Pseudocritical point of the spread-out Ising model with free boundaries at the upper critical dimension in $1 \le d \le 4$

Jiwoon Park

November 20, 2026

Abstract

We study the susceptibility of the spread-out Ising model on a cube of volume $V = n^d$ with free boundary conditions. The interaction $(-\Delta_R)^{d/4}$ places the model at its upper critical dimension for each $d \in \{1, 2, 3, 4\}$. For sufficiently large fixed spread R, we prove that there is an inverse temperature at which the effective potential of the principal mode has zero curvature. At this pseudocritical point the susceptibility has order $\sqrt{V \log V}$. At the bulk critical point, Camia, Jiang and Newman obtain order n^2 for nearest-neighbour Ising boxes with free boundaries in $d \ge 4$ under a bulk two-point hypothesis [12]. The logarithmic factor at our pseudocritical point also occurs for the critical weakly coupled $|\varphi|^4$ model on the four-dimensional torus [31] and in the free-boundary pseudocritical window of the four-dimensional hierarchical model [28].

Version note: you can find earlier versions on DOI:10.5281/zenodo.24167641

v1. November 20, 2026

1 Model and main result

Fix $d \in \{1, 2, 3, 4\}$, and write $\alpha = d/4$, $s = 2\alpha = d/2$. On $\mathbb{Z}^d$ let

$$D_R(x) = \frac{\mathbf{1}_{\{0<|x|_\infty \le R\}}}{(2R+1)^d - 1}, \qquad -\Delta_R = L_R = I - D_R, \qquad \Lambda = \{1, \dots, n\}^d, \qquad V = n^d. \tag{1}$$

Here D_R also denotes the convolution operator. We also let P_Λ is the simple restriction $\ell^2(\mathbb{Z}) \to \ell^2(\Lambda)$ and P_Λ^* be the extension of $\ell^2(\Lambda)$ to $\ell^2(\mathbb{Z})$ by zero. On the full lattice, L_R^α is the Fourier multiplier $(1 - \widehat{D}_R(k))^\alpha$, with value zero at $k = 0$. The physical interaction is $P_\Lambda L_R^\alpha P_\Lambda^*$. Thus the power is taken before restriction, and free boundaries mean deleting interaction bonds to the exterior. The original diagonal is retained and contributes a spin-independent factor. We consider the spread-out Ising model with free boundary condition at inverse temperature τ defined by the measure

$$\mathbb{E}_\tau^F[F(\sigma)] = \frac{1}{Z_\Lambda^F} \sum_{\sigma \in \{-1,1\}^\Lambda} F(\sigma) \exp\left[-\tfrac{\tau}{2} \langle \sigma, P_\Lambda L_R^\alpha P_\Lambda^* \sigma \rangle\right] \tag{2}$$

where Z_Λ^F is the normalisation and the susceptibility

$$\chi_\Lambda^F(\tau) = V^{-1} \mathbb{E}_\tau^F(\mathbf{1}^\mathsf{T} \sigma)^2. \tag{3}$$

For $d < 4$ the interaction is long-range. At $d = 4$ it is the local operator L_R, with R large as required by the Ising initialization. Quartic coupling is marginal in all dimensions because $d = 4\alpha$.

On a cube with free boundaries the principal mode of the Gaussian reference gives the leading contribution to the susceptibility. We fix the reference and its decomposition into scales, then vary the physical temperature until the terminal effective potential of this mode has zero curvature at the origin. This is soon explained in Section 1.0.1.

Theorem 1.1 (Pseudocritical susceptibility)**.** *Fix $d \in \{1,2,3,4\}$. For all sufficiently large fixed R and all $n \geq n_0(R)$, there is an inverse temperature $\tau_\Lambda^{\mathrm{pc}}$ at which the terminal effective potential has zero curvature at the origin, and*

$$c_{d,R}\sqrt{V \log V} \;\leq\; \chi_\Lambda^F(\tau_\Lambda^{\mathrm{pc}}) \;\leq\; C_{d,R}\sqrt{V \log V}. \tag{4}$$

Here $0 < c_{d,R} \leq C_{d,R} < \infty$ are independent of n.

In terms of side length the four orders are $\sqrt{n \log n}$, $n\sqrt{\log n}$, $n^{3/2}\sqrt{\log n}$, and $n^2\sqrt{\log n}$, respectively. In $d = 4$, this is comparable to that with periodic boundary condition at the bulk critical point proved in [31], justifying the 'criticality' of $\tau_\Lambda^{\mathrm{pc}}$. Also, note that the susceptibility at the bulk critical point lacks the logarithmic correction as show in [12], showing a trait of a near-critical Gaussian scaling [31, Theorem 1.9(2)].

1.0.1 The matching condition for $\tau_\Lambda^{\mathrm{pc}}$

Let us first introduce some essential terminologies.

- Let ℓ be the dual principal vector of the full reference. This is defined in Section 2, and as $n \to \infty$, equivalent to v_g, the eigenvector of $P_\Lambda L_\Lambda^\alpha P_\Lambda^*$ with the smallest eigenvalue. Thus, everything below can be interpreted with v_g in place of ℓ.
- C_T is the terminal covariance, defined in Section 2.3.2. It contains the Gaussian fluctuation of the field in the direction of ℓ. The remaining covariance is denoted $C_<$, containing the bulk Gaussian fluctuation of the field, so that $C_< + C_T$ represents the covariance of the Ising model, when we view the Ising model as a Gaussian field restricted to values ± 1. This view is particularly useful in RG implementations.

Let F_T denote the remaining interaction factor after integrating the fluctuation $C_<$ and ϕ for the terminal Gaussian field with covariance C_T. With $a = \ell^{\mathsf{T}}\phi$, we can decompose ϕ as

$$\sigma_T^2 = \ell^{\mathsf{T}} C_T \ell, \qquad u_T = C_T\ell/\sigma_T^2, \qquad \phi = a u_T + Y \tag{5}$$

with Y independent of a. The density of a is proportional to

$$e^{-a^2/(2\sigma_T^2)}\, \mathbb{E}^Y F_T(a u_T + Y) \tag{6}$$

where the expectation is over Y. Noting that this is now defined a distribution on $\mathbb{R}$, the second derivative at zero of the negative logarithm of this density defines the curvature in the theorem. The condition is invariant under a nonzero linear rescaling of a. In dimension 4, the actual implementation differs, only because we first need a quadratic counterterm to tune the reference field.

We also prove in the companion paper [1] that $\Delta_n = \tau_{\Lambda_n}^{\mathrm{pc}} - \tau_c$ is a non-trivial shift from the bulk critical point $\tau_c = \tau_c(d,R) = \inf\{\tau > 0 : \langle\sigma_0\rangle_\tau^+ > 0\}$, where $\langle\cdot\rangle_\tau^+$ is the infinite-volume plus state for the Hamiltonian used in (2). Write $(u)_+ = \max\{u, 0\}$.

Theorem 1.2 (Shift of the pseudocritical point)**.** *[1] For $d \in \{1,2,3,4\}$ and sufficiently large fixed R, $\tau^{\mathrm{pc}}_{\Lambda_n} \to \tau_c$ as $n \to \infty$. For $n \geq n_1(d,R)$ one has $(\Delta_n)_+ \leq C_d R^{-d}$. For $d = 1,2,3$ one also has*

$$(\Delta_n)_+ \leq C_{d,R} \begin{cases} n^{-1/10}(\log n)^{1/10}, & d = 1,3, \\ n^{-1/5}(\log n)^{1/2}, & d = 2. \end{cases} \tag{7}$$

There is $c_{d,R} > 0$ such that, for all sufficiently large n,

$$\Delta_n \geq c_{d,R}\, n^{-s} = c_{d,R}\, n^{-d/2}. \tag{8}$$

1.1 Historical context and related results

Long-range Ising criticality. Dyson proved the existence of a phase transition for the ferromagnetic chain with interaction proportional to $|x-y|^{-1-s}$, $0 < s < 1$ [14]. For an interaction of order $|x-y|^{-d-s}$, Fisher, Ma and Nickel identified $d = 2s$ as the upper critical dimension in their renormalisation group analysis and derived bulk logarithmic corrections at that dimension [16]. Aizenman and Fernández proved mean-field values of bulk critical exponents for a class of reflection-positive long-range Ising interactions when $d < 4$ and $0 < s \leq d/2$ [5]. For the periodic system at $d = 2s$, Luijten and Blöte derived the finite-size susceptibility prefactor $L^{d/2}(\log L)^{1/2}$ and tested it numerically, including for the chain with $s = 1/2$ [26, Eq. (27) and Table V].

Shift and rounding. Fisher and Barber distinguished the shift of the effective finite-volume critical point from the rounding of the transition [15]. Brézin and Zinn-Justin computed an effective Hamiltonian for the homogeneous mode by integrating out the other modes [10]. Above the upper critical dimension the quartic coupling is dangerously irrelevant for the homogeneous mode: although it decreases under rescaling, it determines the scale of the spatial average at the effective critical point. Rudnick, Gaspari and Privman analysed the free-boundary analogue using standing-wave modes and found a boundary-induced shift of the effective critical temperature [33]. Wittmann and Young later found numerical evidence in the five-dimensional Ising model for susceptibility of order L^2 at the bulk critical temperature and $L^{d/2}$ at a free-boundary pseudocritical temperature [36]. In this historical discussion, L denotes the side length of a d-dimensional box and $V = L^d$ its volume.

Bulk scaling and finite-volume observables. Bulk Gaussian scaling and the finite-volume law of the spatial average use different observables and normalizations. Aizenman and Fröhlich proved Gaussianity of continuum limits for high-dimensional spin or scalar field models [2, 18]. Aizenman and Duminil-Copin established marginal triviality at $d = 4$ for the nearest-neighbour Ising model and related scalar models, and Panis extended Gaussianity to critical reflection-positive Ising and scalar models of effective dimension at least four [4, 29]. On periodic lattices, Park's results for weakly coupled $|\varphi|^4$ models give a Gaussian field limit when tested against mean-zero functions and a quartic law for the spatial average at its own normalization [31, Theorem 1.9]. These observables probe different parts of the fluctuation spectrum at criticality.

Self-avoiding walk and percolation. Self-avoiding walk provides a path analogue. Above its short-range upper critical dimension four, the lace expansion gives a Brownian scaling limit on the infinite lattice, first for weak self-avoidance and then for nearest-neighbour strict self-avoidance [11, 22]. For weakly self-avoiding walk on a torus in $d > 4$ with sufficiently small self-repulsion, Michta obtains Brownian motion on the continuum torus when the step count satisfies $m = o(V^{1/2})$ and m/L^2 tends to a positive limit. The diffusion constant agrees with that on the infinite lattice

[27]. At the subcritical edge of the finite-size window, where the critical fugacity is approached on the scale $V^{-1/2}$, the torus susceptibility is of order $V^{1/2}$. The two-point function has a bulk power-law term and a plateau of order $V^{-1/2}$. Slade proved this for weakly self-avoiding walk, and Liu established the corresponding bounds for nearest-neighbour strict self-avoiding walk [35, 24]. The unwrapping picture associates the plateau with the sum over windings around the torus, as illustrated in numerical studies of self-avoiding walk and the Ising model [19].

For short-range percolation above dimension six, the lace expansion gives mean-field critical behaviour and random-walk power-law decay of the bulk connectivity for the models covered by the expansion [21, 20]. Aizenman used mean-field bounds to study the proliferation of incipient spanning clusters in critical boxes [3]. The finite-graph triangle condition of Borgs, Chayes, van der Hofstad, Slade and Spencer yields the Erdős–Rényi scales $V^{-1/3}$ for the critical window and $V^{2/3}$ for the largest clusters [8, 9]. In the high-dimensional torus window, Hutchcroft, Michta and Slade prove a two-point plateau of order $V^{-2/3}$ and susceptibility of order $V^{1/3}$. At the bulk threshold the free-box susceptibility has order L^2 in the setting of Chatterjee and Hanson [23, 13]. For sufficiently high-dimensional nearest-neighbour and sufficiently spread-out models with $d > 6$, Blanc-Renaudie and Nachmias identify the limiting ranked torus cluster volumes, after model-dependent rescaling, with the ranked excursion lengths of Brownian motion with parabolic drift above its past minimum [7]. This is the random-graph law for cluster sizes.

The pseudocritical window. Liu, Park and Slade put the torus plateau estimates into a common framework by comparing the periodic two-point function with the sum of the infinite-lattice two-point function over periodic images. Their comparison theorem gives the scale of the susceptibility and the plateau, and they verify its hypotheses for spread-out long-range self-avoiding walk above its upper critical dimension. They formulate conjectures for the profiles throughout the scaling window and for the corresponding free-boundary pseudocritical behaviour [25]. For spins, the distinction between the bulk critical temperature and the centre of the finite-size window already has rigorous counterparts. Camia, Jiang and Newman obtain the L^2 critical susceptibility for free-boundary nearest-neighbour Ising boxes in $d \geq 4$, assuming that the bulk critical two-point function has order $|x - y|^{2-d}$ [12]. In the hierarchical $|\varphi|^4$ model, Michta, Park and Slade prove $\sqrt{V}$ susceptibility in the periodic and free-boundary windows, each centred at its own effective critical point, with an additional logarithmic correction in $d = 4$. For $d > 4$ the free-boundary shift is of order L^{-2}, larger than the window width $L^{-d/2}$ [28]. This clearly distinguishe the finite-size scalings with different boundary conditions. This distinction and the presence of pseudocritical point also predicted by [17].

Theorem 1.1 addresses the marginal Euclidean box. It gives susceptibility bounds at a point specified by the exact principal-mode curvature. For the same spread-out interaction on $\mathbb{Z}^d$, Park proves critical two-point decay of order $|x|^{-d/2}$ [32, Theorem 1]. Here the principal profile of the finite cube enters both the quartic energy and the coupling to a uniform magnetic source.

1.2 Comparison with other RG formulations

The perturbative flow and the estimates for a single renormalisation group step have counterparts in the lattice constructions of Bauerschmidt, Brydges and Slade [6]. Park's map treats weakly coupled $|\varphi|^4$ interactions with periodic boundary conditions [30]. Slade's long-range analysis treats weak coupling at a positive distance below the upper critical dimension [34]. Here the interaction starts from the Ising spin measure, and the estimates must hold at $d = 2s$ on blocks meeting a physical boundary. The pressure expansion of Section 2.2 and the boundary estimates of Section 3 address these two features.

In the hierarchical model the renormalisation group identifies both the free-boundary critical window and its displacement from the bulk critical point [28]. On the Euclidean cube Theorem 1.1 estimates the susceptibility at a matched temperature. Determining the displacement of this temperature from the bulk critical point is a separate question.

Two comparisons clarify the choice of reference Gaussian measure (equivalently, choice of the counterterm) in dimension four. Appendix C treats the Gaussian measure with a fixed kinetic coefficient. Lemma C.5 gives the normalized fluctuation expansion in an independent background, and Proposition C.6 exhibits second-order terms with support. Using this expansion in the original coordinates would also require a local representation of the transformed background.

1.3 Outline of the proof

We use the renormalisation group flow to control the flow of the coupling constants. The running quartic copuling constant appears as the logarithmic correction in Theorem 1.1—at $d = 2s$, the quartic interaction is marginal both on the full lattice and in the interior of a finite cube.

In [32, Theorem 1 and Sections 7, 10 and 22] the RG is performed on periodic boxes. Translation invariance gives one set of bulk coefficients, and the thermodynamic limit at positive auxiliary mass before computing the limit to the critical two-point function on $\mathbb{Z}^d$. Here, with free boundary condition, we keep the cube finite and tune the temperature to the curvature of its principal mode. Deleting bonds at a face changes the short-range spin pressure already at the initial scale. Integrating a fluctuation then produces a position-dependent quadratic interaction even from a quartic with constant coefficient. The central challenge is to control these terms.

From the spins to a local RG step. We use the positive Gaussian split and the conditional Gaussian representation of [32, Propositions 2.4 and 17.1], with the local split of Lemma 2.1 at $d = 4$. After subtraction of the bulk mass, the difference between the cube and bulk interaction Hessians decays exponentially in the distance to the boundary in units of R (Lemma 2.3). At the initial block size $b_0 \asymp R \log R$, Proposition 2.16 separates the quadratic terms of range $O(b_0)$ from the tails. Proposition 5.1 assigns the tails to a polymer activity so that its value on every polymer avoiding the grouped boundary labels agrees exactly with its bulk value. In [32, Sections 6–7 and 17–18] one can absorb the translation-invariant quadratic correction into the Gaussian covariance. Here we keep the boundary correction in the interaction.

For the Gaussian integration we apply the positive polynomial decomposition of [32, Lemma 2.1] to the finite-range auxiliary precision. A polynomial of degree m in this precision has range $O(mR)$. Thus a fluctuation covariance row agrees with the bulk row whenever its propagation neighbourhood lies inside the cube. This property holds also for the fractional interaction, whose original spin couplings have infinite range.

Reblocking at a face. Let b be a fine block size and $B = Lb$ the next size. A coarse block at a face contains $O(L^{d-1})$ fine blocks within distance $O(b)$ of the boundary. The field radius changes from $\ell_b = b^{-s/2}$ to $\ell_B = L^{-s/2}\ell_b$. For an unrestricted quadratic the resulting factor is $L^{d-1-s} = L^{s-1}$, attained on constant fields (Lemma 3.1). It gives no contraction in dimensions two through four.

The covariance itself supplies the missing power. Set $\rho_b(x) = \min\{b, R + \operatorname{dist}(x, \Lambda^c)\}$. For an ordinary fluctuation slice Γ_j at scale b, the boundary resolvent estimate of Proposition 2.9 gives, with λ the covariance parameter of Section 2,

$$\Gamma_j(x,x) \le C_L \lambda b^{-s} \big(\rho_b(x)/b\big)^{2p}, \qquad 0 < p < s/2. \tag{1.9}$$

A positive supersolution for each face bounds the killed resolvent near the boundary, and there is a freedom of choosing $p \in (0, s/2)$ in the construction of Γ_j. Taking the minimum of the supersolutions covers points near several faces, and applying the estimate on both observation legs gives the square of the weight in (1.9). The annular difference estimates of Section 2.3 give Gaussian moments in the corresponding weighted field norm. On the old boundary layer ρ_B/ρ_b is bounded. Restricting a coarse test field to a fine block in this layer therefore gains L^{-p} in each argument. For a quadratic, the combined factor is

$$L^{d-1}\big(L^{-s/2}L^{-p}\big)^2 = L^{s-1-2p}. \tag{1.10}$$

We choose $\max\{0, (s-1)/2\} < p < s/2$, with $p \neq 1/2$. Lemma 3.6 then gives a strict contraction for large L, including the contributions from edges and corners.

There is a separate estimate for the quadratic created at the current step. Contracting two legs of the quartic gives a coefficient proportional to the diagonal of the filtered slice. Its difference from the bulk value can extend to distance $O(B)$ from a face, so all $O(B^d)$ sites in a coarse block can contribute. Lemma 3.5 bounds this sum using the slice diagonal and the coarse field radius. We retain the resulting quadratic on that block. At the next step its support is already in the old boundary layer, and (1.10) applies. Corollary 3.7 sums these contributions over the scales. The bound is uniform in the side length, with the surface interaction fixed by the microscopic spin measure.

Localisation and the quartic flow. On a boundary block we keep the full quadratic, quartic and sextic Taylor coefficients. In the interior the bulk symmetries allow a smaller set of coefficients. Proposition 3.2 matches the quadratic Taylor coefficient on pairs of spatial polynomials of total degree at most two, and the quartic on constant fields. The resulting Taylor remainder contracts after summation over the fine blocks. When a covariance contraction joins different blocks, Lemma 3.4 assigns its terms to the input labels before reblocking. This keeps the second-order subtraction exact across block interfaces. The bulk quartic consequently has the same leading decrement $\mathfrak{b}_j g_j^2$ on the cube interior. Proposition 2.14 bounds the bubble increment $\mathfrak{b}_j$ above and below by multiples of λ^2 on the ordinary scales away from the terminal cut. The reciprocal increments along the trajectory of Theorem 6.2 give $g_N \asymp_{d,R} (\log V)^{-1}$.

Matching the mass and the kinetic coefficient. For $d < 4$ a bulk quadratic with two spatial differences scales by $L^{s-2} < 1$. Corollary 3.16 controls this coefficient together with the boundary tensors and the polymer activity. We solve the mass recurrence backwards from the equation of zero terminal curvature. The microscopic mass has nonzero thermal derivative, so varying the temperature crosses the required initial mass (Corollary 6.3).

In $d = 4$ the kinetic factor L^{s-2} equals one. Proposition 4.5 estimates the increment of this coefficient, and Lemma 6.6 sums the increments in the backward recurrence. We vary an auxiliary Gaussian reference to match the initial kinetic coefficient as well as the mass. The terminal curvature depends on this choice of reference. To impose the curvature prescribed in Section 1.0.1, we change the microscopic field so that its covariance with the original coordinate is the original column $C_T\ell$. Stopping one scale earlier adds a variance margin of order λn^2 in the Gaussian regression (Proposition 7.12). The conditioned covariance therefore stays positive throughout the matching rectangle. Corollary 7.13 uses it to impose the original curvature condition at the physical temperature.

The magnetic moment. In the natural terminal field write $\phi = au_T + Y$, as in Section 1.0.1. Its profile satisfies $\sum_x u_T(x)^4 \asymp V^{-1}$ and $|\mathbf{1}^{\mathsf{T}} u_T| \asymp V^{1/2}$. The quartic along this direction has size

$g_N a^4/V$, which puts a on the scale $(V/g_N)^{1/4}$. After integration over Y, Proposition 7.4 gives a magnetic second moment of order $V^{3/2} g_N^{-1/2}$. Restoring the uniform source in the earlier slices costs $O_R(V^{3/2})$. The positive-row comparison of Lemma 7.5 transfers the moment to the spin variables. In dimension four this moment is computed in the transformed field at the earlier cut. The uniform source there gives a spin weight bounded above and away from zero at every site, so the same comparison applies (Corollary 7.13). Division by V gives the factor $\sqrt{V \log V}$ in Theorem 1.1.

Contents

1 Model and main result **1**

2 Microscopic and Gaussian estimates **7**

3 Boundary transport and localization **29**

4 The marginal kinetic coefficient in dimension four **59**

5 Physical initial data and terminal covariance **77**

6 A dimension-independent shooting theorem **85**

7 Terminal integration and return to spins **91**

A Polynomial covariance slices and terminal conditioning **106**

B Resolvent derivatives at the local endpoint **108**

C Gaussian integration at a fixed kinetic coefficient **110**

2 Microscopic and Gaussian estimates

The microscopic split has two roles. Its short-range part gives a pressure whose quadratic coefficient can be varied with the temperature. Its positive part gives a Gaussian reference whose principal variance can be prescribed before the temperature is chosen. We first construct the split and control the short-range pressure. We then estimate the principal mode, the remaining covariance, and the bulk bubble that drives the quartic flow.

For $d < 4$, we use the expansion

$$I - L_R^{\alpha} = \sum_{m \geq 1} c_m D_R^{*m}, \qquad c_m = \frac{\alpha\, \Gamma(m-\alpha)}{\Gamma(1-\alpha)\Gamma(m+1)} > 0. \tag{2.1}$$

Use the positive decomposition of [32, Proposition 17.1] and split the series at a sufficiently large fixed even N_{sm}. The short part contains the terms below N_{sm}, and the long part contains the remaining terms. After diagonal regularization and scalar normalization, the long part gives $A > 0$, the short part gives $Q \geq 0$ entrywise, and $A+Q$ has bulk row sum 1. We write $B = I - Q$. Positivity of the long part follows from the beta integral for $\widehat{D}_R(k)^{N_{\mathrm{sm}}}/(1 - u\widehat{D}_R(k))$ with a positive measure on $u \in (0,1)$. For finite matrices we suppress Λ when there is no ambiguity. The short part has finite range $O(R)$, off-diagonal entries $O(R^{-d})$ and row sum bounded away from one. Retain the diagonal of Q in the definitions of A and B. It contributes only a constant to the spin energy and

drops out of the normalized short-range spin pressure. The temperature convention for this split is $\beta = (1 + e^{-R})\tau$.

To obtain a finite-propagation decomposition for the Gaussian reference, we realize the long part on an auxiliary space. On the continuous channel take $\mathcal{D}_u = I - uD_R$ and $\mathcal{L}_u = (\zeta w_\alpha(u))^{1/2} D_R^{N_{\mathrm{sm}}/2}$, where

$$\zeta = (1 + e^{-R})^{-1}, \qquad w_\alpha(u) = \frac{\alpha\, u^{N_{\mathrm{sm}}-\alpha-1}(1-u)^\alpha}{\Gamma(1-\alpha)\Gamma(1+\alpha)}. \tag{2.2}$$

The white channel has diagonal I and coupling $(\zeta e^{-R})^{1/2} I$. With $\mathcal{L}_\Lambda = \mathcal{L} P_\Lambda^*$, the two-block precision has diagonal blocks $\mathcal{D}, B_\Lambda/c$ and off-diagonal blocks $-\mathcal{L}_\Lambda, -\mathcal{L}_\Lambda^*$. The observed covariance is its inverse tested through $\mathcal{O} = (\mathcal{L}_\Lambda^*, 0)$, which gives V_c by the Schur complement. The continuous auxiliary inverse is interpreted as a quadratic form on the observation space.

