positive cycles can be deleted. This gives the graph lower comparison.

Conversely concatenate arbitrarily accurate physical representatives of the finitely many edges of a graph traversal. Local edges stay in their closed declared domains, exterior edges retain their outside restrictions, and any child body stays inside its declared complete modification region. This gives the reverse comparison after taking infima. The argument never tightens a path inside an opaque filling, and all ordinary replacements use the same fixed neighborhood. It is therefore simultaneous over positive fillings, including ties, directed internal traversals, and unattained hidden infima. Any collapsed zero-cost vertex is expanded into its two ordinary arms before a physical-curve estimate is applied. □

B.9 Exact coding buffers and two complete attachments

In this subsection f denotes a *physical* Weyl shift: the length multiplier is e^f. Thus a common floor of depth M means $f \geq -M$. All paths and replacements used in a coding buffer are ordinary paths, outside the hidden filling.

Lemma B.17 (Removing a fixed coding buffer from a wall limit)**.** *Let P be a simple unique geodesic in a fixed ordinary comparison domain, let B be a fixed Euclidean disk with compact closure in that domain, and suppose that no nonconstant D-rectifiable path lies in ∂B. Let $f_n \geq -M$ be smooth physical shifts that equal $-M$ on $P \cup \overline{B}$. The statement permits an additional shift v_n with $\|v_n\|_\infty \to 0$.*

Suppose ordinary query portions Q_n have a common base-length bound and converge uniformly in the Euclidean metric to a curve Q whose image lies in $P \cup \overline{B}$. Their full crossing endpoints lie outside $\overline{B}$. Each compared fragment minimizes the shifted length for $e^{f_n+v_n} \cdot D$ among curves in one fixed comparison region, and the full reference competitor used below is admissible in that region. The region may be a compact set contained in the ordinary open comparison domain. Then the image of Q lies in P. If, in addition, the full shift $f_n + v_n$ is constant on a fixed open neighborhood of the compared coding portion, sufficiently close crossing portions contain its prescribed protected reference segment exactly.

Proof. Parametrize by base arclength, allowing constant extensions. Partition lower semicontinuity makes Q a D-Lipschitz curve. Consider a nondegenerate interval component $[b, c]$ of $Q^{-1}(\overline{B})$. Its endpoints are interior times of the full crossing, and maximality gives $x = Q(b), y = Q(c) \in P$. An interval contained in ∂B is constant. Otherwise, after removing possible initial and final constant boundary intervals, choose $s, t \in (b, c)$ with $Q(s), Q(t) \in B$ and with the two limiting tail lengths tending to zero as $s \downarrow b$ and $t \uparrow c$.

Put $x' = Q(s), y' = Q(t)$. By [3, Lemma 3.2], these nonsingular points have separating loops of arbitrarily small length and arbitrarily small access cost, entirely in B. The connectors also lie in B when sufficiently short, since crossing from a fixed interior point to its boundary costs a positive amount. Choose loops small enough that $Q(b), Q(c)$ lie outside them. Uniform Euclidean convergence implies that Q_n meets the first loop before s and the second after t. Use the first and last contacts if the two loops coincide.

Between these actual contacts, use the following ordinary competitor: a loop arc and connector to x', the reversed limiting tail from x' to x, the unique reference segment $P[x, y]$, the reversed limiting tail from y to y', and the second connector and loop arc. Every part lies on the common floor. The reference segment may leave B, but remains in the fixed ordinary comparison domain.

Minimality and the global lower floor give, after cancelling e^{-M},

$$\operatorname{len}_D Q_n[\text{first contact}, \text{last contact}] \le e^{2\|v_n\|_\infty}\big(\operatorname{len}_D Q[b,s] + \operatorname{len}_D P[x,y] + \operatorname{len}_D Q[t,c] + o(1)\big),$$

where the final error denotes the chosen loop and access costs, and can be made arbitrarily small before sending n to infinity. The comparison is allowed by the whole-fragment replacement hypothesis, since the original fragment and the entire displayed competitor lie in the same comparison region. Lower semicontinuity, followed by removal of the two tails, yields

$$\operatorname{len}_D Q[b,c] \le \operatorname{len}_D P[x,y] = d(x,y).$$

Uniqueness identifies this component with the reference segment, up to constant intervals. Every point of Q outside P would belong to such an interval component, so no such point exists. This proves the first assertion without inferring metric endpoint convergence from Euclidean convergence.