We use the same field normalization for both splits. Write $g_{\mathrm{mic}} = R^{-d}$, $A_{\mathrm{sc}} = g_{\mathrm{mic}}^{-a_*}$, with $1/12 < a_* < 1/10$, and $\lambda = A_{\mathrm{sc}}^{-2}$. The canonical field is defined by $h_{\mathrm{phys}} = A_{\mathrm{sc}} R^{-d/4} \phi$. It gives canonical covariance $a_R V_c$, where

$$a_R = \lambda R^s, \qquad V_c = (A^{-1} - cB^{-1})^{-1}, \qquad E = B^{-1/2} A B^{-1/2}, \qquad c_{*,\Lambda} = \eta_1(E)^{-1}. \tag{2.3}$$

The split is chosen so that $B - A$ is a positive scalar multiple of the compressed L_R^α. Write e for the unit top eigenvector of E with $B^{-1/2}e > 0$, and $\ell = B^{-1/2}e$. For $c = c_{*,\Lambda} - \theta_0 \lambda^2 \bar{g}_N (R/n)^s$ the full principal variance is $n^s/(\theta_0 \lambda \bar{g}_N)$. The parameter $\bar{g}_N = (G_{\mathrm{ref}}^{-1} + \lambda^2 N)^{-1}$, with $G_{\mathrm{ref}} = A_{\mathrm{sc}}^4 g_{\mathrm{mic}}$, depends only on the reference scales, so this construction precedes the temperature variation.

2.1 The microscopic split and the thermal crossing

For a thermal crossing, the short-range reference needs a row sum bounded away from one uniformly in R. In dimension four the binomial series for $I - (I - D_R)^\alpha$ terminates at degree one. Taking the original even-tail split with $N_{\mathrm{sm}} \geq 2$ would leave only the white regularization in A. The short-range row sum would then tend to one and the uniform high-temperature gap would be lost. For the spread-out operator we use the convolution-square construction of [32, Section 3], with the explicit kernel below.

Lemma 2.1 (Positive split for the local power)**.** *Take $d = 4$ and let D_R be uniform on the nonzero sites of the spread-out cube. Fix a positive integer m and let E_R be uniform, including zero, on the cube of radius $\lfloor R/(4m) \rfloor$. Write $F_R = E_R^{*2m}$. For a fixed sufficiently small $\kappa > 0$ and $\epsilon_R = e^{-R}$, put*

$$\begin{aligned} t_R &= \kappa F_R(0) + \epsilon_R, \qquad A = \frac{\kappa F_R + \epsilon_R I}{1 + t_R}, \\ Q &= \frac{D_R - \kappa F_R + \kappa F_R(0) I}{1 + t_R}, \qquad B = I - Q. \end{aligned} \tag{2.4}$$

For sufficiently large R, $A > 0$, Q is nonnegative entrywise and has zero diagonal, and $A + Q$ has row sum one. The bulk row sums $p = \sum_x A(x)$, $q = \sum_x Q(x)$ satisfy $p + q = 1$ and $p \geq \kappa/2$. Furthermore, $\sup_x |Q(x)| \leq C_m R^{-4}$ and

$$B - A = (I - D_R)/(1 + t_R). \tag{2.5}$$

The physical inverse temperatures are related by $\beta = (1 + t_R)\tau$, and the replacement does not change the finite-volume spin measure.

Proof. The convolution F_R has support in the cube of radius $R/2$, row sum one and supremum at most $\|E_R\|_\infty \le C_m R^{-4}$. Every nonzero site in its support has $D_R(x) \ge cR^{-4}$. Taking κ small gives the off-diagonal positivity of Q. The Fourier multiplier of F_R is $\widehat{E}_R^{2m} \ge 0$. The white term makes A strictly positive, including on principal restrictions. Direct addition gives $A+Q = (D_R+t_R I)/(1+t_R)$. Its diagonal correction contributes a constant in Ising variables, since $\sigma_x^2 = 1$. □

The proof also works for a uniform Euclidean spread-out ball by using a smaller inscribed smoothing cube. The smallness R^{-4} is a large-coordination property. It is absent for a fixed nearest-neighbour kernel. The identity $s = 2$ below concerns the local power of the spread-out Laplacian.

For $d < 4$, the original positive binomial coefficients give the same short-range hypotheses. Indeed, the finite head has a row gap equal to the positive tail mass, uniformly at fixed N_{sm}, and its maximum entry is $O(R^{-d})$ because $\|D_R^{*i}\|_\infty \le \|D_R\|_\infty$. Write $Q^\circ = Q - Q(0)I$ for the spin graph. Its normalized pressure equals that of Q. The two walk resolvents differ by $\beta Q(0)(I-\beta Q)^{-1}(I-\beta Q^\circ)^{-1}$, whose absolute exponentially weighted row sum is $O(R^{-d})$. On restricting to the cube, a term in the difference from the bulk must also visit the exterior. The unused walk length weight supplies the boundary factor in Lemma 2.2. Thus that lemma applies with the frozen $B = I - Q$ used in the Gaussian reference.

In dimension four the auxiliary realization can use two copies of the lattice and a finite-range observation coupling

$$\mathcal{L} = \frac{1}{\sqrt{1+t_R}} \begin{pmatrix} \sqrt{\kappa}\, E_R^{*m} \\ \sqrt{\epsilon_R}\, I \end{pmatrix}, \qquad \mathcal{Q}_c = \begin{pmatrix} I & -\mathcal{L}_\Lambda \\ -\mathcal{L}_\Lambda^* & B_\Lambda/c \end{pmatrix}. \tag{2.6}$$

Here $\mathcal{L}_\Lambda = \mathcal{L}P_\Lambda^*$. At $r \ge 0$ the observed bare resolvent is $A_\Lambda/(1+r)$, and the physical Schur complement of $\mathcal{Q}_c + rI$ is $B_\Lambda/c + rI - A_\Lambda/(1+r)$. In particular $\mathcal{O}\mathcal{Q}_c^{-1}\mathcal{O}^* = V_c$ for $\mathcal{O} = (\mathcal{L}_\Lambda^*, 0)$. Polynomials in the block precision therefore give finite-range operators in the observed coordinates. Section 2.3 estimates their covariance up to the boundary.

The Hessian of the short-range pressure fixes the initial mass. The next lemma compares it with the walk resolvent in the bulk and estimates the error caused by deleting exterior bonds. We use the first bound to differentiate the bulk mass with respect to temperature and the second to localize the initial quadratic interaction near the boundary.

Lemma 2.2 (A weighted short-range Hessian estimate)**.** *Suppose the symmetric kernel Q is nonnegative, has zero diagonal and range at most C_0R, and satisfies $\sup_x \sum_y Q_{xy} \le q_0 < 1$ and $\sup_{x,y} Q_{xy} \le C_0 g$. For g sufficiently small there are $a, r > 0$ such that for $|\beta - 1| < r$ and every finite induced set U, with $S_{\beta,U} = D^2 \log Z_{\beta Q,U}(0)$, the error $E_U = S_{\beta,U} - (I - \beta Q_U)^{-1}$ obeys*

$$\sup_{x\in U} \sum_{y\in U} e^{a|x-y|/R} |E_U(x,y)| \le Cg. \tag{2.7}$$

If $U = \widetilde{U} \cap \Lambda$, then, with $\delta(x) = \mathrm{dist}(x, \Lambda^c)$,

$$\sup_{x\in U} \sum_{y\in U} e^{a(\delta(x)+|x-y|)/R} |E_{\widetilde{U}}(x,y) - E_U(x,y)| \le Cg. \tag{2.8}$$

The bounds pass to the bulk limit and give thermal derivative bounds on a smaller complex disc.

Proof. Write $t_{xy} = \tanh(\beta Q_{xy})$. By the support bound, the gap, and $\tanh z = z + O(z^3)$, there are $\rho < 1$ and $a_0 > 0$ such that the matrix

$$T_{xy} = |t_{xy}| \exp\{a_0(1 + |x-y|/R)\} \tag{2.9}$$

has row sum at most ρ and maximum at most Cg. The constants are uniform on a fixed small complex β-disc. A simple cycle C through x has a rooted oriented walk of the same length, so its T-weight summed over all $C \ni x$ is at most $\sum_{k\geq 3}(T^k)_{xx} \leq Cg \sum_{k\geq 3} \rho^{k-1}$.

Let $Z_{\mathrm{ev}}(U)$ be the high-temperature partition sum over even edge sets on U, normalized by stripping the $2^{|U|}\prod_{xy}\cosh(\beta Q_{xy})$ factor. Its polymers are nonempty connected even edge sets, compatible when their vertex sets are disjoint. Each such set decomposes into edge-disjoint simple cycles. Fix one decomposition and a spanning tree of the intersecting cycles, rooted at a cycle through a specified anchor x. Attach each child at one shared site and order the children at each site. The union of edge sets is determined by these data. Thus the sum over ordered rooted cycle trees majorizes the anchored sum. If S bounds the total descendant-tree weight rooted at any site, ordering the children costs at most $(1-S)^{-k}$ at a cycle of length k. For $S \leq Mg$, increasing T by this factor per edge still leaves row sum less than one for small g. Choosing M large and inducting in tree height gives $S \leq Mg$. The same argument with part of a_0 reserved proves an anchored $O(g)$ bound with spare factors exponential in the number of edges and in their total length divided by R.

For any $v \in U$ let $r_U(v) = Z_{\mathrm{ev}}(U)/Z_{\mathrm{ev}}(U \setminus \{v\})$. Separate the unique polymer containing v, if present. Then

$$r_U(v) = 1 + \sum_{\substack{F \text{ connected, even} \\ v\in V(F)\subset U}} t(F)\frac{Z_{\mathrm{ev}}(U \setminus V(F))}{Z_{\mathrm{ev}}(U \setminus \{v\})}, \qquad t(F) = \prod_{e\in F} t_e. \tag{2.10}$$

The quotient on the right is a product of inverse deletion ratios on smaller induced sets. For ratios within Cg of one it costs $\exp(C'g|V(F)|)$, absorbed by the reserve. Induction in the set size therefore gives nonvanishing and $r_U(v) = 1 + O(g)$. Comparing the same recurrence on U and $\widetilde{U}$, the missing polymers reach outside Λ, so their span weight supplies $\exp(-a\delta(v)/R)$. For common polymers replace the inverse deletion factors one at a time. A site w in the polymer costs $\exp(a|w-v|/R)$ to move the distance weight from w to v. The spare length weight pays this cost, and the spare size weight pays $|V(F)|-1$ choices of the changed factor. The resulting anchored bound is Cg times the induction bound on the weighted difference of deletion ratios. Absorb this factor for small g to obtain

$$|r_{\widetilde{U}}(v) - r_U(v)| \leq Cge^{-a\delta(v)/R}. \tag{2.11}$$

Removing the same common vertices from both domains preserves the comparison hypothesis, so this induction applies to all products in the recurrence.

For $x \neq y$, the spin two-point formula sums the connected edge sets F with odd set $\{x, y\}$, times $Z_{\mathrm{ev}}(U \setminus V(F))/Z_{\mathrm{ev}}(U)$. The two odd vertices must belong to the same component. Fix a self-avoiding x-to-y path in F and decompose the remaining even set into cycles. For a path of length k, the product of deletion ratios differs from one by at most $Cg(k+1)e^{Cg(k+1)}$. The contribution with at least one cycle has the same bound times the path weight. Indeed, the ordered cycle forests attached to the $k+1$ path vertices sum to at most $(1-Mg)^{-(k+1)} - 1$, with the deletion product absorbed by a small part of the size weight. Hence subtracting the self-avoiding path sum leaves a weighted row error $O(g)$. For domain comparison, any lost path with a nonzero deletion correction or any lost cycle-bearing graph has total length at least $\delta(x)$ and still has its g factor. For common graphs telescope the deletion products and use the weighted ratio estimate just proved, moving $\delta(w)$ to $\delta(x)$ along the connected graph. This costs its reserved length weight.

To replace the self-avoiding path sum by the walk resolvent, a repeated site at times $i < j$ contributes at most $Cg\rho^{k-1}$ at length k after the weighted row sum. This follows by bounding the return between i and j by $\sup_z (T^{j-i})_{zz} \leq Cg\rho^{j-i-1}$ and summing the two remaining path sections. The sum over k and i, j is $O(g)$. Missing repeated walks have spare length at least $\delta(x)$. Finally,

the expansion of $(I-t)^{-1}-(I-\beta Q)^{-1}$ has one marked edge $t_{xy}-\beta Q_{xy}=O(g^2Q_{xy})$ on each walk, so its comparison obeys the same (in fact stronger) bounds. These arguments include $x=y$ with the trivial path, since $S_{\beta,U}(x,x)=1$. Choose the final a smaller than the length reserve both for $|x-y|$ and for $\delta(x)$. Exhaustion and Cauchy's formula give the last assertion. □

To identify the Hessian in the initial field, we apply the Gaussian reconstruction of [32, Proposition 2.4 and Section 17] to the compressed interaction. Let $P_{\beta,\Lambda}(v)=\log\langle e^{v^{\mathsf T}\sigma}\rangle_{\beta Q_\Lambda}$, and write $h=h_{\mathrm{phys}}$. Completing the square in the joint weight

$$\exp\Big\{-\tfrac12 h^{\mathsf T}A_\Lambda^{-1}h+\sqrt{\beta}\,h^{\mathsf T}\sigma+\tfrac{\beta}{2}\sigma^{\mathsf T}Q_\Lambda\sigma\Big\}\,dh \tag{2.12}$$

gives the spin interaction $\beta(A_\Lambda+Q_\Lambda)$ and the conditional law $h\mid\sigma\sim N(\sqrt{\beta}A_\Lambda\sigma,A_\Lambda)$. Relative to the centered Gaussian of covariance V_c, the negative log density of h, up to an additive constant, is

$$\tfrac{c}{2}h^{\mathsf T}B_\Lambda^{-1}h-P_{\beta,\Lambda}(\sqrt{\beta}\,h). \tag{2.13}$$

In the pressure coordinate $u=B_\Lambda^{-1}h$, its Hessian at zero is $cB_\Lambda-\beta B_\Lambda S_{\beta,\Lambda}B_\Lambda$.

Write $T_{\beta,\Lambda}=cB_\Lambda-\beta B_\Lambda S_{\beta,\Lambda}B_\Lambda$ for this Hessian, $T_\beta^\infty(z)$ for its bulk kernel, and $B_\beta=I-\beta Q$. Separate the bulk row mass from its zero-row remainder by setting $r(\beta)=\sum_z T_\beta^\infty(z)$ and $Z_\beta(z)=T_\beta^\infty(z)-r(\beta)\mathbf{1}_{\{z=0\}}$.

Lemma 2.3 (Bulk mass and localization of the initial Hessian)**.** *Let $g=R^{-d}$ and $|c-1|+|\beta-1|\le C_0 g$ with C_0 fixed. For some $a>0$,*

$$\begin{aligned}\sum_z e^{a|z|/R}|T_\beta^\infty(z)|\le Cg,\quad \sum_z|z|^2|Z_\beta(z)|\le CgR^2,\\ \sup_{x\in\Lambda}\sum_{y\in\Lambda}e^{a(\delta(x)+|x-y|)/R}|T_{\beta,\Lambda}(x,y)-T_\beta^\infty(x-y)|\le Cg.\end{aligned} \tag{2.14}$$

At fixed c, the bulk and cube Hessians satisfy $\partial_\beta T_\beta=-I+O(g)$ in the exponentially weighted row norm. The derivative of the finite-volume difference obeys the same boundary estimate.

Proof. On either domain the identity

$$T_\beta=(c-1)B-(\beta-1)I-(\beta-1)^2QB_\beta^{-1}-\beta BE_\beta B \tag{2.15}$$

follows by multiplying $\beta BB_\beta^{-1}B=B+(\beta-1)I+(\beta-1)^2QB_\beta^{-1}$ by B_β. Apply Lemma 2.2 to the last term and use exponentially weighted convolution. The walk expansion of QB_β^{-1} converges in the weighted row norm. A walk present only in the bulk has length at least $\delta(x)$. The same argument applies to the frozen primary filter B_Λ^{-1} on both sides of the Hessian. The second moment follows by using a smaller exponential weight. Differentiation at fixed c gives $\partial_\beta T_\beta=-I+O(g)$ in the weighted row norm on the stated window. For the pressure error use Cauchy's formula on the fixed disc of Lemma 2.2. The same differentiation and walk comparison give the boundary weight for the derivative of the finite-volume difference. □

Corollary 2.4 (Thermal crossing of the bare mass)**.** *Let $p=\sum_z B(z)>0$ and $h_0=\sum_z B^{-1}(z)=p^{-1}$. Let c be fixed with $|c-1|\le C_0 g$. Given $M<\infty$, take A sufficiently large in terms of M, C_0 and the split. For all sufficiently large R, the bare mass $h_0^2 r(\beta)$ crosses every continuous real target of absolute value at most Mg on $|\beta-1|\le Ag$.*

Proof. With $q = \sum_z Q(z) = 1 - p$, the bulk row sum is

$$r(\beta) = cp - \beta p^2\{(1-\beta q)^{-1} + e_\beta\}, \qquad |e_\beta| + |e'_\beta| \le Cg \tag{2.16}$$

by Lemma 2.2. At $\beta = 1$ this gives $r(1) = O(g)$ with a constant independent of A. Lemma 2.3 gives $r'(\beta) = -1 + O_A(g)$ throughout the interval. Thus the two endpoint values of the primary-coordinate mass $h_0^2 r$ bracket the target values for A sufficiently large, and the conclusion follows by continuity. The prescribed c satisfies $|c-1| = O(g)$ for all sufficiently large n at fixed R, as required for the shooting argument in Section 6. □

2.2 All-field pressure damping and the singleton seed

The polymer expansion requires control of the pressure at every real background field. The all-field forest expansion for the periodic short-range reference in [32, Proposition 5.1] motivates the construction here. Deleted bonds and the frozen filter require estimates uniform over the induced graphs. For one cell we compare the Hessian-subtracted pressure with the onsite quartic and keep the damping in both weights. For a connected set of cells we also have to differentiate in the bonds between cells. Lemma 2.5 gives the one-cell estimate on every induced graph, uniformly in the real interpolated couplings.

Lemma 2.5 (Short-reference pressure on an induced graph)**.** *Let U be a finite set of sites, J a real symmetric matrix with zero diagonal and absolute row sum at most $q < 1$, and $\sup_{x,y} |J_{xy}| \le Cg$. Write $B = I - J$,*

$$P(h) = \log \frac{\sum_\sigma e^{\sigma^{\mathsf T} J\sigma/2 + h^{\mathsf T}\sigma}}{\sum_\sigma e^{\sigma^{\mathsf T} J\sigma/2}}, \qquad S = D^2 P(0). \tag{2.17}$$

If a real matrix L satisfies $\|L - B\|_{1\to 1} + \|L - B\|_{\infty\to\infty} \le Cg$, set

$$\mathcal{U}(u) = \tfrac{1}{2}(Lu)^{\mathsf T} S(Lu) - P(Lu), \qquad V_4(u) = \tfrac{1}{12}\sum_x u_x^4, \qquad \Psi(t) = \frac{t^4}{1+t^2}. \tag{2.18}$$

For g sufficiently small and real u,

$$\mathcal{U}(u) \ge c\sum_x \Psi(u_x), \qquad |\mathcal{U}(u) - V_4(u)| \le C\sum_x (|u_x|^6 + g|u_x|^4). \tag{2.19}$$

For each fixed k and fixed field-radius multiple, if all $\|v_i\|_\infty \le 1$, then at sufficiently small $r > 0$,

$$\begin{aligned} r^k |D^k(e^{-\mathcal{U}} - e^{-V_4})(u)[v_1,\dots,v_k]| \le C_k e^{-c\sum_x \Psi(u_x) + C_k|U| r^4} \\ \times \sum_x \big((|u_x| + r)^6 + g(|u_x| + r)^4\big). \end{aligned} \tag{2.20}$$

The bound for $e^{-\mathcal{U}}$ holds with the last sum replaced by 1. The constants are independent of $|U|$ and of all coordinate values of the real background u.

Proof. We first compare the exact magnetization with its mean-field value, including derivatives in the field. The mean-field variational formula will then give the damping and identify the quartic coefficient. To control the derivatives, we use a Taylor norm. For a power series in field increments $\boldsymbol{\xi}$, use the norm $\|F\|_a = \sum_{\boldsymbol{j}} |[\boldsymbol{\xi}^{\boldsymbol{j}}]F| a^{|\boldsymbol{j}|}$, with $a > 0$ fixed and small. This is a submultiplicative norm and each coordinate increment has norm a. At a real center z_0, the Taylor coefficients of $\tanh(z_0 + \cdot)$ and of its fixed derivatives have uniform Cauchy bounds in a fixed disc, independent of

z_0. In particular $\|(\tanh)'(z)\|_a \le 1 + C_k a$ if $\|z - z(0)\|_a \le ka$. The exact cavity message for an edge $b \in \mathbb{R}$ is

$$f_b(z) = \tfrac{1}{2}\big(\log\cosh(z+b) - \log\cosh(z-b)\big). \tag{2.21}$$

Integrating the derivatives of tanh along the real segment from $-b$ to b bounds its derivative by $|b|(1 + C_k a)$ in this algebra. The same segment formula, with its centered second-order remainder, gives $\|f_b(z) - b\tanh z\|_a \le C|b|^3$.

For a site i, the exact field giving its magnetization as $\tanh z_i$ is obtained by removing i from the graph and taking the ratio of partition functions at the fields shifted by $\pm J_{i,\cdot}$. Telescope the ratio by shifting one neighbor field at a time. Each term is $f_{J_{ij}}(z_j^{(i,j)})$, where the field in the smaller graph has shifts by $\pm J_{ik}$ at other neighbors k. Iteration deletes one ancestor at each depth. In particular the direct field shift and the sum of omitted couplings in a row are both bounded by Cgl at depth l. The mean-field field vector is the unique solution $z^{\mathrm{mf}} = h + J\tanh z^{\mathrm{mf}}$ on the real axis. For the increment series, choose a fixed k with $1 + qk < k$ and then a so small that $1 + q(1 + C_k a)k \le k$. Center the cavity fields and the mean-field fields at their respective real values. The exact recursion and the contraction defining the mean-field solution give the same bound ka for every increment series. The cavity ratio is used first at real fields; its analytic formula gives the germ on the field polydisc. The increment majorant gives a convergent extension on a disc of uniform radius. In particular it gives that extension for the gradient of P at every real center.

Compare an exact cavity field with z^{mf} at the same root. The segment joining their series has real center and increment norm at most ka, so the preceding Lipschitz bound applies even when the centers differ. With $q' = q(1 + C_k a) < 1$, the maximum discrepancy at depth l is at most $Cg(l+1)$ plus q' times the corresponding maximum at depth $l+1$. In this estimate omitted couplings and ancestor shifts cost Cgl, and replacing each f_b by $b\tanh$ costs $\sum_j C|J_{ij}|^3 = O(g^2)$. Backward summation gives the root discrepancy $O(g)$. Applying tanh proves a Taylor-algebra bound $O(g)$ on the difference of each magnetization component, uniformly around every real background.

Define the mean-field pressure and its entropy function by

$$P_{\mathrm{mf}}(h) = \sup_{m\in[-1,1]^U} (h^{\mathsf{T}}m + \tfrac{1}{2}m^{\mathsf{T}}Jm - \textstyle\sum_x I(m_x)), \qquad I(t) = \tfrac{1}{2}\big((1+t)\log(1+t) + (1-t)\log(1-t)\big). \tag{2.22}$$

Its maximizing vector is $m = \tanh z^{\mathrm{mf}}$, since the entropy Hessian on the box is bounded below by the identity and the maximization is strictly concave. In particular $D^2 P_{\mathrm{mf}}(0) = B^{-1}$. The Taylor-algebra bound just proved gives, for every fixed $l \ge 1$, the estimates

$$\sup_{x_1} \sum_{x_2,\dots,x_l} |\partial_{x_1}\cdots\partial_{x_l}(P - P_{\mathrm{mf}})(h)| \le C_l g. \tag{2.23}$$

Indeed ordered index tuples at fixed degree multiply the absolute sum of power-series coefficients by $(l-1)!$. For P alone the same estimate holds with C_l instead of $C_l g$. The bounds hold on a smaller complex field strip by translating the Taylor series and using its spare radius. They also hold after composing with L, whose absolute row and column sums are bounded. The pressure itself extends by integrating its gradient on the strip, with the real normalization above. Its mixed derivatives are symmetric, so any of their indices can serve as the root.

Write $\Delta = (P - P_{\mathrm{mf}}) \circ L$ and $\Delta_{\ge 4}(u) = \Delta(u) - D^2\Delta(0)[u,u]/2$. By spin flip Δ is even, and $\Delta(0) = 0$. For u^s equal to u on $|u_x| \le 1$ and zero elsewhere, Taylor's formula and the fourth rooted bound give $|\Delta_{\ge 4}(u^s)| \le Cg\sum_x |u_x^s|^4$. Here and below, bounds on a symmetric derivative tensor are converted to estimates of sums of powers by the arithmetic–geometric mean inequality in the indices. To compare u with u^s, the gradient of Δ is bounded by Cg in supremum norm, and $D^2\Delta(0)$ has absolute row sum at most Cg. The change of Δ is at most $Cg\sum_{|u_x|>1}|u_x|$. For the change of

the quadratic use the same bound on the cross term, since $\|u^s\|_\infty \le 1$, and $Cg\sum_{|u_x|>1}|u_x|^2$ for the pure large term. Therefore

$$|\Delta_{\ge 4}(u)| \le Cg\sum_x \Psi(u_x). \tag{2.24}$$

For $v = B^{-1}Lu$, the exact variational identity gives

$$\tfrac{1}{2}v^{\mathsf T}Bv - P_{\mathrm{mf}}(Bv) = \inf_{m\in[-1,1]^U}\left\{\tfrac{1}{2}(v-m)^{\mathsf T}B(v-m) + \sum_x [I(m_x) - m_x^2/2]\right\}. \tag{2.25}$$

The right side is bounded below by $c\sum_x \Psi(v_x)$: $B \ge (1-q)I$, and $I(m) - m^2/2 \ge cm^4$ for $|m| \le 1$. If $|v_x| \le 2$, separate $|m_x| \ge |v_x|/2$ from its complement; if $|v_x| > 2$, use $|m_x| \le 1$. The convex function Ψ has $\Psi(at) \le \max\{a^2, a^4\}\Psi(t)$ for $a > 0$. Jensen's inequality shows that a matrix with bounded absolute row and column sums is bounded for the modular $\sum_x \Psi(v_x)$, up to fixed constants. Since $u \mapsto v$ and its inverse have that property by the row gap and the bound on $L - B$, this proves the real damping after using $\Delta_{\ge 4} = O(g\sum \Psi)$.

The mean-field quartic coefficient at $v = 0$ is $\sum_x v_x^4/12$. In fact $m = v + O(\|v\|_\infty^3)$ by the stationarity equation and the invertibility of B, so this is the quartic coefficient of the variational identity. Since $B^{-1}L = I + O(g)$ in absolute row and column norms, its quartic tensor in u differs from that of V_4 by $O(g)$ in rooted norm. The fourth rooted bound on Δ adds $O(g)$. The analytic bounds for P give bounded sixth rooted derivatives of $\mathcal{U}$ on a fixed strip. Together with evenness and Hessian subtraction, Taylor's formula along the segment from zero to ζ gives

$$|\mathcal{U}(\zeta) - V_4(\zeta)| \le C\sum_x (|\zeta_x|^6 + g|\zeta_x|^4) \tag{2.26}$$

for all ζ on a smaller strip.

For the derivative estimate we also need complex damping. For real u, w with $\|w\|_\infty$ small, the linear imaginary shift contributes zero to the real part. To bound its second-order remainder use $D^2\mathcal{U}(0) = D^3\mathcal{U}(0) = 0$ and the integral of the fourth derivative along the segment from zero to $u + itw$. By the fourth rooted bound, after separating $|u_x| \le 1$ from the other sites, one has for any $\varepsilon > 0$ and $0 \le t \le 1$,

$$|D^2\mathcal{U}(u+itw)[w,w]| \le C\|w\|_\infty^2 \sum_{|u_x|>1}|u_x|^2 + C\varepsilon \sum_{|u_x|\le 1}|u_x|^4 + C_\varepsilon \sum_x |w_x|^4. \tag{2.27}$$

Indeed the tensor contracts with two w factors and two $u + itw$ factors. In a term containing a large u_x bound both w factors by their supremum and use the quadratic inequality on the large factors while bounding the small factors by a constant. For terms with small factors only, use Young's inequality, giving an arbitrarily small coefficient to the fourth powers of the small u factors. Sum the resulting one-index powers using the rooted bounds. Take ε and then $\|w\|_\infty$ small to absorb the first two terms into real damping. This proves

$$\mathrm{Re}\,\mathcal{U}(u+iw) \ge c\sum_x \Psi(u_x) - C\sum_x |w_x|^4. \tag{2.28}$$

The same lower bound holds for V_4 by direct expansion and Young's inequality. A real displacement bounded by $C_k r$ is handled by the convexity and scaling inequality for Ψ, at an additional cost $C_k|U|r^4$ in the exponent. On the polydisc of radii a fixed multiple of r in the k prescribed field directions, the difference of the two exponential weights is bounded using its exact interpolation integral and the complex bound on $|\mathcal{U} - V_4|$. Cauchy's formula on this finite polydisc gives the stated directional jet bounds. □

At the two initial radii the factor $\exp(C_k|U|r^4)$ is bounded on a single cell of volume of order b_0^d. The last sum in the singleton estimate is at most

$$C|U|(r^6 + gr^4)(1 + \|u\|_\infty/r)^6. \tag{2.29}$$

Thus the lemma supplies the $b_0^d(r^6+gr^4)$ singleton factor with stable damping and a fixed polynomial loss. It also applies to a fixed real row-filter substitution, with the field-direction radii adjusted by its absolute row bound. For several cells, the derivatives in the forest formula connect the cells of an atom. The next proposition retains the same small factor after summing over its spanning trees.

Proposition 2.6 (The multicell pressure forest)**.** *Let J^0 be the real symmetric, zero-diagonal matrix of microscopic bonds, with range at most C_0R and $\sup_{i,j}|J^0_{ij}| \le Cg$. Assume the absolute row gap of Lemma 2.5 for all real forest interpolants, including J^0. Let the cells have side comparable to $b_0 = M_R R$, with M_R a sufficiently large multiple of $\log R$, and use the frozen thermal chart $L_J = \sqrt{\beta}(I - J/\beta)$, $|\beta - 1| \le Cg$. A fixed diagonal correction to L_J of size $O(g)$ is allowed. This accommodates removing a constant short diagonal from the spin Hamiltonian while keeping the original field filter. For either initial radius r suppose $b_0^d r^4 \le C$ and $\nu_r = b_0^d(r^6 + gr^4) = o(1)$ with $\nu_r \ge R^{-C_0}$. For each atom C of at least two cells, its sum of BKAR tree derivatives of $e^{-\mathcal{U}_J(u)}$ satisfies, through any fixed number of field jets tested on directions of supremum norm at most r, the bound*

$$C_{A,\epsilon}\nu_r \exp\left[-A|C| - c\sum_{x\in\cup C}\Psi(u_x) + \epsilon\sum_{B\in C}\frac{\sup_{x\in B}|u_x|^2}{r^2}\right]. \tag{2.30}$$

Here $A < \infty$ and $\epsilon > 0$ are prescribed, $c > 0$ is fixed, and R is sufficiently large in terms of these constants. The bound is uniform in the real interpolation parameters.

Proof. Fix a tree T on the cells. We first estimate the field and bond coefficients of the pressure. Their spatial decay will pay for the sum over the trees. Write $U = \cup C$. For each microscopic bond $\{i, j\}$ over a cell edge $e \in T$, introduce an independent variable z_{ij}, so the interpolated coupling varies by $J^0_{ij}z_{ij}$. Identifying all such variables above e with θ_e at the end gives the desired cell derivative. For a small fixed $\eta > 0$, use the field envelopes

$$W_B = 1 + \sup_{y\in U} e^{-\eta\,\mathrm{dist}(y,B)/R}|u_y|/r, \qquad W_e = \max(W_B, W_D) \quad (e = \{B, D\}). \tag{2.31}$$

The analytic radius of a microscopic variable above e will be a/W_e for a small fixed $a > 0$.

To estimate the joint coefficients at a root x, use the absolute field and bond Taylor norm with field radii a and the above bond radii, multiplying each monomial by $e^{\mu\mathcal{L}/R}$. Here $\mathcal{L}$ is the minimum length of a connecting network for the root, the field indices, and both endpoints of each microscopic bond in the monomial. The norm is submultiplicative at a fixed root by taking the union of two networks. Moving the root from i to j costs at most $e^{\mu|i-j|/R}$. We take $\mu > 0$ small depending on the row gap and then η small enough that $6\eta < \mu/2$.

In this norm the cavity recursion in Lemma 2.5 gives uniform bounds on each gradient component $\partial_{u_x}(P_J \circ L_J)$, and a bound Cg on its difference from $\partial_{u_x}(P_{\mathrm{mf},J} \circ L_J)$. For clarity, the expansion can be centered at any tu, $0 \le t \le 1$, and at the real interpolated bond values. The change in $L_J tu$ per row has norm Car, since $|u_j| \le rW_e$ at the endpoints of the varied bond. Independent field increments cost Ca. The nonconstant part of the sum of ancestor shifts also costs Ca, by the row sum bound and bounded length of each shifted bond. Thus the increments of all cavity messages and of the mean-field solution are bounded by Ka as in the previous proof. At a real forest interpolant the bond is $s_{ij}J^0_{ij}$, with $0 \le s_{ij} \le 1$. Its increment $J^0_{ij}z_{ij}$ has radius at most $a|J^0_{ij}|$

since $W_e \geq 1$. The integral formula $f_b(z) = \frac{b}{2}\int_{-1}^{1} \tanh(z+tb)\,dt$ and the field derivative bound in Lemma 2.5 give the rooted Lipschitz majorant

$$|J_{ij}^0|(1 + C_K a)e^{\mu|i-j|/R}. \tag{2.32}$$

For $q = \sup_i \sum_j |J_{ij}^0| < 1$ its row sum is at most $q(1 + C_K a)e^{C_0\mu} < 1$ when μ, a are sufficiently small. Expanding the integrand through first order in tb, the odd term integrates to zero. The bounded second derivative of tanh on the smaller field disc gives a weighted error $C|J_{ij}^0|^3$ in replacing $f_b(z)$ by $b\tanh z$. Its row sum is $O(g^2)$, since $\sum_j |J_{ij}^0|^3 \leq q \sup_{k,l} |J_{kl}^0|^2$. Comparing the depleted message to the full mean-field message, the omitted ancestors and their field shifts still have weighted norm at most $Cg(l+1)$ at depth l; each shift uses an actual bond of bounded length. Summation of the same geometric recursion gives the Cg error. Composition by L_J and the derivative chain rule cost a bounded weighted row sum. These arguments prove the asserted joint field/bond coefficient bounds. Fixed field derivatives use only a fixed part of the field radius, so they obey the same bounds with constants depending on their orders. The bounds also apply at $t(u+\xi)$ if $\|\xi\|_\infty \leq C_k r$ is a complex shift: translation of the field series by this constant shift uses a spare part of a and only decreases the network length when an index is removed.

Set $\mathcal{R}_J = \mathcal{U}_J - V_4$, and decompose it as $\sum_x \zeta_x \int_0^1 \partial_x \mathcal{R}_J(t\zeta)\,dt$, where $\zeta = u + \xi$. In the integral use Taylor expansion of $\partial_x \mathcal{R}_J$ through degree four at zero, with the fifth derivative integral remainder. Evenness and exact Hessian subtraction hold as identities in the bond variables. The cubic part has weighted coefficient bound Cg: for the exact-minus-mean-field part this follows from the joint estimate above. For the mean-field part its quartic before subtracting V_4 is $\sum_x ((I-J)^{-1}L_J u)_x^4/12$. The Neumann series gives $(I-J)^{-1}L_J - I = O(g)$ in weighted row and column sums of absolute coefficients. Indeed $L_J - (I-J)$ is a linear combination of I, J with coefficients $O(g)$ and an optional fixed $O(g)$ diagonal, and each additional step of the Neumann series has weighted row sum less than one. Replacing one of the four corresponding matrix factors at a time gives the claimed cubic gradient bound. The fifth gradient derivative has bounded weighted coefficients by the joint estimate for the exact pressure. Consequently a term in the representation of $\mathcal{R}_J$ has either four field insertions and coefficient bound Cg or six field insertions and coefficient bound C, summed over the rooted tensors. All inserted field indices, including x, occur in the weighted network.

We next sum square-free coupling coefficients. For $\varnothing \neq F \subset T$, a microscopic monomial contributing to $\theta_F = \prod_{e\in F} \theta_e$ uses one bond above each edge in F. Choose any edge $e_0 = \{B, D\} \in F$. Every inserted field site y has distance at most $\mathcal{L}$ from B. Thus $|\zeta_y| \leq C_k r W_{e_0} e^{\eta\mathcal{L}/R}$. For a square-free coefficient, take the reduced radius $\rho_e = aW_e^{-7}$. Relative to the analytic radii this pays all four or six factors W_{e_0}, leaving at least half the network weight available for a bound on $|F|$.

A network of length $\mathcal{L}$ intersects at most $C_d(1 + \mathcal{L}/b_0)$ cells containing its marked sites. In fact, a traversal of a spanning tree of the network, subdivided into segments of length at most b_0, covers the marked sites by $O(1 + \mathcal{L}/b_0)$ balls of radius b_0, each meeting a bounded number of cells. Since F is a forest, its edges meet at least $|F| + 1$ distinct cells. There are fixed $m_0, c_0 > 0$ such that $\mathcal{L} \geq c_0 b_0 |F|$ when $|F| \geq m_0$.

Let $t_r = \nu_r^{1/(2m_0)}$. The last half of the network weight pays $t_r^{-|F|}$ for $|F| \geq m_0$ on taking M_R large, by the assumed lower power bound on ν_r. For $|F| < m_0$ this cost is at most $\nu_r^{-1/2}$. Summing the weighted rooted coefficient bounds over $x \in U$ therefore gives

$$\sum_{\varnothing \neq F \subset T} \left(\prod_{e\in F} \rho_e\right) t_r^{-|F|} |[\theta_F]\mathcal{R}_J(\zeta)| \leq C|C|\sqrt{\nu_r}. \tag{2.33}$$

Without $t_r^{-|F|}$ the bound is $C|C|\nu_r$. Both bounds are for absolute sums of coefficients, so they include the sum over different microscopic representatives of each cell edge.

In calculating $[\theta_T]\exp(-\mathcal{U}_J)$ one can work in the algebra $\theta_e^2 = 0$. The quartic V_4 is independent of these variables. The coefficient norm of the nonconstant exponential is bounded by the exponential of the norm of the nonconstant log series minus one. The displayed estimate and the damping of the constant term, which is the pressure at $\theta = 0$, bound the desired coefficient by

$$e^{-c\sum_U \Psi(u_x)+C|C|}\left(\prod_{e\in T}\rho_e^{-1}\right) t_r^{|T|}\big(e^{C|C|\sqrt{\nu_r}}-1\big). \tag{2.34}$$

The complex field shift costs only $C|C|$, since $b_0^d r^4 \le C$. When $|T| \ge m_0$, the first m_0 powers of t_r supply another factor $\sqrt{\nu_r}$. For $|T| < m_0$ use the unweighted estimate instead. The remaining powers of t_r give any prescribed exponential margin by taking R large. This pays both the fixed exponential costs and the number of spanning trees of the bounded-degree cell graph, at most a fixed exponential in $|C|$.

Finally the penalty after removing $\prod \rho_e$ is at most $a^{-|T|}\prod_e W_e^7$. Each cell has bounded degree and

$$W_B \le 1 + \sum_D e^{-\eta\,\mathrm{dist}(B,D)/R}\sup_{y\in D}|u_y|/r. \tag{2.35}$$

The scalar matrix on the right has uniformly bounded row and column sums. Thus $\log(1+t) \le \delta t^2 + C_\delta$ and the ℓ^2 bound on that matrix absorb this penalty in $\epsilon\sum_B \sup_B |u|^2/r^2 + C_\epsilon|C|$. The constant per cell is paid from the exponential size margin. All estimates were uniform for $\|\xi\|_\infty \le C_k r$; Cauchy's formula in the fixed tested field directions proves the jet version. Integration over the BKAR parameters preserves the bound. □

The coefficient ϵ in Proposition 2.6 can be chosen small enough to absorb the quadratic loss into the regulator for a packing of pressure atoms. Indeed the absolute rows of the filters give a bounded ℓ^2 matrix on cell suprema, including their fixed collars. The normalized square of a cell supremum is bounded by the sum of its normalized L^2 square and its quotient square in the large-field recovery estimate. At the small radius one uses the spare growing regulator. Hence a product of disjoint pressure atoms consumes only a fixed spare fraction of those two reserves. Together with Lemma 2.5 this supplies the short-reference pressure input to the packing argument. Proposition 5.1 combines these pressure estimates with the full filter and the quadratic tails.

For thermal matching, fix the width A_β of the window $|\beta - 1| \le A_\beta g$. The part of the quartic defect coming from $L_J - (I - J)$ is $O(A_\beta g)$. At both initial radii $g/r^2 \to 0$, so for sufficiently large R depending on the fixed A_β its contribution is absorbed by the sextic factor $b_0^d r^6$ with a constant independent of A_β. The row gap and constants in the cavity and damping estimates can be fixed in a small temperature neighbourhood before choosing this width.

2.3 Spectral and covariance estimates

We turn to the Gaussian reference on the cube. The principal eigenvalue determines how close the reference can be taken to its pole. After conditioning on the principal coordinate we need estimates on the remaining covariance that retain its decay near each face. For the cube walk of Section 1, write $s = 2\alpha = d/2$, $\varrho = R/n$, and $v = n^d$. The physical observation $\mathcal{O}$ acts on the full auxiliary lattice, while the physical block of $\mathcal{Q}_c$ is restricted to the cube. Below, $a_R = \lambda R^s$. For $d < 4$ use the beta-integral split and for $d = 4$ use the finite-range split of Lemma 2.1, with auxiliary diagonal $\mathcal{D} = I$. Denote the observed bulk auxiliary resolvent by $A_r = \mathcal{L}^*(\mathcal{D} + r)^{-1}\mathcal{L}$ and put $H_{r,\Lambda} = (B_\Lambda/c) + rI - A_{r,\Lambda}$.

2.3.1 Generalized principal eigenvalue

The first two estimates locate the pole and separate it from the next mode. The last two control the principal vector and its overlap with a uniform magnetic field.

Proposition 2.7 (Principal spectral gap and magnetic overlap)**.** *Let $\eta_1 \geq \eta_2$ be the two largest eigenvalues of $E = B_\Lambda^{-1/2} A_\Lambda B_\Lambda^{-1/2}$ and choose the sign of its unit principal eigenvector e so that $B_\Lambda^{-1/2} e > 0$. For n/R sufficiently large,*

$$\begin{aligned} c_{*,\Lambda} - 1 &\asymp \varrho^s, & \eta_2^{-1} - \eta_1^{-1} &\asymp \varrho^s, \\ \|B_\Lambda^{-1/2} e\|_\infty &\lesssim v^{-1/2}, & \sum_{x\in\Lambda} (B_\Lambda^{1/2} e)_x &\asymp v^{1/2}. \end{aligned} \tag{2.36}$$

Moreover, for $1 \leq c < c_{*,\Lambda} = \eta_1^{-1}$, the covariance conditioned on $\ell = B_\Lambda^{-1/2} e$ satisfies $\mathcal{S}_\ell(V_c) \leq C V_1$ and $\|V_1\|_{2\to 2} \leq C\varrho^{-s}$.

Proof. We have $B - A = \zeta(I - D_R)^\alpha$ for a uniformly bounded positive normalization ζ. For the cube walk $1 - \widehat{D}_R(k) \asymp \min\{R^2|k|^2, 1\}$, by the Taylor expansion on $|Rk| \leq \kappa$, the product Dirichlet-sum formula away from zero, and the strict bound $|\widehat{D}_R(k)| < 1$ on the intermediate compact scaled annulus. The latter bound is uniform for large R by the Riemann-sum limit, and for the finitely many smaller $R \geq 2$ by the aperiodicity of the support. For a unit vector supported in the cube, $|\widehat{u}| \leq v^{1/2}$ shows that at least half its Fourier mass lies outside $|k| \leq a/n$ for a small fixed a. The resulting energy lower bound on zero extensions is $c\varrho^s$. For an upper bound on the span of two fixed smooth test bumps at scale n, $\|\tau_z u - u\|_2 \leq C|z|\|u\|_2/n$ for $|z| \leq CR$, so Jensen's inequality for the spectral measure and $0 < \alpha \leq 1$ give $C\varrho^s$. Since B_Λ has uniformly bounded condition number, min–max gives $1 - \eta_i \asymp \varrho^s$ for $i = 1, 2$.

For $\alpha < 1$, the jump kernel also satisfies

$$A(x) \geq cR^s (R + |x|)^{-d-s}. \tag{2.37}$$

To verify this, Fourier inversion gives $D_R^{*m}(x) \geq cR^{-d} m^{-d/2}$ for $m \geq m_0$ and $|x| \leq a_0 R\sqrt{m}$. On $|k| \leq M/(R\sqrt{m})$, take a_0 small enough that the phase has cosine at least $1/2$; the ball of radius $1/(R\sqrt{m})$ gives the stated lower bound. The region up to κ/R outside the first ball of radius $M/(R\sqrt{m})$ costs at most $CR^{-d} m^{-d/2} \int_{|y|>M} e^{-c|y|^2}\, dy$. The remaining integral is at most $CR^{-d} q_\kappa^{m-2}$, using $\int |\widehat{D}_R(k)|^2\, dk = O(R^{-d})$. Sum this lower bound over $m \asymp 1 + (|x|/R)^2$ with binomial weights comparable to $m^{-1-\alpha}$ above the smoothing order.

Set $P = B_\Lambda^{-1} A_\Lambda$. The Neumann series dominates P entrywise by the restriction of the stochastic bulk convolution $P_\infty = B^{-1}A$. In the fractional case P has all entries positive; in the local case it is irreducible because Q includes the nearest-neighbour walk edges and A includes a strictly positive diagonal. Thus its positive eigenvector $u = B_\Lambda^{-1/2} e$ satisfies $\|u\|_2 \asymp 1$. For $m = \lceil \varrho^{-s} \rceil$, Fourier inversion gives

$$\|P_\infty^m(0,\cdot)\|_2^2 \leq CR^{-d}(m^{-d/s} + e^{-am}). \tag{2.38}$$

On $|Rk| \leq \kappa$ the multiplier is bounded by $\exp[-c(R|k|)^s]$. On the complement it is bounded strictly below one, with square integral $O(R^{-d})$ by the fixed smoothing power of the cube walk (or of the smaller cube walk in the local split) and the exponentially small white term. Therefore $\|u\|_\infty \leq C\|P_\infty^m(0,\cdot)\|_2 \leq Cv^{-1/2}$, since $\eta_1^{-m} \leq C$. Positivity gives $\sum u \geq \|u\|_2^2/\|u\|_\infty \gtrsim v^{1/2}$; Cauchy–Schwarz gives the upper bound. The finite-volume rows of B have sum in $[p, 1]$ with $p > 0$, proving the magnetic overlap for Bu. For $\alpha < 1$, $P(x, y) \geq c\varrho^s/v$; hence the stochastic ground-state

transform has rows bounded below by the common mass $c\varrho^s u(y)/\sum u$. Oscillation contraction proves $\eta_1 - \eta_2 \gtrsim \varrho^s$.

For the local case the gap follows from the first two Dirichlet min–max levels of the rescaled walk form. Set $\sigma_R^2 = \sum_z z_1^2 D_R(z) \asymp R^2$ and divide the walk form by $\sigma_R^2/(2n^2)$. Here is a compactness argument valid along every sequence $R/n \to 0$. Unit vectors of bounded rescaled energy satisfy $\|\tau_z u - u\|_2^2 \le C[(|z|/n)^2 + (R/n)^2]$ by splitting Fourier space at $|k| = \kappa/R$. Embed them as piecewise-constant functions with mesh $1/n$ on the zero-extended unit cube. A fractional-cell shift is bounded by the sum of the adjacent integer-shift squared norms, so the translation estimate gives L^2 compactness. On every fixed ball of rescaled Fourier frequencies the normalized symbol converges uniformly to $|\xi|^2$, because the jumps have length $O(R)$. The Fourier liminf gives the Dirichlet energy of the limit, which is in $H_0^1((0,1)^d)$ by zero extension. The denominator satisfies $\langle u, B_\Lambda u\rangle = p\|u\|_2^2 + o(1)$; each shift in the finite-range kernel contributes $O(R/n)$ in norm by the same translation estimate. Conversely, samples of smooth compactly supported functions give the continuum energy by Taylor expansion, uniformly for $R/n \to 0$. Applying these bounds to the two min–max spaces (whose orthogonality passes to the limit by the denominator estimate) gives

$$\frac{1-\eta_i}{\zeta\sigma_R^2/(2pn^2)} \longrightarrow \lambda_i(-\Delta^{\mathrm{Dir}}_{(0,1)^d}), \qquad i = 1, 2. \tag{2.39}$$

These Dirichlet levels are distinct. The resulting gap is uniform when n/R is sufficiently large.

The identity $V_c = B_\Lambda^{1/2}(E^{-1} - cI)^{-1}B_\Lambda^{1/2}$ shows that conditioning on $\ell = B_\Lambda^{-1/2}e$ deletes the principal summand in V_c. Writing $\mathcal{S}_\ell$ for the covariance after conditioning on this coordinate, we obtain $\mathcal{S}_\ell(V_c) \le CV_1$ and $\|V_1\|_{2\to 2} \le C\varrho^{-s}$ for $1 \le c < c_{*,\Lambda}$. □

2.3.2 Finite propagation and terminal conditioning

Let q_i be the positive scalar polynomials of [32, Lemma 2.1], rescaled to a common spectral interval $[0, S]$ for the auxiliary precisions. Put $r_i = L^{-2i}$. Choose the terminal cut $r = L^{-2J} \asymp M_*(R/n)^2$ with M_* a sufficiently large fixed constant. Put

$$C_i = a_R \mathcal{O} q_i(\mathcal{Q}_c)\mathcal{O}^*, \qquad C_< = \sum_{i<J} C_i, \qquad C_T = a_R V_c - C_<. \tag{2.40}$$

Let $X_<, X_T$ be independent centred Gaussian fields with covariances $C_<, C_T$, respectively, and write $a = \ell^{\mathsf{T}}(X_< + X_T)$.