For exact extraction choose a strict reference interval inside the open coding neighborhood, with two unique protecting barriers also contained there. Such barriers are available after Theorem B.14: choose very short accessible side approaches to two strict reference points, approximate them in D by rational points on the prescribed sides, and use strong confluence on the resulting unique deterministic-endpoint geodesics. Make the interval short relative to its positive exit cost from the coding neighborhood, so all these barriers are determined internally there. Uniform Euclidean convergence puts the relevant query fragments in the same neighborhood. There the changed metric is exactly a constant multiple of the old metric. The two-barrier separation argument used in Lemma B.4 forces the exact common segment. All contact replacements stay in this open constant-shift neighborhood. □

The no-boundary-path hypothesis in this lemma is supplied by Lemma 5.1, applied to a countable deterministic family of coding and band circles. The lemma applies to finitely many disjoint coding disks at once. It also applies to a reference arm that is merely rectifiable away from the coding disks, provided every disk attaches only to a unique reference interval and the entire interval between its two contact points is on the common floor. Only those intervals are used in the length comparison.

Lemma B.18 (Two complete attachments in the matched region)**.** *Fix disks $H \Subset D$ and a surrounding ordinary rectifiable loop outside D. All full query endpoints lie outside this loop. For the parameter-a pure law there is a finite positive-probability catalogue of smooth physical shifts supported compactly in $D \setminus \overline{H}$, with the following properties.*

There are two disjoint reference arms, each with a private seed core between two complete coding buffers, and with fixed protected segments in those buffers. A retained ordinary middle contains H, a middle bypass, and all other attachments. Every minimizing query meeting H, this middle, or either specified private core traverses both arms, in opposite directions. It contains the designated inward anchors, one on each arm. Its entire portion between them remains in a fixed compact subset of D and avoids the full private seed supports. These assertions hold for every positive inner filling that retains ordinary local replacements and ordinary positive-loop deletion outside its complete modification region.

These properties hold on a fixed nonempty box of additional smooth amplitudes. Their directions are constant on the complete coding neighborhoods, and their full supports, including transitions, lie inside the designated private cores, with positive clearance from H and every nonincident piece. For arbitrary fillings the corresponding restricted-row identities hold at the level of infima after

ordinary tightening. For an actual ordinary Gaussian filling, the common intrinsic middle is unique, whenever used by a full geodesic; matching the ordinary field on its complete confined domain therefore matches its curve and occupation measure.

Proof. We first describe the ordinary geometry and then the finite exhaustion. Choose nested deterministic disks $H \Subset U \Subset V \Subset D$ and two disjoint corridor regions crossing $V \setminus \overline{U}$. Their axes are used only topologically. Actual finite rectifiable connections in these corridors are supplied by finite rational-disk networks, as in the proof of Theorem B.16; no polygonal axis is assigned a metric length. Erase loops to obtain two simple disjoint arms.

In a small ordinary buffer on each arm, replace a short arc by a shortest internal connection. Choose it short enough that a competitor leaving this buffer pays a strictly larger positive exit cost. Retain a strict middle interval J_i of this connection. It is unique by Theorem B.14. In two disjoint smaller rational disks around distinct strict points of J_i, retain exact protected segments and their barriers. Keep slightly larger open disks at the same constant shift. The private core is a compact region between these two coding buffers; its entire source support is required to lie there. The arms and reference intervals can be shortened and the corridors narrowed so that the two coding buffers are the only complete outlets of each private pocket.

All remaining connections are placed downstream of the inward buffers, in the middle containing U. In particular a bypass connects the two inward arms without entering H. Where access to H is needed, choose finitely many deterministic boundary ports in the corridor template and interior and exterior tangent disks meeting ∂H only at the respective port. By [4, Lemma 2.9], the disk-center-to-prescribed-boundary-point intrinsic distance is finite for every $a > 0$; its normalized pth moment is finite for $0 < p < 2Q(a)/a$. Hence each selected port has a finite exterior access and an almost-sure interior access. The exterior curves and their costs are selected from the exterior metric. Interior accesses are existence witnesses when required for a conditional lower-support argument, and are not exterior observations. This construction uses finitely many specified boundary points, not access from every point of ∂H.