Lemma 2.8 (Terminal Gaussian comparison)**.** *For $1 \le c < c_{*,\Lambda}$, the slices sum to $a_R V_c$ and C_i has range $O_L(RL^i)$. For each fixed integer $m \ge 1$,*

$$\mathcal{S}_\ell(C_T) \le \mathrm{Cov}(X_T \mid a = 0) \le C_m \mathcal{C}_r^{(m)}, \qquad \mathcal{C}_r^{(m)} = a_R r^{2m-1}\mathcal{O}(\mathcal{Q}_c + rI)^{-2m}\mathcal{O}^*. \tag{2.41}$$

The preterminal variance satisfies $\ell^{\mathsf{T}} C_< \ell \asymp \lambda n^s$. The comparison constants are uniform as $c \uparrow c_{*,\Lambda}$.

Proof. Lemma A.1 gives the sum of the slices. Each application of $\mathcal{Q}_c$ propagates at most CR and $\deg q_i \le C_L L^i$, which gives the range bound. The covariance comparison is Corollary A.3, applied to the observed field in the canonical normalization.

By (A.3), $C_< \asymp a_R V_r$, with $V_r = \mathcal{O}(\mathcal{Q}_c + rI)^{-1}\mathcal{O}^*$. Write $Y_r = B_\Lambda/c + rI$. The Schur formulas give

$$V_r = A_{r,\Lambda} + A_{r,\Lambda}H_{r,\Lambda}^{-1}A_{r,\Lambda}, \qquad V_r^{-1} = A_{r,\Lambda}^{-1} - Y_r^{-1}. \tag{2.42}$$

The bulk lower bound $H_r(k) \ge c_1(r + \omega(k))^\alpha$, derived below, survives compression when M_* is sufficiently large. Since A_r is bounded, $\|V_r\| \le Cr^{-\alpha}$. For the reverse estimate, the beta integral split

at $1-u=r$ gives $0 \le A - A_r \le Cr^\alpha A$ as bulk forms and after compression. Indeed $1-u\widehat{D}_R \ge 1-u$, and the integrable density in the difference is bounded by $C(1-u)^{\alpha-1}\min\{1, r/(1-u)\}$ times $\widehat{D}_R^{N_{\mathrm{sm}}}$; this multiplier also occurs with a uniformly positive coefficient in A. The white channel has relative change $O(r)$. In the local model $A_r = A/(1+r)$. Inverting the comparison and using the uniform gap of Y_0 gives

$$0 \le V_r^{-1} - V_c^{-1} \le Cr^\alpha A_\Lambda^{-1} + CrI. \tag{2.43}$$

Put $h = B_\Lambda^{1/2} e$. Then $\ell^{\mathsf{T}} h = 1$, $h^{\mathsf{T}} A_\Lambda^{-1} h = \eta_1^{-1}$, and $h^{\mathsf{T}} V_c^{-1} h = \eta_1^{-1} - c = O(r^\alpha)$. Cauchy–Schwarz for $V_r^{1/2}\ell$ and $V_r^{-1/2}h$ yields $\ell^{\mathsf{T}} V_r \ell \ge cr^{-\alpha}$. Since $a_R r^{-\alpha} \asymp \lambda n^s$, we obtain $\ell^{\mathsf{T}} C_< \ell \asymp \lambda n^s$. □

2.3.3 Boundary diagonal of the higher resolvent

The preceding comparison reduces the boundary estimate for the conditioned Gaussian field to a diagonal bound on the higher resolvent. For a point of depth $\delta(x) = \operatorname{dist}(x, \Lambda^c)$ put $b = Rr^{-1/2}$, $\rho_b = \min\{b, R+\delta\}$, and $w(x) = (\rho_b(x)/b)^{p_0}$.

Proposition 2.9 (Boundary diagonal of the higher resolvent)**.** *Fix $0 < p_0 < \alpha$. For $0 < r \le 1$, $|c-1| \le c_0 r^\alpha$ with c_0 sufficiently small, and all sufficiently large fixed m, we have, uniformly in R, n,*

$$\mathcal{C}_r^{(m)}(x,x) \le C_m \lambda b^{-s} w(x)^2. \tag{2.44}$$

The estimate holds with a constant depending on C if $|c-1| \le Cr^\alpha$ and the finite-volume precision is positive. In the microscopic collar the variance of a fixed k-difference stencil has the additional factor $C_k R^{-2k}$, for sufficiently large fixed smoothing order.

Proof. For $|c-1| \le Cr^\alpha$, take fixed $\kappa \ge 1$ such that $C \le c_0 \kappa^\alpha$. On the positive finite-volume spectrum, $(t+r)^{-2m} \le \kappa^{2m}(t+\kappa r)^{-2m}$, and the scales and weights for r and κr differ by fixed factors.

We first bound the row sum of the physical Schur inverse, then use this bound on each of the two outer resolvent factors. For the bulk Schur multiplier,

$$H_r(k) = \zeta\omega(k)^\alpha + (c^{-1}-1)B(k) + r + A(k) - A_r(k). \tag{2.45}$$

If $\omega \ge r$, the first term absorbs the displacement $|c-1| \le c_0 r^\alpha$. When $\omega \le r \ll 1$, restrict the positive beta integral for $A - A_r$ to $1-u \in [r, 2r]$. There $1 - u\widehat{D}_R \asymp r$, the difference of the two inverse denominators has order r^{-1}, and the density has order r^α. Thus $A - A_r \ge c_1 r^\alpha$ in the fractional case. For $\alpha = 1$ the term r suffices. When r is bounded below the same term absorbs the displacement by reducing c_0. Together with the upper bound on $A - A_r$ above this proves $H_r(k) \asymp (r + \omega(k))^\alpha$.

We shall also use the inverse bound for the full bulk precision. Put $h = r + \omega(k)$ and $D_r(k) = \mathcal{D}(k) + rI$. For the fractional split, $D_r(k) \ge chI$, and the squared norm of $D_r(k)^{-1}\mathcal{L}(k)$ is bounded by

$$C\left(1 + \int_0^1 \frac{t^\alpha}{(h+t)^2}\, dt\right) \le C' h^{\alpha-1}. \tag{2.46}$$

Indeed $t = 1-u$, $1 - u\widehat{D}_R(k) + r \ge c(t+h)$, and the remaining factors in the density and observation are bounded. For the local split the same squared norm is bounded uniformly. Block inversion, using the Schur bound just proved, now gives

$$\|(\mathcal{Q}_c(k) + rI)^{-1}\| \le \|D_r(k)^{-1}\| + H_r(k)^{-1}(1 + \|D_r(k)^{-1}\mathcal{L}(k)\|^2) \le Ch^{-1}. \tag{2.47}$$

The entries of $H_{r,\Lambda}^{-1}$ are nonnegative. Moreover

$$(H_{r,\Lambda}^{-1}\mathbf{1})(x) \le Cr^{-\alpha}w(x). \tag{2.48}$$

To see this, consider each face separately. In the fractional case any fixed $0 < p_0 < \alpha$ is available. Use $\min\{1, ((C_1R + x_1)/b)^{p_0}\}$ for $x_1 > 0$, zero outside, where C_1 is fixed and sufficiently large. We verify the required generator sign. For the one-coordinate projected walk write S_m for its m-step sum and $\sigma_R^2 = \operatorname{Var} S_1 \asymp R^2$. For $y/R \to \infty$, the binomial coefficients $c_m \asymp m^{-1-\alpha}$, with $c_m m^{1+\alpha}$ converging to a positive limit, give

$$\left(\frac{y}{\sigma_R}\right)^{2\alpha} \sum_{m\ge 1} c_m\, \mathbb{E}[1 - (1 + S_m/y)_+^{p_0}] \;\longrightarrow\; c_\alpha \int_0^\infty t^{-1-\alpha}\mathbb{E}[1 - (1 + \sqrt{t}Z)_+^{p_0}]\, dt, \qquad Z \sim N(0,1). \tag{2.49}$$

The convergence is uniform along such sequences of y, R. In fact pair opposite steps in the expectation. The paired test is $O(|S_m/y|^2)$ at small displacement and $O(1+|S_m/y|^{p_0})$ at large displacement. Thus the part $m \le \epsilon(y/\sigma_R)^2$ contributes $O(\epsilon^{1-\alpha})$ after scaling, and the part $m \ge M(y/\sigma_R)^2$ contributes $O(M^{-\alpha+p_0/2})$. In the middle range the bounded standardized third moments give a uniform central limit theorem, uniform integrability follows from the second moments, and the rescaled sum is a Riemann sum. Integrating the Gaussian density in t gives a positive constant times

$$\int_0^\infty (2 - (1+u)^{p_0} - (1-u)_+^{p_0})u^{-1-2\alpha}\, du. \tag{2.50}$$

This integral vanishes at $p_0 = \alpha$: integration by parts leaves a constant times the difference of the integrals of $(1-u)^{\alpha-1}u^{-2\alpha}$ on $(0,1)$ and $(1+u)^{\alpha-1}u^{-2\alpha}$ on $(0,\infty)$. The substitution $v = u/(1-u)$ identifies them; when $\alpha \ge 1/2$, the difference of lower cutoffs contributes $O(\epsilon^{2-2\alpha})$ and vanishes in the paired integral. For $p_0 < \alpha$, strict concavity of the power p_0/α bounds its generator from below by the tangent-line multiple of the generator at exponent α, with strictly positive integral. Taking C_1 large therefore gives $(I - D_R)^\alpha y_+^{p_0} \ge c(R/y)^s y^{p_0}$ for $y \ge C_1R$. The cutoff at height one and killing on $x_1 \le 0$ only increase this bound at points where the candidate is less than one. By the identity

$$H_r = c^{-1}\zeta(I - D_R)^\alpha + (c^{-1} - 1)A + rI + (A - A_r) \tag{2.51}$$

in bulk, the possibly adverse correction is at most $C|c-1|$ times the candidate, using the projected jump p_0-moment bound for A. For $C_1R + x_1 \le b$, this is absorbed by the positive term by choosing c_0 small. At $C_1R + x_1 > b$ the candidate is one. This proves the required bound by cr^α; the harmless shift C_1R changes the claimed weight by a fixed factor.

For $d = 4$, use $R + x_1$ in the barrier and any fixed $0 < p_0 < 1$. There is a uniform positive component of the projected step kernel comparable to the projected D_R away from its diagonal, with projected step range at most R. Near $x_1 = 0$, say $x_1 < c_1R$ with fixed small c_1, a fixed step slab exits and gives the lower bound by a loss of one term. All crossing pairs of opposite jumps have nonnegative contributions, since the ratio of the largest forward barrier to its value at x_1 is at most $2^{p_0} < 2$. At $x_1 \ge c_1R$ and $y = R + x_1 < b$, use a smaller fixed slab of projected increments of size comparable to R that does not exit. The concave second difference of the barrier on this slab is at least $c(p_0)(R/y)^2(y/b)^{p_0}$ in magnitude. Truncation at height one only improves the lower bound. This again yields cr^α with $\alpha = 1$. The bulk mass gives the lower bound when the barrier is one. Minimum over all faces preserves supersolutions by positivity of the jumps. This proves the row estimate. The weight is preserved by convolution with A_r by its uniform moment in $(1 + |z|/R)^{p_0}$.

To reach the full diagonal bound we use a high resolvent power rather than the first derivative of the observed resolvent. Set $v_x = \mathcal{O}^*\delta_x$ and $F_x = (\mathcal{Q}_c + r)^{-1}v_x$. These are nonnegative in the auxiliary

and physical variables. The physical component of F_x is $H_{r,\Lambda}^{-1}A_{r,\Lambda}\delta_x$ and has spatial mass $Cr^{-\alpha}w(x)$. In the auxiliary channel the component is $(\mathcal{D}+r)^{-1}(v_x+\mathcal{L}F_{x,\mathrm{phys}})$. The L^2 norm over the channel variable of its spatial sum is at most $Cr^{-(1+\alpha)/2}w(x)$, using $\int_0^1(1-u)^\alpha(r+1-u)^{-2}\,du\le Cr^{\alpha-1}$. The direct v_x term fits since $1\le r^{-\alpha}w(x)$. For the local split the same bound follows from $\mathcal{D}=I$ with $\alpha=1$. After extending the physical component by zero, we thus have $\sup_k\|\widehat{F}_x(k)\|\le Cr^{-(1+\alpha)/2}w(x)$.

The resolvent powers in the finite domain are bounded entrywise by the bulk powers. This follows by expanding the positive-path semigroup of a sufficiently large constant minus $\mathcal{Q}_c$ and integrating in time with e^{-rt}. Consequently $\langle F_x,(\mathcal{Q}_c+r)^{-(2m-2)}F_x\rangle$ is bounded by the quadratic form with the bulk multiplier. The low-frequency operator norm is $C_m(r+R^2|k|^2)^{-(2m-2)}$ by the Schur spectral bound. Integration there, and the two factors of $\widehat{F}_x$, give $CR^{-d}r^{d/2-2m+2}r^{-1-\alpha}w(x)^2$. On high frequencies use the bounded multiplier and $\|F_x\|^2\le CR^{-d}r^{-1-\alpha}$. The last estimate follows by entrywise bulk comparison on the two resolvent factors. At low frequency use $\langle v(k),(\mathcal{Q}_c(k)+r)^{-2}v(k)\rangle\le Cr^{-1-\alpha}$, by the spectral and Schur bounds, and integrate over a ball of volume $O(R^{-d})$. At high frequency the spectral gap is bounded below and $\|v(k)\|^2$ has integral $O(R^{-d})$ by smoothing. The factor r^{2m-1} absorbs this crude high-frequency estimate for large m, since $w^2\ge r^{p_0}$. The resulting scale is $a_RR^{-d}r^{d/2-\alpha}=\lambda b^{-s}$, as asserted.

Taking the absolute value at each differentiated observation leg gives $C_k\lambda b^{-s}R^{-2k}w(x)^2$ for the variance of a k-difference stencil in the microscopic collar. For the walk part, move each difference onto a separate step in $D_R^{*N/2}$. Here N denotes the fixed smoothing order N_{sm}; in the local split the observation leg has the corresponding convolution of E_R. The absolute marked step has mass $O(R^{-1})$ and range $O(R)$, and the unmarked steps retain an arbitrarily large fixed smoothing power. The absolute white term has mass at most $C_k\sqrt{\varepsilon_R}$ and support in a fixed stencil. After taking R^{-k} out of each leg the spatial sum, moment, and high-frequency integral bounds in the proof are unchanged. □

2.3.4 Intermediate annuli for fractional powers

For fractional powers the field norm also tests differences on regions where $R\ll\delta(x)\lesssim b$. The next estimate supplies the derivatives used in that norm.

Lemma 2.10 (Annular stencils for fractional powers)**.** *Let $d<4$, $0<r\le1$, $|c-1|\le c_0r^\alpha$, and $X\sim N(0,\mathcal{C}_r^{(m)})$. Given a fixed derivative order P and $0<p<\alpha$, choose the smoothing and resolvent orders sufficiently large. For stencils in the cube,*

$$\|\rho_b(x)^{|\gamma|}\nabla^\gamma X(x)\|_{L^2}\le C(\lambda b^{-s})^{1/2}(\rho_b(x)/b)^p,\qquad |\gamma|\le P+1. \tag{2.52}$$

For each scale-b cell B, the corresponding supremum has a uniform Gaussian tail. More precisely, for $t\ge1$,

$$\mathbb{P}\left(\sup_{x\in B;\ |\gamma|\le P}\frac{\rho_b(x)^{|\gamma|}|\nabla^\gamma X(x)|}{(\lambda b^{-s})^{1/2}(\rho_b(x)/b)^p}>t\right)\le Ce^{-ct^2}. \tag{2.53}$$

Proof. We first prove

$$\mathrm{Var}(\nabla^\gamma X(x))\le C_\gamma\lambda b^{-s}\rho_b(x)^{-2|\gamma|} \tag{2.54}$$

for $|\gamma|>s$. We use interior bounds for bulk Green functions and then apply them to the difference between bulk and cube solutions. On a fixed complex disc $|z-r|\le\epsilon r$ write the bulk Green kernel $G_z=H_z^{-1}$. The estimates from the beta integral and the Fourier symbol bound give for $|y|\ge C_kR$

$$|\nabla^kG_z(y)|\le C_kR^{-s}|y|^{s-d-k},\qquad |\nabla^kA_z(y)|\le C_kR^s|y|^{-d-s-k}. \tag{2.55}$$

Indeed the derivatives in $\theta = Rk$ of the low-frequency symbols on a window $T \geq \sqrt{r}$ are bounded by $C_j T^{-s-j}$ and $C_j T^{s-j}$, respectively, for G_z and $A_z - A_z(0)$ (the value subtracted here is at zero Fourier frequency). This follows by differentiating the beta integral in $\widehat{D}_R$. Low-frequency window integrals are $CR^{-d-k}T^{d+k\mp s}(1+T|y|/R)^{-J_1}$ for any fixed J_1, by integration by parts. At high frequencies subtract a finite-range partial Neumann sum of the short kernel from G_z. The remainders and their finitely many derivatives have rapid fixed decay by the cube Dirichlet-sum bounds and arbitrarily large fixed smoothing order. The zero-exclusion term is $O(R^{-d})$, and hence $O((1+|\theta|)^{-1})$ on the scaled torus; the same single inverse power bounds all fixed derivatives by the product Dirichlet-sum formula. Taking a sufficiently large fixed smoothing power supplies the required integrability of the differentiated remainders. Summing the windows gives the two kernel bounds.

Put $d_x = \min\{b, \delta(x)\} \gg R$. For a bounded H_z-harmonic function u on the cube, take a cutoff $\chi = 1$ near x with support in a ball of radius $d_x/2$. Set $h = G_z * \nabla^{\gamma,*}\delta_x$. The identity $\nabla^\gamma u(x) = \langle H_z((1-\chi)h), u\rangle$ uses the bilinear pairing. The tail bounds give

$$\|(1-\chi)h\|_1 \leq CR^{-s}d_x^{s-|\gamma|}, \qquad \|\nabla((1-\chi)h)\|_1 \leq CR^{-s}d_x^{s-|\gamma|-1} \tag{2.56}$$

when $|\gamma| > s$. The kernel tail gives a jump bound $\sum_y |K_z(y)| \min(1, |y|/d_x) \leq C(R/d_x)^s$ for $s < 1$; for $s \geq 1$ use the centered second difference instead. The same tail computation gives $\sum_y |K_z(y)| \min(1, |y|^2/d_x^2) \leq C(R/d_x)^s$ and the cutoff kernel has second difference norm $CR^{-s}d_x^{s-|\gamma|-2}$. The mass contribution is bounded by using $|H_z(0)| \lesssim r^\alpha$, with zero denoting Fourier frequency. Thus $|\nabla^\gamma u(x)| \leq Cd_x^{-|\gamma|}\|u\|_\infty$. The same estimate holds for $A_z * u$ by using this bound on the near part of the convolution and the differentiated A_z tail on the far part.

For a unit source in the auxiliary-plus-physical Hilbert space compare the bulk and cube solutions, extending the physical source by zero. Their physical difference is H_z-harmonic with exterior data from the bulk solution. Its supremum is bounded by the exterior supremum by the substochastic exit kernel for the real part of z. The exterior physical-source row omits the onsite term and has squared L^2 norm $CR^{-d}(1+\log(1/r))$. The auxiliary-source row has squared norm $CR^{-d}r^{\alpha-1}$, by integrating the low-frequency majorant $(r+R^2|k|^2)^{-1-\alpha}$. On the high frequencies both row bounds follow from the smoothing estimates. Specifically the physical-source bound uses the Neumann difference $G_z - (c^{-1}+z)^{-1}I$, whose high-frequency square is integrable at order R^{-d} by Plancherel for the short kernel. The auxiliary-source bound uses the squared norm of $\mathcal{L}(k)$ on high frequencies and the auxiliary beta density on low frequencies. Since $\alpha < 1$, the logarithm is dominated by $Cr^{\alpha-1}$. The observed difference of the solutions is $A_z * u$. Cauchy differentiation on the disc for the $(m-1)$-st derivative multiplies its bound by $C_m r^{1-m}$. Squaring the resulting dual norm and multiplying by $a_R r^{2m-1}$ gives $C\lambda b^{-s}d_x^{-2|\gamma|}$. The bulk contribution has the same bound with $d_x = b$ by Fourier integration of the bulk higher resolvent, so the triangle inequality completes the estimate.

To interpolate and take suprema, on a cube of diameter d_x use the Hilbert space of Gaussian random variables, the improved diagonal bound, and stencils through order $k_0 > \max\{P+1, s\}$ (as large as needed). In the tensor-product discrete Taylor formula use mesh $\max\{1, \lfloor cd_x(d_x/b)^{p_0/k_0}\rfloor\}$, with stencils directed into the cube. The fixed-order Vandermonde inverse and Taylor remainder bound give

$$\|\rho_b(x)^{|\gamma|}\nabla^\gamma X(x)\|_{L^2} \leq C(\lambda b^{-s})^{1/2}(\rho_b(x)/b)^{p_1}, \qquad |\gamma| \leq P+1, \quad p_1 = p_0(1-(P+1)/k_0). \tag{2.57}$$

Here we choose $p < p_0 < \alpha$ and then k_0 large enough that $p < p_1$. If the real-valued optimal mesh is less than a fixed constant, use the immediate finite differences of the diagonal estimate. For each derivative and each cube of diameter comparable to the shell depth, the extra derivative gives a

Lipschitz L^2 metric on the normalized stencil field, with scale $(\lambda b^{-s})^{1/2}(\rho_b/b)^{p_1}$. Dyadic spatial grids prove a subgaussian supremum bound with this scale: at level j there are $O(2^{dj})$ increments whose standard deviations are $O(2^{-j})$ times the scale; thresholds proportional to $t2^{-j/2}$ give a summed tail Ce^{-ct^2} for $t \geq 1$. A shell of depth $2^{-i}b$ in a scale-b cell needs at most $C2^{i(d-1)}$ such cubes. For a field weight $p < p_1$, the shell tail after normalization is bounded by $C2^{i(d-1)} \exp[-ct^2 2^{2i(p_1-p)}]$. Its sum gives the exponential field-norm moment uniform in the cube and the scale, including edges and corners. □

2.3.5 The local annular estimate in dimension four

For the local split of Lemma 2.1, the auxiliary diagonal is the identity and every block has range $O(R)$. Consequently the interior equation can be iterated on a ball contained in the cube. This yields the annular derivative estimate needed for the four-dimensional field norm.

Lemma 2.11 (Local resolvent derivatives)**.** *Take $d = 4$, $0 < r \leq 1$, $|c - 1| \leq c_0 r$ for a sufficiently small fixed c_0, $b = Rr^{-1/2}$, and $\rho_b(x) = \min\{b, R + \operatorname{dist}(x, \Lambda^c)\}$. For each fixed derivative multiindex γ and all sufficiently large fixed integers m, the covariance*

$$\mathcal{C}_r^{(m)} = a_R r^{2m-1} \mathcal{O}(\mathcal{Q}_c + rI)^{-2m} \mathcal{O}^* \tag{2.58}$$

satisfies

$$\operatorname{Var}(\nabla^\gamma X(x)) \leq C_{m,\gamma} \lambda b^{-2} \rho_b(x)^{-2|\gamma|}, \qquad X \sim N(0, \mathcal{C}_r^{(m)}). \tag{2.59}$$

Near the boundary the differences are taken on stencils lying in the cube. The short smoothing split can be chosen once for any finite family of derivative orders.

The proof is given in Appendix B.

Proposition 2.9 and Lemma 2.11 give the two inputs to the discrete Taylor interpolation and shell sum in the proof of Lemma 2.10. Thus the same weighted stencil and cell-supremum estimates hold in dimension four for any fixed $0 < p < 1$.

For the merged slice set $M_R = b_0/R \asymp \log R$. Measured at scale b_0, a stencil of order at most P at a constituent scale $b \leq b_0$ gains at most $(b_0/b)^{P+p}$ from the change of weights. A merged cell contains $O((b_0/b)^d)$ cells at that scale. The subgaussian tail on each cell therefore bounds their maximum at an additional factor $C\sqrt{1 + \log(b_0/b)}$. Summing the geometric scales, the squared fluctuation parameter in the $b_0^{-s/2}$ field units is at most $C\lambda M_R^{2P+2p+s}(1 + \log M_R)$, which tends to zero. For the diagonal alone, summing the variances gives $C\lambda R^{-s} = C\lambda M_R^s b_0^{-s}$.

Corollary 2.12 (Profile of the terminal principal coordinate)**.** *At the prescribed enhancement let $\sigma_T^2 = \ell^\mathsf{T} C_T \ell$ and $u_T = C_T \ell/\sigma_T^2$. For $1 \leq d \leq 4$, any fixed $0 < p < \alpha$, and each of the finitely many tested stencils,*

$$|\nabla^\gamma u_T(x)| \leq C v^{-1/2} \rho_n(x)^{-|\gamma|} (\rho_n(x)/n)^p. \tag{2.60}$$

Moreover, as $n \to \infty$ at fixed R, $\|u_T - B_\Lambda^{1/2} e\|_2 = o(1)$, $\mathbf{1}^\mathsf{T} u_T \asymp v^{1/2}$, and $\sum_x u_T(x)^4 \asymp v^{-1}$.

Proof. Put $\sigma^2 = \sigma_T^2 + \sigma_L^2$. By regression, the covariance of X_T conditioned on $a = 0$ contains the positive rank-one term $\sigma_T^2 \sigma_L^2 u_T u_T^\mathsf{T}/\sigma^2$, where $\sigma_L^2 = \ell^\mathsf{T} C_< \ell$. By Lemma 2.8, $\sigma_L^2 \asymp \lambda n^s$ and the conditioned covariance is bounded by $C_m \mathcal{C}_r^{(m)}$. Since $\sigma_T^2/\sigma^2 \to 1$ under enhancement, testing the rank-one term against a stencil and applying the weighted stencil estimate at $b \asymp n$ gives the stated pointwise bound. With $u_\star = B_\Lambda^{1/2} e$ the exact identity

$$u_T - u_\star = \frac{\sigma_L^2 u_\star - C_< \ell}{\sigma_T^2} \tag{2.61}$$

has right side $o(1)$ in ℓ^2 by the regularized inverse bound. Proposition 2.7 and Cauchy–Schwarz now give $\mathbf{1}^{\mathsf{T}} u_T \asymp v^{1/2}$. Finally $\|u_T\|_2 \asymp 1$, so the fourth-moment bounds follow from $v^{-1}\|u_T\|_2^4 \le \sum_x u_T(x)^4 \le \|u_T\|_\infty^2 \|u_T\|_2^2$. □

2.3.6 The integrated uniform source

A uniform magnetic source induces the deterministic field $C_<\mathbf{1}$ after the earlier slices are integrated. We estimate it in the raw weighted test norm $F_{n,p}$ of Section 3, taken at scale n.

Corollary 2.13 (Integrated uniform source)**.** *For an ordinary slice C_j of side length b, for any fixed derivative order and $0 < p < p' < \alpha$ one has*

$$\|C_j\mathbf{1}\|_{F_{n,p}} \le C\lambda n^p b^{s-p}. \tag{2.62}$$

Consequently $\|C_<\mathbf{1}\|_{F_{n,p}} \le C_R n^s$ at fixed R.

Proof. The variance bounds at weight p' and covariance Cauchy–Schwarz give

$$|\nabla^\gamma (C_j\mathbf{1})(x)| \le C\lambda b^{d-s}\rho_b(x)^{-|\gamma|}(\rho_b(x)/b)^{p'}. \tag{2.63}$$

Indeed a row and its fixed stencil have support of volume $O(b^d)$, and the other covariance leg has diagonal variance at most $C\lambda b^{-s}$. For positive derivative order the left side vanishes more than Cb from the boundary by finite propagation and the constant bulk row sum. In the collar $\rho_n(x) \le C\rho_b(x)$. Everywhere $(\rho_b/b)^{p'}/(\rho_n/n)^p \le (n/b)^p$. Together with $d = 2s$ these inequalities give the bound. Choose $p < s$. Summation over the preterminal scales gives $\|C_<\mathbf{1}\|_{F_{n,p}} \le C_R n^s$, with a scale-independent constant at fixed R also for the merged slice. The same bound holds in any test norm controlled by this raw norm. □

2.4 The filtered bulk bubble

The filtered bubble gives the leading decrease of the bulk quartic coupling at each scale. The logarithm in the marginal flow requires a positive contribution of order λ^2 on the ordinary scales outside a fixed number of terminal scales. Write $C = \Gamma_j^{\mathrm{bulk}}$, $W = \sum_{i<j}\Gamma_i^{\mathrm{bulk}}$, $D_j = \overline{T}C\overline{T}^*$ and $w_j^T = \overline{T}W\overline{T}^*$.

Proposition 2.14 (The filtered bulk bubble)**.** *For the prescribed reference and the normalized finite filter $\overline{T}$, the subtraction in the compensated quartic coordinate is $\mathfrak{b}_j g^2$, where*

$$\mathfrak{b}_j = 9\sum_x ((w_j^T + D_j)(x)^2 - w_j^T(x)^2). \tag{2.64}$$

On ordinary scales $0 \le \mathfrak{b}_j \le C\lambda^2$, with $\mathfrak{b}_j \ge c\lambda^2$ outside a fixed number of terminal scales. The merged initial slice satisfies $0 \le \mathfrak{b}_0 \le C\lambda^2 \log(2 + b_0/R)$.

Proof. To verify the coefficient, periodize the kernels on a torus larger than the contraction ranges and test on a constant field. For a symmetric kernel S, write $P_S = \exp(\frac{1}{2}\sum_{x,y} S(x,y)\partial_x\partial_y)$ and $X_S(U) = \frac{1}{2}[P_S((P_{-S}U)^2) - U^2]$. The heat semigroup identity

$$X_{W+C}(P_C U) = P_C X_W(U) + \tfrac{1}{2}[P_C(U^2) - (P_C U)^2] \tag{2.65}$$

expresses the second-order subtraction as the outgoing localized X_{W+C} minus the incoming localized X_W evolved by P_C. Localization preserves the quartic evaluation at a constant field since $\overline{T}\mathbf{1} = \mathbf{1}$.

Write G for the local quartic coefficient before mass compensation. Two cross contractions between filtered quartic vertices give $\frac{1}{2}(G/4)^2\binom{4}{2}^2 2! = 9G^2/4$ times the square of the filtered covariance per ordered pair of sites. With the quartic convention $G/4$, the difference gives the displayed coefficient. For the single cross contraction with the quadratic vertex, write u for its dimensionful row mass, u_+ for the row mass of $P_C U$, and $s_j = \sum_x w_j^T(x)$. The resulting quartic increment is $-4G(u_+ s_{j+1} - u s_j)$, because both the mixed and filtered covariances have the unfiltered row sum. It cancels to second order in the coordinate $g = G(1 + 4u s_j)$ used in Section 3.9. A zero-row quadratic vertex has zero contribution on a constant field. In an interior cube cell all contraction stencils agree with their bulk counterparts. We estimate the coefficient in Fourier space, first from below on a frequency band and then from above by the pointwise slice bound.

For the observation vector $v(k) = \mathcal{O}(k)^* 1$ write $V_\rho(k) = \langle v(k), (\mathcal{Q}_c(k) + \rho)^{-1} v(k)\rangle$. Let $\omega(k) = 1 - \widehat{D}_R(k)$. Near zero, $\omega(k) \asymp R^2|k|^2$, with isotropic quadratic leading term. With a regularization $\rho \geq 0$, the physical Schur complement is

$$S_\rho(k) = B(k)/c + \rho - A_\rho(k), \qquad V_\rho(k) = A_\rho(k) + A_\rho(k)^2/S_\rho(k) \tag{2.66}$$

for scalar bulk filters. In the fractional case A_ρ is obtained by replacing $1 - u\widehat{D}_R$ in the beta integral by $1 - u\widehat{D}_R + \rho$. It is bounded below near zero frequency, uniformly for $0 \leq \rho \leq 1$. In the local split of Lemma 2.1 use $A_\rho = A/(1+\rho)$.

Let r be the polynomial scale. At $\omega(k) \asymp \delta r$ and $|c-1| \leq \eta r^\alpha$, with η sufficiently small depending on δ, $S_0(k)$ is positive. The estimates

$$V_{\delta r}(k) \geq c\delta^{-\alpha} r^{-\alpha}, \qquad V_r(k) \leq C r^{-\alpha} \tag{2.67}$$

hold with constants independent of small δ. To see this for $\alpha < 1$, use $B - A \asymp \omega^\alpha$ and the upper bound $A_0 - A_\rho \leq C\rho^\alpha$, obtained by splitting the beta integral at $1 - u = \rho$. For $\omega \leq \delta r$, integration over $1 - u \asymp r$ gives $A_0 - A_r \geq c r^\alpha$. In dimension four the corresponding bounds follow from $A_0 - A_\rho = A\rho/(1+\rho)$ and $s = 2$. The added positive ρ term in the Schur complement gives the same estimates. This is a calculation on the full-lattice bulk fibers.

Let ν_k be the spectral measure of the observation vector on this band. For the positive polynomial q_i of scale $r = L^{-2i}$, choose the fixed cutoff $c_* > 0$ in $q_i(t) \geq c/r$ for $0 \leq t \leq c_* r$. The part of $V_{\delta r}$ from $t > c_* r$ is at most $(1 + c_*^{-1})V_r$. Choose δ small enough that this is at most half the lower bound on $V_{\delta r}$. Hence $\nu_k([0, c_* r]) \geq c_\delta r^{1-\alpha}$ and

$$\langle v(k), q_i(\mathcal{Q}_c(k)) v(k)\rangle \geq c_\delta r^{-\alpha}. \tag{2.68}$$

The normalized finite filter is bounded below there by its row sum one and second moment $O(R^2)$. The band has measure $\asymp_{\delta,d} R^{-d} r^{d/2}$; multiplication by a_R and Plancherel give

$$\|D_i\|_2^2 \geq c\lambda^2 R^{2s-d} r^{d/2-2\alpha} = c\lambda^2. \tag{2.69}$$

Here D_i denotes an individual filtered polynomial slice in this estimate. The nonnegative cosine coefficients of the scalar construction in Appendix A.1 give positivity also on the negative bulk spectrum. Thus the two contributions to $\mathfrak{b}_j$ are nonnegative by Plancherel, and a merged slice inherits the lower bound of any one component. The estimate $|c-1| = O((R/n)^s)$ puts all but a fixed number of terminal scales in the required displacement range.

For the upper bound the ordinary pointwise estimate is $\|\Gamma_j^{\mathrm{bulk}}\|_\infty \leq C\lambda b_j^{-s}$. Here is its direct bulk proof, which also covers preterminal scales with a fixed enlarged displacement constant $|c-1| \leq C r^\alpha$. Shift the auxiliary precision by $\kappa' r$ to make it nonnegative, using the Schur lower bound at that

regularization, and choose $\kappa > \kappa'$ fixed. The spectral majorant of $q_i(t)$ is $C_m r^{2m-1}(t+\kappa r)^{-2m}$. On $t \geq 0$ this follows from the positive-spectrum polynomial bound. On $-\kappa' r \leq t < 0$, the cosine coefficients, $q_i(0) = O(r^{-1})$ and the degree $O(r^{-1/2})$ bound the hyperbolic cosines by a constant. At $|Rk| \leq c_1$, the Schur spectral lower bound for the shifted precision reduces the observed inverse power to an observed resolvent times $C(r + R^2|k|^2)^{-(2m-1)}$. The observed resolvent is bounded by $C(r+R^2|k|^2)^{-\alpha}$. Integration yields $CR^{-d}r^{d/2-\alpha}$ before the factor a_R. On the remaining frequencies the gap is bounded below and the square integral of the observation vector is $O(R^{-d})$, by the smoothing in the split. The factor r^{2m-1} thus controls this range as well. If κr is bounded below, the same inequalities hold with constants depending on κ and fixed spectral bounds. Positivity of the bulk multiplier controls every entry by its diagonal, and $a_R R^{-d} r^{d/2-\alpha} = \lambda b_j^{-s}$.

Finite range and bounded absolute filter rows now give $\|D_j\|_1 \leq C\lambda b_j^s$ and $\|w_j^T\|_1 \leq C\lambda b_j^s$. It follows that $\mathfrak{b}_j \leq C\lambda^2$. For a merged slice, the same ordinary bound and finite range give the envelope $C\lambda(R+|x|)^{-s}$ on its support of radius $O(b_0)$ before filtering. Squaring and summing over radial layers costs $C\lambda^2 \log(2 + b_0/R)$ because $2s = d$. The filter is bounded in ℓ^2, giving the asserted first-step upper bound. This calculation isolates the bulk covariance input needed for the logarithmic flow. □

2.5 Relative form bound on the boundary Hessian

Subtracting the bulk row mass leaves a zero-row convolution and a correction concentrated near the boundary. We first bound both terms relative to the kinetic energy. The next subsection estimates them in the weighted field norm required by the polymer expansion.

Lemma 2.15 (Form bound on the initial boundary Hessian)**.** *Write the zero-extension kinetic form as*

$$\mathcal{E}_\Lambda(u) = \langle P_\Lambda^* u, L_R^{d/4} P_\Lambda^* u \rangle. \tag{2.70}$$

Suppose D_Λ is symmetric and $\sum_y |D_\Lambda(x,y)| \leq Cge^{-a\delta(x)/R}$. For $d = 1,2,3,4$, uniformly in the cube,

$$|\langle u, D_\Lambda u \rangle| \leq C_a g \mathcal{E}_\Lambda(u). \tag{2.71}$$

The same relative form bound holds for a compressed even zero-row convolution Z with $\sum |Z(z)| \leq Cg$ and $\sum |z|^2 |Z(z)| \leq CgR^2$.

Proof. Symmetry and $2|u_x u_y| \leq |u_x|^2 + |u_y|^2$ reduce the boundary bound to $\sum_x e^{-a\delta(x)/R}|u_x|^2$. For $d < 4$, the zero-extension jump form includes the killing row $r_\Lambda(x)|u_x|^2$. The jump lower bound gives $r_\Lambda(x) \geq c[R/(R+\delta(x))]^s$ by summing an exterior box across a closest face at distance comparable to $R + \delta(x)$. This dominates the required exponential.

For $d = 4$, the Fourier lower bound $1 - \widehat{D}_R(k) \geq c \min\{R^2|k|^2, 1\}$ controls the zero-extension energy of the one-dimensional uniform walk with jumps $\pm 1, \ldots, \pm R$ in any chosen coordinate. Consider the lower face of that coordinate and a fixed step $\ell \in [R/4, R/2]$. If x_i is the normal distance in lattice units, telescope from x to the first point outside the face along the negative ℓ-steps. Cauchy–Schwarz gives

$$|u_x|^2 \leq C(1 + x_i/R) \sum_{\substack{j \geq 0 \\ x_i - j\ell \geq 1}} |u_{x - j\ell e_i} - u_{x-(j+1)\ell e_i}|^2. \tag{2.72}$$

With weight $e^{-ax_i/R}$, summing the right side costs only a fixed multiple of each edge with interior upper endpoint t, whose coefficient is bounded by

$$\sum_{j \geq 0} e^{-a(t_i + j\ell)/R}(1 + (t_i + j\ell)/R). \tag{2.73}$$

This bound is uniform in ℓ in the specified range and in $t_i \geq 1$. Average over those ℓ, whose number is comparable to R, to obtain the normalized axis energy. Repeat for the upper faces. The exponential of the distance to the complement is bounded by the sum of the face exponentials, up to a harmless lattice convention.

For the bulk term the zero row and evenness give

$$|\widehat{Z}(k)| \leq \sum_z |Z(z)|\,|1-\cos(k\cdot z)| \leq Cg\min\{R^2|k|^2,1\} \leq C'g(1-\widehat{D}_R(k))^{d/4}. \tag{2.74}$$

Use Plancherel on the zero extension. □

The affine-completed field norm defined in Section 3 admits constant test directions at the boundary. A shrinking boundary coefficient in that norm therefore requires a separate estimate. Absorbing the perturbation into the Gaussian reference would also require a new principal pole and regulator estimates for the modified slices.

2.6 The initial bulk and boundary quadratic heads

We now split each initial quadratic interaction into a head, whose two endpoints lie within a fixed multiple of the initial cell size, and a tail that enters the polymer activity. Writing the zero-row bulk form as a sum of squared differences gains two powers of the jump length. For rows lost at the physical boundary, the gain comes from the boundary weight in the field norm.

Proposition 2.16 (Initial quadratic heads in the cube chart)**.** *Let $d \in \{1,2,3,4\}$, $b = b_0 = M_R R$, and choose $\max\{0,(s-1)/2\} < p < s/2$, with $p \neq 1/2$. Subtract the bulk row mass from the canonical Hessian of Section 2.1. On a fixed enlarged cube of side comparable to b, the resulting bulk difference-square form and the boundary form have normalized bilinear F_b-norms, at radius ℓ_b, bounded respectively by*

$$C\lambda G_0(M_R^{s-2} + M_R^{s-1-2p}), \qquad C\lambda G_0 M_R^{s-1-2p}. \tag{2.75}$$

The interior bulk estimate is $C\lambda G_0 M_R^{s-2}$. The forms admit exact decompositions into heads of span $O(b)$ and exponentially small tails. Anchors of the boundary head are within $O(b)$ of the physical boundary. The tail row bounds and the corresponding local tensor bounds gain a factor $C\exp(-cb/R)$. All bounds hold on the complex thermal window $|\beta - 1| \leq A_0 g_{\mathrm{mic}}$ at fixed A_0.

Proof. In the primary field, let Z denote the translation-invariant zero-row Hessian after the scalar subtraction and let D denote the finite-volume comparison. Put $\theta = \lambda G_0$. Lemma 2.3 and the exponentially weighted Neumann norm of the primary pullback give

$$\sum_z |Z(z)|e^{a|z|/R} \leq C\theta R^{-s}, \qquad \sup_{x\in\Lambda}\sum_{y\in\Lambda}|D(x,y)|e^{a(\delta(x)+|x-y|)/R} \leq C\theta R^{-s}, \qquad \delta(x) = \mathrm{dist}(x,\Lambda^c). \tag{2.76}$$

Here the canonical conversion multiplies the physical Hessian by $\lambda^{-1}R^{-s}$, and $\lambda^{-1}g_{\mathrm{mic}} \asymp \theta$. All norms in these two inequalities use absolute values, so the argument applies also to complex coefficients. For the compressed bulk form use the exact identity

$$\frac{1}{2}\sum_{x,y\in\Lambda}\phi_x Z(x-y)\phi_y = -\frac{1}{4}\sum_{x,y\in\Lambda}Z(x-y)(\phi_x-\phi_y)^2 + \frac{1}{2}\sum_{x\in\Lambda}m_x\phi_x^2, \qquad m_x = -\sum_{y\notin\Lambda}Z(x-y). \tag{2.77}$$

The missing row m_x obeys $|m_x| \leq C\theta R^{-s}e^{-a\delta(x)/R}$. We anchor the two-site summands at a site of the pair and use the same assignment for the head and tail. For a local test the chart is taken on a

fixed enlarged tile containing the anchor and the endpoints of the head summands. When testing long summands on several tiles we use directions satisfying the stated field bound at each endpoint.

Write $w_b(x) = (\rho_b(x)/b)^p$. For $F_b(v) \le \ell_b$, one has $|v_x| \le \ell_b w_b(x)$ and $w_b(y) \le C w_b(x)(1 + |x - y|/R)^p$. For the boundary forms, absorb the latter polynomial in a fraction of the exponential row weight. Each physical normal layer intersects an enlarged tile in at most Cb^{d-1} sites, including tiles with several faces. Consequently

$$\sum_{x \in U} e^{-a\delta(x)/(2R)} w_b(x)^2 \le C b^{d-1} R (R/b)^{2p}. \tag{2.78}$$

Multiplying by $\theta R^{-s} \ell_b^2$ proves the boundary estimate.

For span $|x - y| \le c\rho_b(x)$, summing first differences in the enlarged tile gives

$$|v_x - v_y| \le C \ell_b b^{-p} \rho_b(x)^{p-1} |x - y|. \tag{2.79}$$

For longer spans one uses the endpoint bound and $\sum_{|z|>c\rho} |Z(z)|(1 + |z|/R)^{2p} \le C\theta R^{-s} e^{-c'\rho/R}$. These two estimates bound the difference-square sum on the tile by

$$C\theta \ell_b^2 b^{-2p} R^{2-s} \sum_{x \in U} \rho_b(x)^{2p-2}. \tag{2.80}$$

For the long spans the same bound follows from $t^2 e^{-c't} \le C$ for $t = \rho_b(x)/R$. On interior tiles $\rho_b \asymp b$. In general, summing the layers with $\delta(x) < b$ and then the interior gives

$$\sum_{x \in U} \rho_b(x)^{2p-2} \le C b^{d-1} (b^{2p-1} + R^{2p-1}). \tag{2.81}$$

Since $\ell_b^2 = b^{-s}$ and $d = 2s$, these are the desired powers of M_R. The same estimates bound the bilinear form by applying the pointwise bounds to each direction.

Take the bulk head to have $|x - y| \le \vartheta b$, the D-head to have $|x - y| + \min(\delta(x), \delta(y)) \le \vartheta b$, and the missing-row head to have $\delta(x) \le \vartheta b$, for any fixed $\vartheta > 0$. Cutting the bulk difference square preserves its zero row in the infinite lattice. On a discarded term, the unused exponential row weight supplies the factor $Ce^{-cb/R}$; for a long bulk span one can use the long-span estimate without taking a derivative along that span. One can put back every discarded term using the same pair anchor. This proves the tail bounds and the exact decomposition. □

The leading thermal term before the pullback is $-(\beta - 1)I$. Its pullback has a bulk constant part, a bulk zero-row part, and a boundary comparison with the same weights as in Proposition 2.16. After division by the initial mass radius $\Lambda_R G_0$, the two nonscalar terms change by $O(A_0(M_R^{-2} + M_R^{-1-2p}))$ across the window. All other thermal changes in the weighted microscopic Hessian are $O((A_0 + A_0^2) g_{\mathrm{mic}}^2)$ before canonical conversion, by the exact formula for T_β and its weighted derivative estimates. Thus the nonscalar initial coordinates can use constants fixed before choosing the thermal window, by taking R large after fixing A_0. The boundary estimate decreases with M_R in all four dimensions. The bulk estimate leaves an initial kinetic term of size $O(\lambda G_0)$ in dimension four.

3 Boundary transport and localization

Passing from an interval to a cube changes the boundary contribution to a single renormalization-group step. A coarse face contains many fine tiles. To sum their contributions we test fields with a weight that decays towards the physical faces and retain the local quadratic forms supported in a

collar of those faces. An *anchor* specifies the site or tile to which a polynomial is assigned. Several polynomials may depend on fields at the same site while having different anchors. Keeping track of the anchors makes the covariance subtraction compatible with the reblocking of polymers.

Two results connect this construction to the shooting argument of Section 6. Corollary 3.16 gives the one-step estimates in dimensions one through three. In dimension four the bulk kinetic coefficient is marginal. Proposition 4.7 gives the corresponding estimates with the Gaussian reference fixed and a separate equation for the residual kinetic coefficient.

3.1 Face geometry and the change in the RG chart

Set $s = d/2$, $\ell_b = b^{-s/2}$, and $B = Lb$, where b is the side length of a fine tile. Fixed enlargements of the supports are absorbed in the constants independent of L. On a cube touching a physical face, let

$$\rho_b(x) = \min(b, R + \operatorname{dist}(x, \Lambda^c)), \qquad F_b(v) = \max_{|\gamma|\le P} \sup_x \frac{\rho_b(x)^{|\gamma|}|\nabla^\gamma v(x)|}{(\rho_b(x)/b)^p}. \tag{3.1}$$

Write $E_b(v) = \inf_{a \text{ affine}}\{F_b(v-a) + b^{-d/2}\|a\|_2\}$ for the norm obtained by adjoining affine fields. Two estimates for unrestricted boundary tensors explain our choice of test space.

Lemma 3.1 (Unweighted sum of face tensors)**.** *Let a coarse face tile of side $B = Lb$ meet a full layer of fine face tiles of side b. On each fine tile let a homogeneous k-linear form have absolute value at most one whenever all its arguments satisfy $E_b(v) \le \ell_b$. The norm of their sum on directions with $E_B(v) \le \ell_B$ is at most $CL^{d-1-ks/2}$. The power of L is attained up to a constant.*

Proof. A new face tile of side B meets order L^{d-1} old face tiles of side b. Restriction of a coarse E_B-unit direction to a fine tile has bounded fine E_b-norm. This follows by restricting the affine decomposition, since affine coefficients on the coarse cube have size $O(1)$ in supremum and slope $O(B^{-1})$, while $\rho_b \le \rho_B$ and $(\rho_B/B)^p \le (\rho_b/b)^p$ for stencils in the tile. Thus a homogeneous degree-k form of normalized fine norm one has coarse norm at most $CL^{-ks/2}$ on one fine tile. The sum of the old face forms has norm at most $CL^{d-1-ks/2}$. This power is sharp on unrestricted tensors. Take a site x_i in each tile and forms $\ell_b^{-k} v(x_i)^k$. Their fine norms are uniformly bounded since E_b controls the supremum. Evaluate their sum on the constant coarse direction $\ell_B \mathbf{1}$, whose E_B-norm is at most ℓ_B. The value is of order $L^{d-1-ks/2}$. □

In particular, the quadratic power

$$L^{d-1-s} = L^{s-1} \tag{3.2}$$

equals one in $d = 2$ and grows with L in $d = 3, 4$. Gaussian integration of a quadratic changes only its constant term. Thus the same power occurs in the linearized RG map. In dimension one there are only two endpoints and the factor is $L^{-1/2}$. In higher dimensions our boundary test space gives an additional gain for each field direction.