Fix the complete candidate network and the final finite lowering depth M, including any access lowering already prescribed by the source construction. The subsequent profiles are bounded below by $-M$ and equal $-M$ on the reference arms, their comparison intervals and the complete coding disks. They may vanish near H and outside D. All transitions to these required zero regions lie downstream of the inward buffers or outside the matched band. The floor is fixed before the following shrinking-tube and wall choices.

For each n, take a rational tube of width tending to zero about the arms, together with the fixed closed coding disks, and boundary neighborhoods of widths tending to zero about ∂U and ∂V. Put a smooth wall of height tending to infinity off this set in the closed matched band. Include the analogous side walls for each seed pocket, leaving precisely its two coded outlets. The coding disks, and slightly larger open constant-shift disks, are part of the allowed set from the outset. Where the latter enlargement is used, regard its larger deterministic boundary circle as the coding circle in Lemma B.17. Profiles vanish near H and outside D; every part of their support is therefore ordinary and observed.

The surrounding-loop replacement bounds the cost of any relevant minimizing passage by a finite cap U_0, uniformly in the inner filling and in the wall height. A fixed small amplitude box changes this cap and the floor by bounded factors. Consequently its ordinary base length is at most $e^{M+1}U_0$. Suppose no finite wall profile and no positive amplitude box forces the stated protected segments. At profile index n, choose a violating locally minimizing query with amplitudes of norm at most $1/n$. Pass to a fixed violated band or pocket, trim at the last departure-boundary contact before the first arrival-boundary contact, and use Euclidean path compactness.

Every limit lies on the reference arms, the closed coding disks and the two boundary circles. Otherwise it traverses a fixed small ordinary annulus a positive distance from this set, whose positive base crossing cost is multiplied by a wall height tending to infinity, contrary to the cap. A nonconstant rectifiable subpath cannot travel along either boundary circle by Lemma 5.1. It therefore cannot change between the two disjoint arms along a limiting boundary neighborhood. Each proper band crossing follows one arm, apart from excursions in its coding disks. Lemma B.17 removes those excursions. Inside each open constant-shift coding neighborhood, the protecting barriers force the exact prescribed segment, in the prescribed order and orientation. This contradicts the chosen failure.

For a private-core visit, apply the same argument to its pocket. The private core has positive Euclidean clearance from both coded outlets and all walls. A limiting visit therefore reaches the middle of its reference interval. If the visit returned through the same outlet, the approximating locally minimizing query would contain the same protected anchor on both sides of a nonconstant excursion. Deleting this positive ordinary loop is allowed and is strictly shorter. Thus every private-core visit is a full passage through that pocket, and reaches the retained middle on its inward side. The same contradiction treats full supports, including their transitions, since their compact containment was fixed before the walls.

We have proved the claimed traversal property: every full crossing of $V \setminus \overline{U}$ by an admitted minimizing query traverses one of two protected segments, and every private-core visit forces such a traversal. Each protected segment can occur at most once, because a second occurrence repeats an ordinary point with positive cost in between. A query entering U from outside V and subsequently leaving must therefore use the two different protected segments, exactly once each. Between them it cannot leave V and return: doing so would create additional full crossings. Its inward-anchor middle stays in $\overline{V} \Subset D$. It cannot visit either private support again, since the two-outlet pocket property would repeat an inward anchor. This proves both confinement to the matched region and independence of this middle from the private amplitudes. All other attachments were placed in this same middle by construction.