Loss of the affine product bound. In dimensions $d \ge 2$, on a cube of side b touching $x_1 = 1$, take $v(x) = x_2/b$. As v is affine, $E_b(v) = O(1)$. For any affine a, evaluate $v^2 - a$ on three equally spaced tangential sites of spacing $\lfloor (b-1)/3 \rfloor$ in the layer $x_1 = 1$. Its second difference across these sites has absolute value bounded below by a positive constant, since the affine part has zero second difference there. Therefore some one of these values is bounded below in modulus, whereas its boundary weight is $O((R/b)^p)$. Taking the infimum over a gives

$$E_b(v^2) \ge c(b/R)^p. \tag{3.3}$$

Thus in $d \geq 2$ the product estimate for E_b fails for every $p > 0$. This affects the large-field regulator and the matching of local spatial jets. At the level of the quotient restriction estimate, the direct summing argument also encounters order L^{d-1} tiles in each boundary layer. Indeed a raw coarse direction on the layer at distance ib has local affine quotient bounded by $CL^{-p}(1+i)^{p-2}$ after subtracting its first-order Taylor polynomial. Summing the second differences along coordinate segments of length $O(b)$ gives the stated bound. For the first layer use direct raw restriction and the factor L^{-p}. Summing the squares gives CL^{d-1-2p} if $0 < p < 1$, to which the large-field radii contribute the factor L^{-s} at comparable couplings. The resulting power L^{s-1-2p} requires $2p > s-1$ to yield a small factor.

Ordinary spatial localisation. For mass localisation, cancelling evaluation on constant fields alone gives a factor $O(b/\rho_B(x_i))$ from the first-order variation of the field across a fine placement x_i. Even on a coarse cube at distance at least B from the physical boundary, summing over order L^d placements gives the power $L^d L^{-s} L^{-1} = L^{s-1}$. The order in this estimate cannot be improved for arbitrary jets from this cancellation: the diagonal quadratic difference $v(x_i+te)^2 - v(x_i)^2$, with $t \asymp b$, has zero constant-field evaluation and zero relative second moment, yet is of order b/B on affine coarse fields in a region of fixed rescaled width. An additional first-spatial-moment cancellation gives $O((b/B)^2)$ in the interior, hence $O(L^{s-2})$, which suffices there for $d = 1, 2, 3$. Proposition 3.2 implements the moment cancellations also for marks crossing coarse faces. A fixed coarse face can meet order L^{d-1} translates of a mark crossing the interface. Even with $O(L^d)$ placements, the nine-direction Taylor remainder gains $L^{d-9s/2} = L^{-5s/2} = L^{-5d/4}$, since $d = 2s$.

A weighted sunset bound for $d = 1, 2, 3$. If a cubic sunset difference-square kernel obeys

$$|S(x,y)| \leq Cg^2\lambda^3(R + |x-y|)^{-d-s}, \qquad S(x,y) = 0 \quad (|x-y| > Cb), \tag{3.4}$$

then its quadratic form on an enlarged face tile has bounded normalized E_b-norm whenever $2p > s-1$. To see this, write a direction of E_b-norm at most one as affine plus a raw component of bounded F_b-norm. For the affine component use $|v_x - v_y| \leq C|x-y|/b$; the summed kernel row is $O(b^{-s})$ since $s < 2$. For a raw component f, for $|x-y| \leq c\rho_b(x)$ its difference is at most $C(\rho_b(x)/b)^p |x-y|/\rho_b(x)$, by summing first differences on the tile. For longer spans it is bounded by $C(\min(b, \rho_b(x) + |x-y|)/b)^p$. For $0 < 2p < s$, shell summation at fixed x gives an energy row bounded by $Cg^2\lambda^3 b^{-2p}\rho_b(x)^{2p-s}$. The assumption $2p > s-1$ then yields $\sum_x \rho_b(x)^{2p-s} \leq Cb^{d+2p-s}$, by summing normal layers each having at most Cb^{d-1} sites. After multiplying by the two small-field radii $\ell_b^2 = b^{-s}$ the estimate is $O(g^2\lambda^3)$, since $d = 2s$. For $2p \geq s$ reduce p to any value with $s - 1 < 2p < s$, which weakens the weight. Polarization supplies the symmetric tensor norm. The estimate uses only the kernel envelope and the weight. In dimension three it requires $p > 1/4$. The product estimate for the field regulator will use the raw norm F_b.

Kinetic coordinates at $d = 4$. A normalized coefficient of a bulk quadratic with two spatial differences transports with factor L^{s-2}, which is one at $s = 2$. Thus it cannot be placed in the contracting stable block of the shooting theorem. Also, for a cumulative Green envelope of order $(R + |x|)^{-s}$, the absolute second moment of its cube is bounded by a constant times

$$\sum_{m=1}^{Cb} m^{d+1}(R+m)^{-3s}. \tag{3.5}$$

This is $O(b^{2-s})$ for $s < 2$ and logarithmic at $s = 2$. At this endpoint the kinetic coefficient and its derivatives must therefore be followed over the full trajectory. A source of order $O(g_j^2)$ is summable

when $g_j \asymp (g_0^{-1} + cj)^{-1}$. Proposition 4.5 bounds the successive increments, and Proposition 4.7 treats their coupling to the collar.

The boundary field norm and its affine quotient. The raw norm F_b is a uniform algebra by the discrete Leibniz rule, since the weights are at most one and comparable on fixed stencils. For a tile U of side comparable to b, let $\mathcal{A}_U$ be the affine fields vanishing on the planes of all physical faces within distance $2b$ of U. Other fixed dilation factors work as well, using nested enlargements for fine and coarse tiles. If there is one such face, an element of $\mathcal{A}_U$ is linear in the distance from that plane. If there are two distinct face planes, it is zero. If there are none, the raw weights on U are one. These observations and scaling of affine polynomials on a cube give

$$F_b(a) \le Cb^{-d/2}\|a\|_{2,U} \quad (a \in \mathcal{A}_U,\ 0 < p < 1). \tag{3.6}$$

For one face, $|a(x)| \le Cb^{-d/2}\|a\|_2 \operatorname{dist}(x, \text{face})/b$; when $\rho_b(x) < b$ the nearest physical face must be among the included faces. The derivative estimates follow from the bounded normalized slope. In particular, with $q_U(v) = \inf_{a\in\mathcal{A}_U} F_b(v-a)$ one has

$$F_b(v) \le C\{q_U(v) + b^{-d/2}\|v\|_{2,U}\}. \tag{3.7}$$

Indeed an approximating a has normalized L^2-norm at most $b^{-d/2}\|v\|_2 + CF_b(v-a)$. This quotient therefore suits quadratic damping by the local L^2-norm and does not require an affine-completed algebra. Coarse affines in this definition restrict to fine affines of the allowed type.

In the raw norm, restriction from a coarse unit direction to fine tiles in the outer $O(b)$ physical face layer gains CL^{-p} in each direction, since the two distance functions are comparable on that layer. Therefore full degree-two face forms with bounded fine norms give coarse norm $O(L^{s-1-2p})$, a power strictly below one when $2p > s-1$. The bound contracts for any $p > 0$ in dimensions one and two, and for $p > 1/4$ in dimension three. The same computation gives the summed quotient estimate, since the fine quotients annihilate the restricted coarse affine. On the remaining layer $i \ge 1$ at distance comparable to ib use fine affine Taylor subtraction, giving the local quotient factor $CL^{-p}i^{p-2}$.

For ordinary quadratic residuals, suppose the local projection cancels the constant evaluation and the evaluation with one constant and one linear coordinate. Taylor expansion of the coarse raw directions on a layer at depth ib then costs $C((i+1)/L)^{2p}(i+1)^{-2}$ per placement, before multiplication by the ratio $\ell_B^2/\ell_b^2 = L^{-s}$. For this estimate, terms with two linear increments cost the square of the first-order variation, and a single Taylor remainder has the same bound. There are at most CL^{d-1} placements at any fixed layer. Summation gives $O(L^{s-1-2p})$ for $p < 1/2$, including the interior contribution $O(L^{s-2})$. For the quartic, constant evaluation cancellation alone gives $C((i+1)/L)^{4p}(i+1)^{-1}$ per placement before the factor L^{-d}, and hence $O(L^{-1})$ for $p > 0$. These powers give contraction for $d = 1, 2$ and, with the stated boundary weight, for $d = 3$. To use the estimates in a polymer activity we must allocate the moments on the actual cell labels and control the Gaussian translation in the regulator. The allocation is given in Proposition 3.2. The fluctuation moments of Section 2.3 enter the product estimate of Lemma 3.10.

Activity norm and circle product. At scale b let X be a set of disjoint cell labels, each with a fixed enlarged field support P^*, clipped to the physical cube. A field direction on X is tested in $\|v\|_X = \max_{P\in X} F_{b,P^*}(v)$. For a field-dependent activity $K(X)$ use the finite jet seminorm

$$|K(X)|_{r,\phi} = \sum_{k=0}^{M} \frac{r^k}{k!} \sup_{\|v_i\|_X \le 1} |D^k K(X,\phi)[v_1, \dots, v_k]|, \qquad M \ge 9. \tag{3.8}$$

Here $g_b > 0$ is the reference coupling used to set the radii $\ell_b = b^{-s/2}$ and $h_b = g_b^{-1/4}\ell_b$. The two regulators are

$$\begin{aligned} G_b(X,\phi) &= \exp\left(c_\ell \sum_{P\in X} F_{b,P^*}(\phi)^2/\ell_b^2 \right), \\ \mathcal{R}_b(X,\phi) &= \exp\left(\sum_{P\in X} \left[-\kappa \frac{\|\phi\|_{2,P}^2}{b^d h_b^2} + c_h \frac{q_{P^*}(\phi)^2}{h_b^2} \right] \right). \end{aligned} \tag{3.9}$$

The mixed norm on one polymer is

$$\|K(X)\|_{\mathrm{mix},b} = \max\left\{ \sup_\phi \frac{|K(X)|_{\ell_b,\phi}}{G_b(X,\phi)},\ g_b^{9/4} \sup_\phi \frac{|K(X)|_{h_b,\phi}}{\mathcal{R}_b(X,\phi)} \right\}. \tag{3.10}$$

The activity norm takes the supremum over connected X, with the size weight $\varrho^{(|X|-N_s)_+}$, $\varrho > 1$, of Section 3.6. The anchored sum used at the initial scale bounds this supremum for any fixed size weight. A direction on a union restricts to each factor with norm at most one. The finite-jet product rule and the product over disjoint cell labels therefore apply even when their stencils overlap.

Write $I(X) = \prod_{P\in X} I(P)$, extend K by products over the connected components of its argument, and set $K(\varnothing) = 1$. The weight on X is the circle product

$$(I \circ K)(X) = \sum_{Y\subseteq X} I(X\setminus Y) K(Y). \tag{3.11}$$

A step integrates this weight against the fluctuation covariance, regroups the fine labels by their coarse supports, and allocates the retained Taylor coefficients by Proposition 3.2.

The corrected cubic row. The three-line covariance term contains a logarithmic row sum when $2s = d$. Subtracting the local diagonal from the current slice gives a bound uniform in the scale under the following hypotheses. Suppose the cumulative kernel W and the current Gaussian covariance Γ have range Cb, $|W(x,y)| \le C\lambda(R + |x-y|)^{-s}$, and the diagonal and first-difference standard deviations of Γ are bounded by $C(\lambda b^{-s})^{1/2} w(x)$ and $C(\lambda b^{-s})^{1/2} w(x)/\rho_b(x)$, respectively, with $w(x) = (\rho_b(x)/b)^p$ and $p > 0$. The Cauchy–Schwarz inequality for Γ, applied to a difference stencil and then telescoped, gives

$$|\Gamma(x,y) - \Gamma(x,x)| \le C\lambda b^{-s} w(x)^2 |x-y|/\rho_b(x) \qquad (|x-y| \le c\rho_b(x)). \tag{3.12}$$

The sum of $(R + |x-y|)^{-2s}|x-y|/\rho_b(x)$ over this near region is bounded by a constant, since $2s = d$. On the far region use $|\Gamma(x,y)| \le C\lambda b^{-s} w(x) w(y)$ and $w(y) \le C(\min(b, \rho_b(x) + |x-y|)/b)^p$. There the integral of $r^{-1}(r/b)^p$ on $[c\rho_b(x), Cb]$ is bounded. The subtracted diagonal contributes at most $C\lambda b^{-s} w(x)^2 \log(2 + b/\rho_b(x))$, also bounded by $C_p \lambda b^{-s}$. Thus

$$\sup_x \sum_y |W(x,y)|^2 |\Gamma(x,y) - \Gamma(x,x)| \le C_p \lambda^3 b^{-s}. \tag{3.13}$$

The same argument treats the symmetric diagonal average, with the far y-diagonal integral using $r^{-1}(r/b)^{2p}$. The other two increment terms $W\Gamma^2$ and Γ^3 give the same row bound by finite range and the pointwise estimates. This proves the corrected row inequality in all four dimensions under the stated slice bounds.

Connected covers in the initialization. In higher dimension a rectangular hull can have cardinality of order the d-th power of its diameter, so its exponential polymer weight cannot generally be paid by a linearly exponential filter-tail bound. For a tail edge between two fine tiles choose a grid path of length at most $C(1 + \operatorname{dist}/b_0)$. A boundary mark can likewise be joined to a closest physical face by a normal path. Include the actual occupied labels and these paths, rather than filling a rectangle. For fixed root tile D, the number of pairs of separation at most k in grid units whose path contains D is at most $C(1+k)^{2d}$, since both endpoints must be within $C(1+k)$ of D. For a boundary promotion use the corresponding count with k equal to the pair span plus the distance to the selected face, in grid units. The boundary-defect weight contains the exponential in this additional distance. Consequently these weights pay $e^{A|\mathrm{path}|}$ for any fixed A when b_0/R is large, and leave a summable edge weight at each root. Averages over coordinate-order paths and over tied closest faces preserve the cube symmetries and change the rooted bounds by at most fixed constants. Proposition 5.1 uses this count to regroup the filter tails together with the microscopic pressure terms in the activity norm.

Large-field reserve for the adapted quotient. For $d = 1, 2, 3$ and $2p > s - 1$, the adapted quotient controls the field on the complement of the occupied labels. Assume a finite primary filter T with uniformly bounded absolute rows and columns and a finite forward filter B_{rec} of range $O(R)$ such that $\|I - B_{\mathrm{rec}}T\|_{2\to 2} \le \epsilon$. Use the local decaying regulator with tile volume b^d. In a coarse tile, normalize its squared background norm Z^2, the squared norm A^2 on occupied sites, and the squared filtered norm S^2 on unoccupied rows by $B^d h_B^2$. Write $Q = q_{C^*}(\phi)/h_B$, where C^* is its enlarged tile.

Outside strips of width $O(R)$ along the fine occupied interfaces and the coarse faces, forward recovery on an unoccupied site uses only unoccupied rows of $T\phi$. The total strip volume is $O((R/b)B^d)$. The norm bound $F_B(\phi)/h_B \le C(Z + Q)$ follows as above with the affine norm on the smaller tile in the right-hand side; normalized L^2 on a tile controls the affine norm on its fixed enlargement. The recovery identity, squared on the unoccupied sites, yields

$$(1-\delta)Z^2 \le A^2 + CS^2 + \delta Q^2, \qquad \delta = C(R/b + \epsilon^2). \tag{3.14}$$

If S uses a rowwise convex combination of $T\phi$ and $T(\phi + \zeta)$, add $C\|T\zeta\|_{2,C}^2/(B^d h_B^2)$ on the right. The quartic bound on the unoccupied labels supplies any fixed negative multiple of S^2, since $g_b B^d h_B^4$ is bounded above and below. The occupied decay supplies a multiple of A^2 of order $\kappa B^d h_B^2/(b^d h_b^2) \asymp \kappa L^s$. Using the summed quotient gain and then taking $R/b + \epsilon^2$ small yields a strict negative Z^2 margin and positive Q^2 coefficient strictly below the output allowance. With uniform Gaussian exponential moments for F_b, finite-range coloring bounds the additional fluctuation cost by a constant per fine label. Fixed field jets and small quadratic perturbations are absorbed by the strict margins in the two squared norms, since any fixed power is bounded by a Gaussian exponential with arbitrarily small positive coefficient. Lemma 3.10 applies this estimate to products whose occupied labels and field stencils can differ.

3.2 An exact cube projection for the low field jets

For a quadratic Taylor coefficient, the spatial moments of order two test a bilinear form in two ways. Both directions may be linear, or one may be constant and the other quadratic. Matching both, as well as the lower moments, gives a small remainder on each tile when $2p > s - 1$, including at $d = 4$. The bulk kinetic coefficient retained by this projection is marginal at $s = 2$.

Proposition 3.2 (Local moment allocation on cube tiles)**.** *Use an integer scale ratio L, fine tiles of side b, and coarse tiles of side $B = Lb$. A small connected set X of fine tiles has at most N_s tiles*

and its field support lies in an enlargement of diameter $O(b)$. Fix a boundary collar of a sufficiently large constant number of fine layers. Declare X interior if it meets no tile of this collar. Assume the ordinary filtered quartic uses a finite-range filter of row sum one and bounded absolute row sum on interior rows. There is a linear projection on even Taylor jets through field order eight, with an allocation to the labels in X, such that:

1. *The allocation is bounded in the radius-ℓ_b jet norm using the raw field space F_b, uniformly in b and the cube. It preserves the vacuum. At boundary placements it extracts the full tensors of degrees two, four and six.*

2. *At interior placements, the quadratic residual has zero evaluation on any pair of polynomial directions whose spatial degrees sum to at most two. The quartic residual vanishes on constant fields.*

3. *Suppose $0 < p < 1$, $2p > s - 1$, and $P \geq 3$ in the field norm. For the sum of the residuals of small marks with a fixed coarse closure, the normalized zero-field jet norm is bounded by the sum's maximal fine input norm times*

$$C\big(L^{s-1-2p} + L^{s-3} + L^{-1} + L^{-s}\big). \tag{3.15}$$

   Here and in this estimate the small-set threshold and all stencil enlargements are fixed before L.

The allocation is covariant under translations and cube symmetries respecting the collar. In the uniform bulk quadratic potential its additional spatial-moment densities reduce to finite-range kinetic densities and lattice divergences.

Proof. Write an even jet as $f_0 + f_2 + f_4 + f_6 + f_8$, and write $f_2(\phi) = A(\phi, \phi)$ with A symmetric bilinear. Use $z_i(x) = (x_i - o_i)/b$ for coordinates on the fixed support enlargement of a small set. On interior sets all these coordinates and their products of degree at most two have F_b-norm bounded by a constant. The quadratic moments to reproduce are

$$A(1,1), \qquad A(1,z_i), \qquad A(z_i,z_j), \qquad A(1,z_iz_j) \quad (i \leq j) \tag{3.16}$$

with all i in the first-moment family. Moment matching in this list does not depend on the choice of o.

Let e_i be the lattice basis vectors, set $D_i\phi(x) = \frac{b}{2}(\phi(x + e_i) - \phi(x - e_i))$, and write $\langle\cdot\rangle_P$ for averaging over the sites of an ordinary tile P. On this tile use

$$\begin{aligned} M_P(\phi) &= \langle\phi^2\rangle_P, & J_{P,i}(\phi) &= \langle\phi D_i\phi\rangle_P, \\ Z_{P,ij}(\phi) &= \langle D_i\phi D_j\phi\rangle_P, & H_{P,ij}(\phi) &= \langle\phi D_iD_j\phi\rangle_P. \end{aligned} \tag{3.17}$$

The inverse of their moment matrix is uniformly bounded. In fact M_P has mass moment one. Each $J_{P,i}$ has mass moment zero and first-moment vector $\frac{1}{2}e_i$. The $Z_{P,ij}$ and $H_{P,ij}$ have both lower moments zero. In the two remaining moment families, $Z_{P,ij}$ gives the elementary symmetric (i,j)-matrix in $A(z_k,z_l)$ and zero in $A(1,z_kz_l)$. The $H_{P,ij}$ gives zero in the first of these families and the elementary symmetric matrix in the second. The nonzero diagonal entries in these last two blocks are 1 or 1/2. All lower-order cross entries are bounded averages of products of the z_i. This proves the inverse bound and gives, for each P, a unique quadratic combination reproducing the moments of f_2. Its coefficient norm is bounded by that of f_2, because D_i and D_iD_j are bounded evaluation operators on the interior stencils in the fine F_b units.

Multiply this combination by $|P|/\sum_{Q\in X}|Q|$ and assign it to P. With the same weights assign the vacuum f_0 and the quartic $f_4(1)\langle(T\phi)^4\rangle_P$, where T is the fixed ordinary filter. The latter density equals one on the unit constant field. Summing the allocations gives a projection because it preserves all the input moments and depends only on them. On a boundary set use the full homogeneous degrees zero, two, four and six, dividing the tensor equally among the tiles of X in the collar and retaining its actual field stencils. Fine field restriction to the selected stencils has uniformly bounded norm in a fixed enlargement. The construction is symmetric in cases of ties between boundary tiles. In the interior the moments and representing span transform covariantly under coordinate reflections and permutations, so uniqueness of the per-tile combinations gives covariance of the allocation.

For the spatial bound, all ordinary supports assigned to a fixed coarse closure have diameter $O(b)$ each, and the closure uses a bounded number of coarse tiles. Group the supports in layers indexed by $i \geq 1$ with $\rho_B \asymp ib$ for $i < L$, and put the region $\rho_B = B$ in the last layer. A layer below B contains $O(L^{d-1})$ fine tiles, including near edges and corners. The interior part with $\rho_B = B$ contains $O(L^d)$ tiles. Each fine tile belongs to only a bounded number of small sets. For a coarse unit direction, its spatial Taylor polynomial of degree two and the remainder on an ordinary support have fine bounds, respectively,

$$C(i/L)^p i^{-a} \quad (a = 0, 1, 2), \qquad C(i/L)^p i^{-3} \quad \text{for the remainder.} \tag{3.18}$$

To obtain these bounds take the discrete Newton polynomial using $\prod_j \binom{x_j - o_j}{\gamma_j}$ through total degree two. Repeated telescoping on coordinate segments within the enlargement bounds its remainder after $|\beta| \leq 2$ differences by $Cb^{3-|\beta|}\max_{|\gamma|=3}|\nabla^\gamma \phi|$. Higher differences are those of ϕ itself. Because the fine weights here are one and the coarse depth on the segments is comparable to ib, these bounds give the asserted F_b estimates. Polynomial terms of total degree at most two cancel in testing the residual bilinear form. Every remaining term costs $C(i/L)^{2p} i^{-3}$. The ratio of quadratic radii is L^{-s}. Summing the layer estimate gives $CL^{d-1-s-2p}$, since $\sum_{i\geq 1} i^{2p-3} < \infty$. The last layer costs CL^{d-s-3}. Both exponents have the stated form because $d = 2s$. The quartic uses constant cancellation and costs $C(i/L)^{4p} i^{-1}$ before the radius factor L^{-d}. The resulting layer and interior sums are $O(L^{-1})$ for $p > 0$. Boundary placements have no residual degree below eight. Everywhere, all residual degrees at least six gain L^{-s} after summing $O(L^d)$ terms. This proves the zero-field bound.

Finally, the necessary extra densities can be retained literally during the allocation. At common coefficients on identical tiles, $J_{P,i}$ and $Z_{P,ij} + H_{P,ij}$ sum to zero in the bulk by skew symmetry of D_i. With the summation restricted to the ordinary sites S the two identities are

$$\begin{aligned} \sum_{x\in S} \phi(x) D_i \phi(x) &= \tfrac{b}{2} \sum_x \phi(x)\phi(x+e_i)(\mathbf{1}_S(x) - \mathbf{1}_S(x+e_i)), \\ \sum_{x\in S} (D_i\phi(x) D_j\phi(x) + \phi(x) D_i D_j \phi(x)) &= -\sum_x [D_i, \mathbf{1}_S]\phi(x) D_j \phi(x). \end{aligned} \tag{3.19}$$

Their stencils meet the interface of the ordinary region and the collar. For a fine tile on this interface, the number of crossing bonds is $O(b^{d-1})$, and the commutator costs $O(b)$. Before the tile normalization, ϕ and $D_j\phi$ have bounded evaluation norms at this positive distance from the physical boundary. Consequently the corresponding shifted collar tensor has bounded fine norm after dividing by the tile volume. An allocation which agrees with the bulk outside the fixed halo has its nonuniform coefficients supported in the halo in addition to these interface terms. Thus the reassignment can be made without identifying distinct block densities pointwise. □

In particular, on a small set the allocation can be applied to the finite jet quotient $K(X)/I^X$. Suppose the normalized jet of $I(P)$ has vacuum bounded away from zero and, after division by

its vacuum, equals $1 + O(g)$. The finite jet inverse has bounded norm and is analytic on the corresponding complex reference box. If the reference is analytic there with bounded first derivatives in the normalized potential directions and $\tau = \|K\|$ in the mixed norm, the allocation satisfies

$$\|q_P\|_{\ell_b,0} \le C\tau, \qquad \|D_K q_P[\dot{K}]\|_{\ell_b,0} \le C\|\dot{K}\|, \qquad \|D_V q_P[\dot{V}]\|_{\ell_b,0} \le C\tau\|\dot{V}\|. \tag{3.20}$$

Indeed the first two bounds follow by composing the finite jet inverse with the bounded allocations and summing the bounded list of small sets through P. The last follows from $D(I^{-1}) = -I^{-1}(DI)I^{-1}$ in the same algebra. Writing $q_P = \sum_{X \ni P} q_{X,P}$, the usual reapportioning $J_P(X) = q_{X,P}$ for $X \neq \{P\}$, adjusted at the singleton to have zero sum through P, is therefore exact. Changing a block factor I to $\widehat{I}$ in a circle product uses the multiplicative difference with factors $I(P) - \widehat{I}(P)$. This identity follows by the binomial expansion on the uncovered labels. It requires no pointwise inverse of the reference.

Size of the coarse closure. For any connected fine set X, with fixed bounded graph range in fine units, its coarse closure satisfies

$$|\overline{X}| \le C_d(1 + |X|/L). \tag{3.21}$$

Here the closure may include a fixed fine-stencil enlargement. Color the coarse lattice modulo three in every coordinate and choose a fine site of X within fixed fine distance of each coarse tile counted. Two chosen sites of the same color have distance at least cL for L large. A spanning-tree walk in the fine graph visits every site in at most $2(|X| - 1)$ steps, each of bounded length. Consecutive first visits to the chosen sites of that color therefore cost at least cL in fine distance. Summing over colors proves the estimate. Choose N_s and then L large. This gives $|\overline{X}| \le \vartheta|X|$ with $\vartheta < 1$ for $|X| > N_s$, and hence the large-set gap in the size exponent $(|X| - N_s)_+$. It does not use a rectangular hull, which can expand a connected sparse fine set by an arbitrarily large factor.

Proposition 3.2 also covers small marks meeting several coarse tiles, since their coarse closure has bounded size. Section 2.3 supplies the Gaussian weights, including $1/2 < p < 1$ when $d = 4$. We now use the projection to bound the term linear in one marked activity. To pass to the full derivative we will also compare the reference factors on the unoccupied labels.

3.3 The cube single-mark estimate

A single *mark* here is one insertion of a fine activity F_X. Its high Taylor remainder gains a power from each field direction. For the lower coefficients, the gain comes from the spatial moments removed by the projection. The proposition combines these estimates with the quartic damping on the other labels, including when the mark spans more than one coarse tile.

Proposition 3.3 (Contraction of the leading small-mark operator)**.** *Use the raw test norms and adapted quotients in Section 3, and choose $0 < p < 1$, $2p > s - 1$. Let U be the fixed coarse closure of a small fine mark. Run over all small marks X with this closure. Let F_X be an even fine activity on the enlarged stencils of X, with field jets through a fixed order $M \ge 9$ and fine mixed norm at most ν. Use the growing small-field regulator at $\ell_b = b^{-s/2}$, and the decaying and quotient regulator at $h_b \asymp g_b^{-1/4}\ell_b$, with large-field weight $g_b^{9/4}$. The coarse radii are defined using $B = Lb$ and $g_B \asymp g_b$. Assume a Gaussian exponential moment at radius ℓ_b on the fixed fine stencils, with a reserve for the growing small-field regulator there. Put*

$$P_X = T_8\,\mathbb{E}F_X(\phi + \zeta), \qquad R_X = (1 - T_8)\mathbb{E}F_X(\phi + \zeta), \tag{3.22}$$

where T_8 is the Taylor jet through degree eight.

Let I_X^{comp} carry the stable filtered quartic on the unoccupied fine labels, with positive real part of its quartic coefficient comparable to g_b. Use the recovery filter of Section 3. One may include vacuum- normalized perturbative factors on those labels formed from quadratic exponent shifts of radius-ℓ_b size $O(g_b)$ and polynomial factors $1+W$ of degrees at most six with $\|W\|_{\ell_b,0} = O(g_b^2)$. Set

$$\mathcal{M}_U F = \sum_{X:\overline{X}=U} I_X^{\mathrm{comp}}((1-\Pi_X)P_X + R_X), \tag{3.23}$$

with the projection of Proposition 3.2. For sufficiently large L, small filter recovery error and R/b, and then small g_b, this operator satisfies

$$\|\mathcal{M}_U F\|_{\mathrm{mix},B} \le C\big(L^{s-1-2p} + L^{s-3} + L^{-1} + L^{-s} + L^{-s/2} + C_L g_b^{1/4}\big)\nu. \tag{3.24}$$

The coefficient of the inverse powers of L is independent of L, and the final constant C_L may depend on L.

Proof. A fixed U for a small mark has at most a bounded number of coarse labels. There are $O(L^d)$ placements, by the bounded number of shapes through any fine tile. For $C \in U$, write $Z_C = \|\phi\|_{2,C}/(B^{d/2}h_B)$ and $Q_C = q_{C^*}(\phi)/h_B$ for the background field quantities, and put $Z^2 = \sum_C Z_C^2$, $Q^2 = \sum_C Q_C^2$. The adapted quotient yields $F_{B,C^*}(\phi)/h_B \le C(Z_C + Q_C)$.

We first establish a uniform reserve from the complementary reference. In the recovery inequality of Section 3, the normalized occupied squared norm in a coarse tile is now at most $CL^{-d}(Z_C^2 + Q_C^2)$, since the mark contains only a fixed number of fine tiles. Consequently it can be absorbed into the recovery error. With S_C denoting the normalized filtered norm on the unoccupied rows,

$$(1 - O(R/b + \epsilon^2 + L^{-d}))Z_C^2 \le CS_C^2 + C(R/b + \epsilon^2 + L^{-d})Q_C^2, \tag{3.25}$$

where ϵ is the recovery error. The quartic is bounded below in real part by cS_C^4, since $g_b B^d h_B^4 \asymp 1$. Quartic Young's inequality then gives any fixed negative quadratic exponent in Z_C with an arbitrarily small positive quadratic exponent in Q_C, by the choice of L and recovery errors. Up to fixed field-jet order, the quartic factor has derivatives equal to its exponential times polynomials bounded by $C(1+Z+Q)^C$ in coarse radius units. This uses bounded filter row sums and $g_b B^d h_B^4 \asymp 1$. These polynomials are absorbed by a fixed fraction of the reserve.

For the optional perturbative factors, each normalized large-field jet of a quadratic exponent shift is bounded by $Cg_b^{1/2}(1+Z+Q)^2$. Its exponential has this same bound in the exponent, times a polynomial in the bound for its nonzero jets. Each polynomial factor has jet norm at most $1 + Cg_b^{1/2}(1+Z+Q)^6$, by radius conversion of its coefficient. For fixed L the number of factors is bounded by a fixed C_L. Their product, including the tested derivatives, can use an arbitrarily small fraction of the reserve with a bounded prefactor. Indeed, for fixed degree a,

$$\sup_{t\ge 0} \frac{\log(1+\delta(1+t)^a)}{1+t^2} \longrightarrow 0 \quad (\delta \to 0). \tag{3.26}$$

Apply this with the tolerance divided by the number of factors, using the submultiplicative finite-jet seminorm for products. This gives the large-radius bound by the coarse regulator with spare fractions and a constant independent of L after choosing g_b. At coarse small radius and zero field the vacuum-normalized complement differs from 1 by $O_L(g_b)$ in its tested jets.

On each fine mark $\|P_X\|_{\ell_b,0} \le C\nu$, by the small-field bound at ζ and its exponential moment. Thus the sum of the localized low jets is bounded by Proposition 3.2 at zero field. The varying

complementary multipliers add only $C_L g_b \nu$ there, using their difference from 1 and the product rule. At the large radius, monomials of degree $a \le 8$ are bounded by $C_L \nu g_b^{-a/4}$ times a fixed polynomial in Z, Q. The reserve and coarse weight $g_B^{9/4}$ bound their contribution by $C_L g_b^{1/4} \nu$.

For R_X, test the field derivative of order $k \le 8$ by Taylor's integral remainder at $\zeta + t\phi$, $0 \le t \le 1$. There are $9 - k$ background directions and k tested directions. At $k \ge 9$ no Taylor subtraction remains. The large-radius fine cost of each restricted direction is $CL^{-s/2}$, with background directions carrying additionally a factor $C(Z + Q)$. This uses the raw restriction inequality on the enlarged stencils and the ratio h_B/h_b. The fine decaying regulator is at most one. For the quotient at the translated field, use its triangle inequality, then the Gaussian moment for the ζ term. The background term costs $CL^{-s}Q^2$, because there are only a bounded number of marked tiles and the coarse admissible affine fields restrict to fine admissible affine fields. This exponent and all fixed powers of $Z + Q$ fit in the coarse reserve. The Gaussian moment and this same envelope also justify passing the tested derivatives through the expectation. The input mixed cost $\nu g_b^{-9/4}$ is canceled by the output weight. Summing the remainder estimate over the $O(L^d)$ possible placements gives

$$CL^{d-9s/2}\nu = CL^{-5s/2}\nu \le CL^{-s/2}\nu, \tag{3.27}$$

as required by the stated operator bound. This argument includes derivatives of the product with the complement, since the derivative of R_X, of any order used by the product rule, retains at least nine directions from the Taylor formula or the test arguments alone.

At zero field R_X has no jets through order eight. For each higher tested jet, use the fine large-radius input bound at ζ, with scale ratio $\ell_B/h_b \asymp L^{-s/2} g_b^{1/4}$. After multiplying by $g_b^{-9/4}$, its degree $k \ge 9$ again gives $O(L^{-9s/2})$, before summation. Finally the zero-field small-radius and full large-radius bounds imply the full mixed bound. To see this for any field, Taylor-expand each small-radius derivative through total order eight using the zero-field jets. Control the integral remainder by the large-radius bound, using $(\ell_B/h_B)^9 \asymp g_B^{9/4}$. The resulting polynomial in the small-field background norm and the quotient exponential with coefficient $O(g_B^{1/2})$ are absorbed by the growing small-field regulator. □

The annular fluctuation estimates and boundary weights give the hypotheses of the single-mark bound in the adapted chart. To estimate the difference between this operator and the derivative of the full map, Proposition 3.12 accounts for the terms with additional insertions and for the change of reference. The latter comparison must hold also under variation of the boundary and kinetic coefficients.

3.4 Anchorwise covariance cancellation on the cube

The projection assigns each local Taylor coefficient to its anchor. A second-order covariance has two input anchors, and its subtraction is performed before the localized polynomial is replaced by the outgoing one. The next lemma assigns each contribution to an anchor so that their sum reproduces the covariance, even for overlapping field stencils.

Lemma 3.4 (Oriented covariance allocation)**.** *Write a polynomial potential at each anchor i as $U_i = A_i + B_i$. Under $P_C = \exp(\Delta_C/2)$, use as its two outgoing parts $P_C A_i, P_C B_i$, recording constants separately. For any symmetric covariance parameter S, define the polynomial bilinear covariance*

$$\mathcal{D}_S(u,v) = P_S(uv) - (P_S u)(P_S v), \qquad E^S_{ik}(U) = \tfrac{1}{2}\mathcal{D}_S(A_i, A_k) + \tfrac{1}{2}\mathcal{D}_S(B_i, B_k) + \mathcal{D}_S(A_i, B_k). \tag{3.28}$$

Set $X_{S,i}(U) = \sum_k E^S_{ik}(P_{-S}U)$. Then for each individual anchor,

$$P_C X_{S,i}(U) + \sum_k E^C_{ik}(U) = X_{S+C,i}(P_C U). \tag{3.29}$$

In a coarse cell D, the part of the fresh covariance assigned to D is the sum of the i-rows for $i \in D$. For each pair of different coarse cells the two oriented allocations sum to the full pair covariance. Both assertions hold before estimating the polynomial remainders.

Proof. For any polynomial parts u, v, commutation of the heat operators gives

$$\mathcal{D}_{S+C}(P_{-S}u, P_{-S}v) = P_C \mathcal{D}_S(P_{-S}u, P_{-S}v) + \mathcal{D}_C(u, v). \tag{3.30}$$

It applies to each of the three summands defining E_{ik}. This proves the row identity. Bilinearity and symmetry give $E^C_{ik} + E^C_{ki} = \mathcal{D}_C(U_i, U_k)$ for $i \neq k$ and $E^C_{ii} = \frac{1}{2}\mathcal{D}_C(U_i, U_i)$. Summation over the two disjoint sets of labels gives the asserted pair allocations. When finite-range fluctuation variables on the two actual field stencils are independent the corresponding term is zero. The algebra uses the labels for the allocations and allows overlapping stencils. □

For the filtered primary interaction one can put the uniform mass, quartic, and first filtered quadratic head in A_i, and the ordinary kinetic density and fixed quadratic head in B_i. The local density of the first filtered head is the rowwise difference of the filtered square and the uniform square. Thus the quartic first heat splits into the uniform mass and that same shape, at the same anchors. The mixed quartic–kinetic two-line term is charged to the quartic anchor by the orientation in Lemma 3.4. In particular the position of the surviving external filtered square is fixed even at a coarse interface. This identity requires that the outgoing densities on those anchors be kept exactly. The comparison of the resulting polynomial with the outgoing reference is made in Section 3.7.

Lemma 3.5 (Boundary quadratic generated by the quartic)**.** *Write $b = b_j$, $B = Lb$, and $G = g_j^{\mathrm{loc}}$ for the raw quartic coefficient. Use the finite filter and frozen covariance of Sections 2.3 and 2.4. For a coarse tile D, let $\mathcal{T}_{j,D}$ be the quadratic from the Gaussian average of the quartic after assigning its common bulk term. The contributing anchors lie within $O(B)$ of the boundary, and*

$$\|\mathcal{T}_{j,D}\|_{\ell_B,0} \le C L^s \lambda_j^{\mathrm{tad}} |G|, \qquad \lambda_j^{\mathrm{tad}} = \begin{cases} O_L(\lambda), & \text{on an ordinary step,} \\ O_L(\lambda M_R^s), & \text{on the merged step.} \end{cases} \tag{3.31}$$

On the shooting box its value and its derivative in the normalized quartic coordinate are bounded in the outgoing weight $\Lambda_R x_{j+1}$, uniformly in j and n. Its normalized mass derivative is $o(1)$ at fixed L as $R \to \infty$.

Proof. The heat of an incoming quadratic adds only a constant. The heat of the quartic produces a quadratic proportional to $G\gamma(x)(T\phi)_x^2$, where $\gamma(x) = (T\Gamma_j T^*)_{xx}$ uses the actual finite filter. The cube and bulk rows agree beyond the propagation range of this slice and the filter stencils, hence outside a layer of width $O(B)$. To estimate the two diagonal entries, covariance Cauchy–Schwarz gives

$$(T\Gamma_j T^*)_{xx} \le \left(\sum_y |T_{xy}| \sqrt{\Gamma_j(y,y)} \right)^2. \tag{3.32}$$

The filter has bounded absolute row sum for both the cube and the bulk. Thus the diagonal bounds of Sections 2.3 and 2.4 give $C\lambda_j^{\mathrm{tad}} b^{-s}$ in both cases. One may take $\lambda_j^{\mathrm{tad}} = O_L(\lambda)$ at ordinary scales.

At the merged step the sum of the diagonals from scale R to b_0 is $O_L(\lambda R^{-s})$, so $\lambda_0^{\mathrm{tad}} = O_L(\lambda M_R^s)$ suffices.

In the polarization of $\mathcal{T}_{j,D}$, each actual or bulk filtered field direction costs at most $C\ell_B$ on the enlarged stencil, since the raw test norm controls the supremum. This bound applies to each of the two squares in a filter restriction difference and to the bare square subtracted on a promoted anchor. There are at most CB^d sites in D, including at edges and corners. Their total cost is at most $C|G|\lambda_j^{\mathrm{tad}} b^{-s} B^d \ell_B^2$. Since $\ell_B^2 = B^{-s}$ and $d = 2s$, this is (3.31).

The outgoing quadratic has weight $\Lambda_R x_{j+1}$. On the shooting box, $|G| + |\partial_{\gamma_j} G| \le C x_j$ for the normalized quartic coordinate $\gamma_j = (g_j - x_j)/x_j$. At fixed covariance and filter the source is linear in G. Its value and its γ_j-derivative in the outgoing weight are therefore bounded by

$$CL^s \frac{\lambda_j^{\mathrm{tad}}}{\Lambda_R} \frac{x_j}{x_{j+1}} \le C_L. \tag{3.33}$$

Here $\lambda_j^{\mathrm{tad}} \le C_L \Lambda_R$, including the merged step, and x_j/x_{j+1} is uniformly bounded. The source is independent of the incoming collar tensors. Its dependence on the normalized mass $\mu_j = \nu_j/(\Lambda_R x_j)$ comes through G. Lemma 3.14 gives $|\partial_{\nu_j} G| \le C\lambda x_j$. Thus its normalized μ_j-derivative is at most $C_L \lambda_j^{\mathrm{tad}} \lambda x_j = o(1)$ for fixed L as $R \to \infty$. The full quadratic is retained on its anchors in the outgoing boundary potential. At the following step its $O(B)$ support is a collar of the incoming scale B. □

Contractions in the Gaussian integral lower the degree of a retained collar tensor. The next estimate measures each such term in the weight of its output degree, including a quartic contribution to the quadratic coefficient.

Lemma 3.6 (Heat and reblocking on the collar)**.** *Let $s = d/2$, $2p > s - 1$, and let $P_{U,k}$, $k = 2, 4, 6$, be homogeneous polynomials on fixed enlarged stencils of fine tiles U in a collar of width $C_0 b$. Let a_k bound their fine multilinear norms at radius ℓ_b. Suppose the fixed fluctuation covariance has*

$$\sup_U \mathbb{E}\,(F_{b,U^*}(\zeta)/\ell_b)^q \le C_q \epsilon_j^{q/2} \qquad (q = 2, 4, 6). \tag{3.34}$$

Sum the heated polynomials with fine anchors in a given coarse tile and remove the vacuum. The coarse norm of degree $r = 2, 4, 6$ is at most

$$CL^{d-1-r(s/2+p)} \sum_{k \ge r} \epsilon_j^{(k-r)/2} a_k. \tag{3.35}$$

The same estimate bounds the operator on variations of the incoming polynomials at fixed covariance. The outgoing nonvacuum polynomials remain on stencils in a physical collar of width $O(b)$ covered by the enlarged coarse tiles.

Proof. In the collar both depth weights are comparable on a fixed fine stencil, so $F_{b,U^*}(v) \le CL^{-p} F_{B,D^*}(v)$ on a coarse direction. This follows by taking the ratio of the two weights in each difference seminorm and using $\rho_b \le \rho_B \le C\rho_b$. The ratio $\ell_B/\ell_b = L^{-s/2}$ contributes the remaining factor for each coarse direction. For the symmetric k-linear form of $P_{U,k}$, the degree-r term of the heat substitutes r coarse directions and $k - r$ copies of the fluctuation field. Its multinomial factor depends only on k, and the fluctuation copies cost at most $C\epsilon_j^{(k-r)/2}$. The odd terms vanish by centered Gaussian symmetry. The number of fine anchors in the collar of a coarse tile is $O(L^{d-1})$, by covering the collar with the strips from the physical faces. This count includes anchors at edges and corners. The heated tensor has its original deterministic arguments, so its support does not increase. Applying the same calculation to the variation proves the operator bound. □

Corollary 3.7 (Accumulation of the boundary quadratic)**.** *Use the comparison sequence of Section 6. Give the quadratic collar tensor weight $\Lambda_R x_j$ and the quartic and sextic tensors weight x_j^2. Under the hypotheses of Lemmas 3.5 and 3.6, suppose $\epsilon_j = O_L(\lambda)$ on ordinary steps and $\epsilon_0 = o(1)$ on the merged step as $R \to \infty$ at fixed L. For sufficiently large L and then sufficiently large R, the linear transport of the inherited collar tensors is strictly contractive in these weights. In particular, the sum of the quadratics in* (3.31) *transported from all preceding steps is bounded in units of $\Lambda_R x_j$, uniformly in the scale and in the side length of the cube.*

Proof. For the original reference, the merged-scale stencil estimate in Section 2.3 gives $\epsilon_0 \le C_L \lambda M_R^{2P+2p+s}(1+\log M_R) = o(1)$. The smaller diagonal estimate $O_L(\lambda M_R^s)$ suffices for the newly generated term in (3.31). For inherited tensors the extra factor in the degree-six to degree-four block is $O(\epsilon_j)$, and in the degree-four to degree-two block it is $O(\epsilon_j x_j/\Lambda_R)$. Both are uniformly small since $x_j \le x_0$ and $x_0/\Lambda_R = o(1)$. The diagonal ratios from successive comparison couplings tend to one as $R \to \infty$. Thus the linear transport on these weights is strictly contractive by taking L large and then R large. In particular the inherited quadratic has diagonal factor $CL^{s-1-2p}x_j/x_{j+1}$. Fix L so that this factor is bounded by $\kappa < 1$ for all sufficiently large R. A source as in (3.31), generated at each step and then transported as an inherited quadratic, contributes at most $C_L \sum_{m\ge 0}\kappa^m = C_L/(1-\kappa)$ in the normalized quadratic coordinate. □

The quartic derivative of this source is a bounded entry from the quartic to the stable coordinate in Assumption 6.1. The radii and relative weights in the path contraction may be chosen after L is fixed. A change of outgoing reference can introduce further coefficient increments, for which this linear estimate gives the transport only after the increments have been bounded in the same weights.

3.5 A local estimate for the moving reference

Extracting a Taylor polynomial from a polymer changes the reference on each label to which a coefficient is assigned. The linear term in this change is used in the extraction. The estimate below controls the remainder at both field radii.

Lemma 3.8 (Remainder in a local reference change)**.** *Let $I = e^{-A}(1+W)$ be a normalized reference on one label, with the stable filtered quartic and the allowed quadratic term in A. Assume $0 < g \le 1$, $\deg W \le 6$, and $\|W\|_{\ell_b,0} \le Cg^2$. Let q_A, q_B be the extracted polynomials, each of norm at most $C\tau$ at ℓ_b, where $\tau = \|K\| \le cg^{9/4}$. On a boundary label q_A has degree at most two and q_B has terms of degrees four and six. On an ordinary label q_A has degree at most four, with quartic term in the primary coefficient, and $q_B = 0$. Suppose $\deg \Delta W \le 6$, $\|\Delta W\|_{\ell_b,0} \le Cg\tau$ and put*

$$\widehat{I} = e^{-A+q_A}(1 + W + q_B + \Delta W), \qquad H = \widehat{I} - I - I(q_A + q_B). \tag{3.36}$$

Assume the primary quartic remains in its stable sector along the segment $A - tq_A$, $0 \le t \le 1$, with the field reserve of the single-mark estimate or of the joint product containing the label. Then the contribution of H to the mixed norm is at most $Cg\tau$. If the coefficients are holomorphic and satisfy these bounds on the enlarged activity and potential domains, then on fixed smaller concentric domains the activity derivative is bounded by Cg, and the derivative in a unit normalized potential direction by $Cg\tau$.

Proof. The exact difference is

$$\begin{aligned} H = e^{-A}\big[&(1+W)(e^{q_A} - 1 - q_A) + (e^{q_A} - 1)q_B \\ &- Wq_B + e^{q_A}\Delta W\big]. \end{aligned} \tag{3.37}$$

Apply the integral Taylor formula to the two differences of exponentials. Each resulting exponential is evaluated on the given segment, so the field reserve absorbs the polynomial background powers and their fixed derivatives. At the small radius the resulting bound is $C(g\tau + \tau^2)$.

At the large radius $h_b \asymp g^{-1/4}\ell_b$, a polynomial of degree k costs at most $Cg^{-k/4}$ times its small-radius coefficient bound. In particular W costs $Cg^{1/2}$, and Wq_B and ΔW together cost at most $C\tau g^{-1}$. The remaining terms contain either q_A^2 or $q_A q_B$. Both products have degree at most eight under the stated allocation, so they cost $C\tau^2 g^{-2}$. Multiplication by the mixed weight $g^{9/4}$, and the bound $\tau \le cg^{9/4}$, give $Cg\tau$. Cauchy's formula on a disc of radius comparable to $g^{9/4}$ in the activity variable gives the first derivative bound. Applying the same formula at fixed K on the enlarged normalized potential domain gives the second. □

We also use a Cauchy estimate to pass from the derivative at zero to a nonlinear contraction. If the exact connected output $\Psi_j(V,K)$ is norm holomorphic for $\|K\| < r = cg^{9/4}$ and satisfies $\|\Psi_j(V,K)\| \le C(g^3 + \|K\|)$ there, two Cauchy integrals give $\|D_K^2 \Psi_j(V,K)\| \le C'(g^3 + r)/r^2 \le C''/r$ on the half-size ball. Integration along the segment from zero to K therefore gives

$$\|D_K \Psi_j(V,K) - D_K \Psi_j(V,0)\| \le C''\|K\|/r. \tag{3.38}$$

A strict derivative bound at zero persists on a sufficiently small concentric ball. The next two subsections establish the connected size count and the estimates on its factors.

3.6 The connected size sum on cube labels

The next lemma sums the configurations contributing to one connected coarse polymer. Its input is a bound for each product after the field regulators have been estimated. This distinction matters when two factors use overlapping field stencils.

Lemma 3.9 (Size estimate for a connected output)**.** *Write $P(X) = (|X| - N_s)_+$. Consider configurations of coarse closure U consisting of m disjoint connected fine inputs X_i, q local extraction covers and r ordinary decorations. The latter two types have bounded fine support, a bounded number of possible objects per fine centre and bounded multiplicity at a centre. Take N_s sufficiently large, in particular larger than the size of the coarse closure of any local cover, and then take L sufficiently large. Set $t = \sum_i |X_i| + q + r$. Suppose the cost of a configuration at each field radius, before the output mixed weight, is at most*

$$C_L^t a^{m+q} \epsilon^r \prod_{i=1}^{m} \varrho^{-P(X_i)}. \tag{3.39}$$

At the small radius take $(a,\epsilon) = (\tau, C_L g)$ and at the large radius take $(a,\epsilon) = (\tau g^{-9/4}, C_L g^{1/4})$. For each $A > 1$, choose ϱ sufficiently large and then g and $\tau g^{-9/4}$ sufficiently small. With the output size weight $(A\varrho)^{P(U)}$, the mixed norm of the sum with $m + q \ge 1$ is at most $C_{L,A,\varrho}\tau$. The corresponding sum with $m + q = 0$ and $r \ge 3$ is bounded by $C_{L,A,\varrho} g^3$. Both bounds are uniform in U and the volume.