For an arbitrary compatible filling, start from a finite-cost near minimizer. Only finitely many visits reach the strict coding cores or seed cores, since each such visit crosses a fixed ordinary positive-cost buffer. Before tightening, choose compact closed seed-pocket buffers inside the ordinary open domain, disjoint from H, and containing the whole relevant reference intervals, their coding neighborhoods, the access loops and connectors, and the limiting tails used in the comparison. Tighten each episode in its fixed compact buffer. In the wall-limit contradiction every fragment between the two actual coding-loop contacts lies strictly inside that same pocket buffer. The length comparison in Lemma B.17 uses only these ordinary minimizing fragments; the rest of the query is unchanged. The same finite wall construction therefore gives exact anchor traversals for the tightened representatives. Remove repeated ordinary-anchor loops and pass to infima. Hidden minimizers need not exist. Directed or tied internal fillings cause no difficulty, because all altered restrictions and all internal states lie strictly inside their complete modification regions, away from these ordinary anchors.

For an actual ordinary Gaussian field, every used anchor-to-anchor middle is a strict subsegment of a longer ordinary geodesic in a fixed compact comparison region. Its uniqueness follows from the simultaneous internal conclusion of Theorem B.14, transferred through the finite deterministic shift catalogue by Cameron–Martin equivalence. If two normalized fields agree on the confined middle domain, their internal metric agrees there by locality. They consequently have the same middle curve and its same restricted arclength measure, with the prescribed common scalar change of units. This does not apply deterministic-endpoint uniqueness to random anchors. For directed

fillings retain the two orientations of the same confined intrinsic subproblem separately.

Finally, all disks, corridor neighborhoods, tangent-disk port templates, rational tubes, smooth plateau profiles, integer wall heights and rational amplitude radii range over countable catalogues. Strict separation and the local bounded-length bad-path relation make their defining events completed-measurable, by the coding above. The pathwise construction covers a full-probability ordinary event. At least one finite initial subcatalogue therefore has positive probability. Cut off its positive ordinary crossing costs, finite candidate lengths, finite Dirichlet energies and Gaussian scores. This retains positive probability and gives a finite deterministic catalogue. Conditional Cameron–Martin use of this catalogue compares the score-bounded source event with its shifted image; it does not assert a lower success probability for every harmonic mean. All observed geometric data are outside H. □

References

[1] J. Pfeffer, *Weak Liouville quantum gravity metrics with matter central charge* $\mathbf{c} \in (-\infty, 25)$, Probability and Mathematical Physics 5 (2024), 545–608. arXiv:2104.04020v3.

[2] E. Gwynne and J. Pfeffer, *Bounds for distances and geodesic dimension in Liouville first passage percolation*, arXiv:1903.09561v3.

[3] J. Ding and E. Gwynne, *Regularity and confluence of geodesics for the supercritical Liouville quantum gravity metric*, arXiv:2104.06502v4.

[4] J. Ding and E. Gwynne, *Uniqueness of the critical and supercritical Liouville quantum gravity metrics*, arXiv:2110.00177v3.

[5] E. Gwynne and J. Miller, *Confluence of geodesics in Liouville quantum gravity for* $\gamma \in (0, 2)$, arXiv:1905.00381v3.

[6] E. Gwynne, J. Pfeffer, and S. Sheffield, *Geodesics and metric ball boundaries in Liouville quantum gravity*, arXiv:2010.07889.

[7] E. Gwynne, *Geodesic networks in Liouville quantum gravity surfaces*, arXiv:2010.11260v2.

[8] M. Ang, *Comparison of discrete and continuum Liouville first passage percolation*, arXiv:1904.09285.

[9] M. Beck and R. Sanyal, *Combinatorial Reciprocity Theorems*, Graduate Studies in Mathematics 195, American Mathematical Society, 2018.

[10] M. Bhatia and K. Kavvadias, *Strong confluence of geodesics in Liouville quantum gravity*, arXiv:2512.09219v2, 24 August 2026.

[11] W.-K. Chen, N. Dafnis, and G. Paouris, *Improved Hölder and reverse Hölder inequalities for Gaussian random vectors*, Advances in Mathematics 280 (2015), 643–689. arXiv:1306.2410.

[12] J. Ding and E. Gwynne, *Tightness of supercritical Liouville first passage percolation*, arXiv:2005.13576v2.

[13] E. Gwynne and J. Miller, *Existence and uniqueness of the Liouville quantum gravity metric for* $\gamma \in (0, 2)$, arXiv:1905.00383v3.