Proof. For $m + q \ge 1$, the union bound for the supports and the definition of P give

$$P(U) \le \sum_i P(\overline{X}_i) + N_s(m + q - 1) + Cr, \tag{3.40}$$

where C bounds the closure size of a decoration. For a connected fine input X, a traversal of a spanning tree visits all its labels in at most $2|X|$ steps. Each segment of at most L steps meets only a bounded number of coarse tiles, since its diameter is $O(Lb)$. Thus $|\overline{X}| \le C_d(1 + |X|/L)$. When

$|X| > 2N_s$, the choices of N_s and L give $P(X) - P(\overline{X}) \geq c|X|$. The resulting inverse power of ϱ pays for $A^{P(\overline{X})}$ and an arbitrarily large fixed exponential in $|X|$ when ϱ is sufficiently large. Each smaller input and each cover or decoration costs at most a fixed factor in the size weights.

To count configurations with fixed U and t, there are at most $C_L|U|$ fine positions for occupied sites and centres, counting the allowed multiplicities. Since $|U| \leq Ct$, choosing these positions costs at most $C_L'^{\,t}$. Choosing an edge subset of the induced bounded-degree graph specifies a partition into connected inputs, at an additional cost exponential in the number of occupied labels. The choices of objects at their centres satisfy the same bound. Part of the exponential gain on large inputs pays these costs; a further part sums their sizes. The remaining factors are bounded by a fixed constant per input, cover or decoration. Keep one factor a when $m + q \geq 1$; summation of the remaining factors is geometric for small a, ϵ. At the large radius the output multiplier $g^{9/4}$ converts the retained a to τ.

For the pure configurations, $|U| \leq Cr$, so the count and the weight cost at most a fixed exponential in r. Keeping three factors of ϵ gives $O(g^3)$ at the small radius and $O(g^{3/4})$ at the large radius. The large-radius multiplier $g^{9/4}$ completes the estimate. □

The pure terms with one or two decorations still require cancellation on their assigned labels. Section 3.7 proves the product estimate and performs this subtraction.

3.7 Joint bounds for the reblocked reference expansion

There are three steps in estimating the connected output. Lemma 3.10 bounds products under the Gaussian integral. An extraction cover may enlarge the connected support, but only the originally occupied labels lose their primary reference. This leaves quartic damping on all other labels, even when two factors depend on the same field values. Proposition 3.11 uses this bound after the first two orders of the reference have been subtracted. Proposition 3.12 adds the terms containing an activity and estimates the resulting derivative.

Lemma 3.10 (Joint regulator for reblocking)**.** *Use the adapted quotient and recovery filter of Section 3. Suppose $0 < p < 1$, $2p > s - 1$, and $g_B \asymp g_b = g$. Let the Gaussian fluctuation have range $O(B)$ and, on a fixed enlargement of each fine tile P, satisfy*

$$Y_P = F_{b,P^*}(\zeta)/\ell_b, \qquad \sup_P \mathbb{E} \exp(aY_P^2/\sigma) \leq C_0 \tag{3.41}$$

for a fixed $a > 0$. The enlargement may contain all the stencils of a finite menu of local insertions at P. Let U be a set of coarse labels, each the disjoint union of its fine labels. On pairwise disjoint originally occupied fine polymers use their large-radius activity bounds. On the other fine labels use the stable quartic with real row coefficient at least cg. Individual factors may be evaluated at $\phi + t\zeta$, $0 \leq t \leq 1$, with different parameters on different factors. The quartic exponent may also be a rowwise convex combination of two such quartic exponents. Allow quadratic exponent perturbations of small-radius jet size Cg per label, and a bounded number of polynomial factors per fine label. Assume each polynomial insertion, after factoring out its stated coefficient cost, has fixed degree in the normalized coarse field and in the Y_P on its stencil. The same bounds are to hold through the fixed tested field-jet order of the mixed norm.

Choose L large, then the recovery and strip errors small, and then the upper bounds on σ and g small. The expected product in the coarse large-radius jets is bounded by the product of the component and insertion costs times

$$C_L^{|U|} \exp\left[-2\kappa \sum_{C \in U} Z_C^2 + \tfrac{1}{2} c_h \sum_{C \in U} Q_C^2\right], \qquad Z_C = \frac{\|\phi\|_{2,C}}{B^{d/2} h_B}, \quad Q_C = \frac{q_{C^*}(\phi)}{h_B}. \tag{3.42}$$

The analogous zero-background bound at ℓ_B has prefactor $C_L^{|U|}$ and uses the fine small-radius activity bounds and insertion costs. A factor Y_P^k in an insertion gives the additional cost $C_{L,k}\sigma^{k/2}$ when retained in the bound. The moment and reserve hold uniformly over the segment parameters.

Proof. Let A_C denote the normalized L^2 norm of the background on the occupied sites of C, and S_C the normalized L^2 norm of the unoccupied filtered rows in their shifted backgrounds. For a rowwise convex combination apply Jensen's inequality to the real quartics, using the lower bound on their coefficients, so that the resulting effective shift is again in $[0,1]$. The local recovery estimate in Section 3 gives, with $\delta = C(R/b + \|I - B_{\mathrm{rec}}T\|_{2\to 2}^2)$,

$$(1-\delta)Z_C^2 \le A_C^2 + CS_C^2 + \delta Q_C^2 + C\|T\zeta\|_{2,C}^2/(B^d h_B^2). \tag{3.43}$$

Indeed, the exceptional strips around all fine label boundaries have relative volume at most CR/b, even when the occupied set is disconnected. The error of the inverse filter is estimated on a fixed enlargement of C, by its operator norm and $F_{B,C^*}(\phi)/h_B \le C(Z_C + Q_C)$. The fluctuation term, summed on the coarse tiles, is bounded by $C_L \sum_P Y_P^2$ using the bounded filter row norm.

On occupied labels, $\|\phi + t\zeta\|_{2,P}^2 \ge \|\phi\|_{2,P}^2/2 - \|\zeta\|_{2,P}^2$. Their damping supplies a negative coefficient of order κL^s on $\sum_C A_C^2$. The triangle inequality for each adapted quotient and the summed restriction estimate give the background quotient cost $Cc_h L^{s-1-2p}\sum_C Q_C^2$. The remaining occupied fluctuation costs are bounded by $C_L \sum_P Y_P^2$. On each coarse tile the stable quartic is bounded below by cS_C^4, by Cauchy–Schwarz and $gB^d h_B^4 \asymp 1$. It supplies any prescribed negative multiple of S_C^2 with a constant cost per coarse tile. Take the coefficients on A_C^2 and S_C^2 large enough to get strictly more than 2κ on Z_C^2. By the choice of L and then δ, the positive coefficient on Q_C^2 is strictly below $c_h/2$. A complex quadratic term of normalized size Cg uses at most $C_L\sqrt{g}\sum_C(Z_C^2 + Q_C^2) + C_L\sum_P Y_P^2$ in the exponent. Its background part fits in the strict reserves for small g.

The fine radius of a restricted coarse unit direction is bounded by C_L, at both radii. By the product rule and bounded degrees, the polynomial and field-jet factors, after taking out their costs, are bounded by a polynomial in Z_C, Q_C, Y_P per fine label, of degree bounded uniformly per label. There are at most C_L fine labels in a coarse tile and the enlargement graph has bounded degree. Thus the background polynomials are absorbed by the remaining quadratic reserves, at fixed cost per coarse tile. Equivalently, one can first use the submultiplicative finite-jet seminorm with factorial weights, and then absorb each local polynomial. An activity on a large polymer uses its tested seminorm on that entire polymer. It contributes no extra degree per site, and the direction restriction bound still holds.

Color the fine stencil graph by grid residue classes with spacing larger than the covariance range, increased by twice the stencil radius. Within each of the $\chi = O_L(1)$ colors the Gaussian vectors are independent. For bounded k_P, Hölder and the single-tile moment give

$$\mathbb{E}\prod_P Y_P^{k_P} e^{C_L Y_P^2} \le C_L^{\#\{P\}}\sigma^{\sum_P k_P/2}. \tag{3.44}$$

In fact factor out the powers of $\sqrt{\sigma}$, absorb the rescaled polynomial by $\exp(aY_P^2/(2\chi\sigma))$, and use $\chi C_L \le a/(2\sigma)$. The number of fine labels in the Gaussian bound is at most $C_L|U|$, proving the asserted prefactor. At zero background the stable quartic has nonnegative real part, and all derivatives, quadratic terms and growing activity regulators have the same Gaussian majorant. This proves the small-radius bound. The bounds also justify Gaussian integration of the tested derivatives. □

Centered differences and the pure reference. After centering, a first-order quartic difference contains at most three background fields. At the large radius this degree bound leaves enough powers of the coupling for a cubic remainder. The subtraction uses the exact assignment to anchors in Lemma 3.4.

Proposition 3.11 (Cubic residual for a bounded anchored split)**.** *Normalize the incoming vacuum. Write its reference as a product over fine labels with factors $I_i = e^{-U_i}(1 + W_i)$, where U_i is the stable quartic on its rows and a quadratic term. Assume $\|U_i\|_{\ell_b,0} \le Cg$ and $\|W_i\|_{\ell_b,0} \le Cg^2$, with $\deg W_i \le 6$. A bounded product of such polynomial factors is also allowed with W_i their sum to first order in their coefficients. Put*

$$P_i = \mathbb{E}\theta U_i, \qquad R_i = \mathbb{E}\theta W_i + \sum_k E^{\Gamma}_{ik}, \qquad J_i = e^{-P_i+D_i}(1 + R_i - D_i), \tag{3.45}$$

where θ is Gaussian translation and E^{Γ}_{ik} is the oriented row of Lemma 3.4. Suppose D_i has degree at most four, size $C_L g^2$ and quartic part proportional to the primary rows on i. Constants can equivalently be removed on the labels. Use the joint bound of Lemma 3.10. Then on zero independent activity the reblocking expansion about J_i, followed by the zero-sum covariance reapportioning, has connected remainder $O_{L,\varrho}(g^3)$ in the mixed norm, with a strict coarse field reserve. The bound holds on the enlarged complex coefficient domain whenever the displayed size and sector assumptions hold there.

Proof. Use $\xi_i = \theta U_i - P_i$ and put $\gamma = g^{1/4}$ for the large-radius estimate. In normalized fine small-radius jets ξ_i has cost $C_L g$. In large-radius jets it has cost $C_L\gamma$ under the product bound. In fact

$$\xi_i(\phi, \zeta) = [U_i(\phi + \zeta) - U_i(\phi)] - \mathbb{E}[U_i(\phi + \zeta) - U_i(\phi)]. \tag{3.46}$$

Each monomial of the difference has at least one fluctuation direction measured at ℓ_b, leaving at most three powers of the large background or tested directions at h_b. The coefficients of the centered expectation follow by the same Gaussian moments. The translated W_i and all R_i, D_i have small-radius costs $C_L g^2$ and large-radius costs at most $C_L\gamma^2$: their degree is at most six, including for the covariance of two quartics, since every surviving term has a cross pairing. The constants are uniform in scale because the number of fine labels interacting with an anchor is $O_L(1)$.

Taylor's integral formula for the exponential, used along the segment between P_i and θU_i, gives

$$\theta I_i = e^{-P_i}(1 - \xi_i + \tfrac{1}{2}\xi_i^2 + \theta W_i) + \mathcal{E}_i, \qquad J_i = e^{-P_i}(1 + R_i) + \mathcal{F}_i. \tag{3.47}$$

The errors have costs $C_L g^3, C_L\gamma^3$ at the two radii. For the first identity the errors contain either three ξ_i, one ξ_i and one W_i, or two polynomial factors of order g^2. For the second, use the segment between P_i and $P_i - D_i$ and keep two factors of order g^2. The exponent remains stable on each segment, and Lemma 3.10 applies after multiplying by the other factors. In particular, a difference $\Delta_i = \theta I_i - J_i$ has the leading term $-e^{-P_i}\xi_i$ and a residual of costs $C_L g^2, C_L\gamma^2$. These conclusions also follow directly from the first-order integral formula when no cancellation is being used.

For one difference, its expected coefficient relative to the common stable factor, through order two, is

$$\tfrac{1}{2}\mathbb{E}\xi_i^2 + \mathbb{E}\theta W_i - R_i = -\sum_{k \ne i} E^{\Gamma}_{ik}. \tag{3.48}$$

For two differences it is $\mathbb{E}\xi_i\xi_k = E^{\Gamma}_{ik} + E^{\Gamma}_{ki}$. The error in the latter identity for the actual product keeps one first-order factor and one second-order factor or two second-order factors. In either calculation the deterministic perturbative fillers may be replaced by their common stable parts for

the displayed second-order terms. A filler difference has second-order cost, so the resulting error still has at least cubic order.

For a coarse cell C, the pairs internal to C cancel the internal singleton row contributions. Between two coarse cells C, D assign to the ordered cover $(C, \{C, D\})$ the sum $\sum_{i\in C, k\in D} E^{\Gamma}_{ik}$, and to its singleton the negative sum of those rows. The row has zero sum for each C. Using the circle identity for that row cancels the displayed second-order coefficient on each connected output polymer; the sign is selected to subtract that coefficient. The replacement of the stable filler by the perturbative reference adds a product of two second-order factors. Multiplicative factors commute on labels, so this subtraction uses the original assignments even on stencils shared by two labels.

All terms with at least three fine differences have at least three factors of g at small radius and γ at large radius. The product bound and the pure connected count of Lemma 3.9 sum them with exponential size reserve. The one- and two-difference errors are on bounded-size coarse polymers with $O_L(1)$ choices, hence have the same estimates. In the exact zero-sum reapportioning all further terms have either two inserted covariance rows or a row and an already estimated residual. The same-scale connected count pays the finite menu of row covers: they have uniformly bounded size and the extra g^2, respectively γ^2, pays their size weight. Existing residuals use a fixed part of their stronger exponential weight for the partition count. At common scale, products of the residual bounds on disjoint occupied labels concatenate their regulators, and the stable reference on uncovered labels supplies any fixed quadratic damping, by the local recovery with $b = B$ and the small strip error. The spare field reserve pays for the polynomial insertions, of bounded multiplicity per label. Including the output weight $g^{9/4}$ in the large-radius bound gives the asserted g^3. □

Conversion to the outgoing reference. A local comparison follows directly once an outgoing polynomial representative is available on the same anchors. Suppose the outgoing exponential is the literal sum of $P_i - D_i$ on a coarse label, after recording the vacuum. On a boundary label this formula can be implemented by retaining the full quadratic in the inner collar and using the common scalar quartic coefficient on its actual filter rows. Assign the difference of that coefficient from the incoming one to D_i and leave the remaining quartic and sextic terms in $R_i - D_i$. This is admissible when the coefficient difference is $O_L(g^2)$. The sum of the remaining polynomials can be stored as a single factor per collar, while factors on the ordinary margin use their original anchored formulas.

In multiplying the fine $1 + (R_i - D_i)$ factors, the linear terms sum to the retained polynomial. Every discarded product contains at least two degree-at-most-six factors of size $O_L(g^2)$, hence has costs $C_L g^4$ and $C_L g$ at the two radii under the reserve. If a bounded bilinear ordinary reference is recomputed after an $O_L(g^2)$ change of its potential, the coefficient difference is $O_L(g^3)$ of degree at most six. Its large-radius cost is $C_L g^{3/2}$. Thus the literal cellwise reference error is $O_L(g^3)$ in the mixed norm. The estimate extends to moved inputs when their reference sizes and the anchor equality hold on the enlarged extraction disc. It uses neither a pointwise inverse of a reference polynomial nor cancellation between different coarse labels.

The localization–Gaussian commutator. On a single mark, the linearized extraction contains a difference between applying the local projection before and after Gaussian integration. With ν its mixed norm, the Gaussian shift estimate on each zero-field derivative of order $k \le 6$ is

$$\left|D^k(\mathbb{E}F_X(\zeta + \cdot) - F_X)(0)\right|_{\ell_b} \le C\sqrt{\sigma}\,\nu. \tag{3.49}$$

Here and below the norm denotes the supremum over normalized test directions. Integrate the $(k+1)$-st derivative along $t\zeta$, use the fine growing regulator, and retain the factor Y_P in the Gaussian

bound on the bounded collection of mark stencils. The same shift estimate holds for the polynomial $\Pi_X T_8 F_X$, by its bounded coefficients. The projection uses only degrees through six, so

$$\|\Pi_X T_8 \mathbb{E} F_X(\zeta + \cdot) - \mathbb{E}(\Pi_X T_8 F_X)(\zeta + \cdot)\|_{\ell_b,0} \le C\sqrt{\sigma}\,\nu. \tag{3.50}$$

The resulting polynomial has degree at most six. After insertion in the complementary reference, passage to the large radius costs at most $g^{-3/2}$, which is paid by the mixed weight. Summing over the $O(L^d)$ small-mark placements gives $C_L\sqrt{\sigma}\,\nu$ in the coarse mixed norm. This estimate includes the merged step whenever its raw fluctuation amplitude obeys the stated moment assumption.

We now combine the commutator estimate with the connected count. The hypotheses below specify the comparison of reference factors on a single coefficient domain. The resulting derivative estimate is the one used in the shooting argument.

Proposition 3.12 (Extraction disc with a bounded outgoing split)**.** *Suppose the fine reference and the moved reference on $\tau = \|K\| \le cg^{9/4}$ have the bounds of Proposition 3.11 on an enlarged complex coefficient domain. Suppose their outgoing splits satisfy the literal matching and the bounded bilinear ordinary-reference comparison above on that domain. If extraction affects the ordinary reference through such a bilinear factor, use its difference bound $C_L g\tau$. Include among the boundary labels a fixed halo containing the nonbulk supports. Use the same nested grid and definitions of ordinary data in the finite-volume and bulk maps.*

For small c the connected output $\Psi(V,K)$ of the extraction, reblocking, covariance redistribution and outgoing reference change is norm-holomorphic and satisfies

$$\|\Psi(V,K)\| \le C_{L,\varrho}(g^3 + \|K\|). \tag{3.51}$$

At zero activity its derivative differs from the leading localized small-mark operator by $C_{L,\varrho}(g^{1/4} + \sqrt{\sigma}) + \varepsilon_{\varrho,L}$, where the large-set term $\varepsilon_{\varrho,L}$ tends to zero as $\varrho \to \infty$ for fixed sufficiently large L. The extraction change of each full boundary tensor is $O_L(\tau)$. The extraction change of each ordinary margin polynomial is $O_L(g\tau)$. Their derivatives on smaller domains obey the same bounds in a unit normalized potential direction and, with τ replaced by $\|\delta K\|$, in an activity direction δK.

Proof. On each originally uncovered fine label i, insert the constant $1 - \sum_{X \ni i} J_i(X) = 1$ from the zero-sum extraction, with $J_i(X)$ assigned to the cover X. For this formula only, denote the perturbative filler by J_i^{pt} and the moved incoming reference by $\widehat{I}_i$. After Gaussian translation the resulting uncovered factor is exactly

$$\theta I_i\Big(1 - \sum_{X \ni i} \theta J_i(X)\Big) = J_i^{\mathrm{pt}} + (\theta\widehat{I}_i - J_i^{\mathrm{pt}}) + \theta(I_i - \widehat{I}_i) - \sum_{X \ni i}(\theta I_i)\theta J_i(X). \tag{3.52}$$

In each selection group the covers and the originally occupied polymers into their coarse connected components. Inside a selected cover away from its selected centre and the occupied labels, the first three terms sum to θI_i; summing them first implements the original reference fill on that cover. Outside the selected covers, the same three choices give the ordinary multiplicative reference change and reblocking. Gaussian expectation factors on separated coarse components by the range condition. Thus the expansion has the exact circle product of the original input. Its occupied-label damping comes from the original activities, and every other fine label keeps one stable exponential, possibly evaluated along a segment.

The local extraction costs are $C_L\tau$ and $C_L\tau g^{-3/2}$. For the cover terms this is the degree-six bound on the extracted jet. For the reference differences use a coefficient segment with a stable primary quartic and one differentiated polynomial or exponent, of respective degrees at most six and

four. The ordinary differences have costs $C_L g$ and $C_L g^{1/4}$, by the centered expansion. All insertions have bounded multiplicity on a fine stencil because their centres do and their assigned sizes are bounded. Lemma 3.10 proves the joint product estimate required by Lemma 3.9. At the large radius each fine component has coefficient cost at most $\tau g^{-9/4}$ times its inverse size weight, and each extraction mark costs an additional $g^{3/4}$ relative to this bound. The prefactor $C_L^{|U|}$ is allowed since the number of coarse labels is at most a fixed multiple of the total occupied size and number of decorations. This gives $C_{L,\varrho}\tau$ for all connected terms with an activity or extraction mark. The pure terms at each fixed extracted reference are handled by Proposition 3.11. At common scale the covariance and outgoing-reference changes use only bounded-size covers and retain the same strict field reserve as in that proposition, with a fixed part of the stronger size reserve. Their pure reference errors have mixed size $O_L(g^3)$. This gives the claimed bound.

For the derivative, hold the extracted reference fixed while assigning each term to its independent marked component. The derivative through the pure remainder at the extracted reference is $C_{L,\varrho}g^{3/4}$, by Cauchy's formula on the disc of size $cg^{9/4}$ and its pure $O(g^3)$ bound. A linear term with an ordinary decoration retains $g^{1/4}$ at the large radius and a factor g in its small-radius jets. Later same-scale decorations have second-order cost or smaller. A single large mark without decorations has the large-set size gap and gives the stated $\varepsilon_{\varrho,L}$.

The remaining first variation is a single small mark. At the fine scale the zero-sum extraction and the singleton change have the linear combination $F_X - I^X q_X$, where $q_X = \Pi_X(F_X/I^X)$ in the finite jet algebra. Indeed, the nonsingleton row is $-I^X \sum_{i\in X} q_{X,i}$. On a singleton i the negative singleton row and the reference change combine to $-I_i q_{\{i\},i}$. The remainder from the literal reference change has derivative $O_L(g)$ in the mixed norm, by Lemma 3.8. In the small-radius jets, replacing $I^X q_X$ by $\Pi_X T_8 F_X$ costs $O_L(g\|F_X\|)$ since the normalized reference is $1 + O(g)$. At the large radius both inserted polynomials cost at most $C_L g^{-3/2}\|F_X\|$ before the mixed weight, with the reserve supplied by the small-mark complement. This gives an error bounded by $C_L g^{3/4}\|F_X\|$ there. The localization–Gaussian commutator then replaces the incoming projection by the projection after integration at cost $C_L\sqrt{\sigma}\|F_X\|$. The surviving operator is the one in Proposition 3.3, with the unoccupied perturbative factors as complement. Its local quadratic and polynomial perturbations satisfy the sizes in that proposition. The small-mark argument needs a quartic-only complement, so all perturbations are absorbed after L is fixed and g is small.

For the retained differences, the jet inverse gives $C\tau$ per extracted fine coefficient. Gaussian heat on a bounded collection of stencils is a bounded linear operator on their finite jet norms. The covariance correction is bilinear, so its change is bounded by $C_L(g\tau + \tau^2)$. The complete degree-two, four and six differences in the inner collar are finite sums of these coefficients. A margin has the ordinary bilinear polynomial, and its outgoing ordinary potential varies by at most $C_L\tau$ under the bounded split. This gives the additional g there. A boundary variation does not change an ordinary factor or an ordinary margin read outside its support and propagation neighborhood. On an enlarged domain all these differences are holomorphic, including their vacuum normalization. A one-variable Cauchy integral in each normalized potential direction preserves the factor τ. On a smaller activity ball a Cauchy integral uses numerator $O_L(g^{9/4})$ and denominator of order $g^{9/4}$, preserving the extra g for margin differences.

The finite algebraic maps are analytic on their domains. Gaussian integration admits the uniform reserve majorant on every smaller complex ball. Cauchy coefficients of the integrated field jets are estimated using the same mixed norm bounds, so their series converge in the output norm. The connected counts give uniform bounds in the volume. This proves norm holomorphy and justifies the activity derivatives used above. □

Iteration also requires the retained coefficients to lie in the allowed domain at the next scale.

For $d = 1, 2, 3$, Proposition 3.15 checks that transport. Its first input is a uniform bound for the quadratic and quartic contractions on an ordinary stencil.

The ordinary second-order polynomial for $d = 1, 2, 3$. On an ordinary stencil all coefficients are translation invariant and the covariance and filter agree with the bulk ones. This makes it possible to bound the cumulative reference with constants independent of the scale. In the three-line quadratic contraction the relevant spatial cancellation is evenness in the internal displacement. A first-order absolute difference would be insufficient for this purpose in dimension three.

Lemma 3.13 (Bilinear ordinary reference bound below $s = 2$)**.** *Let $d = 2s$, $s < 2$, and use the normalized, even bulk finite filter T with row sum one. Let S be a cumulative bulk covariance supported at distance $O_L(b)$, with, for the orders used below,*

$$|\nabla^k S(x)| \le C_{L,k}\lambda(R + |x|)^{-s-k}. \tag{3.53}$$

Use the exponentially weighted absolute row bounds of T at length R. Write $u_x = (T\phi)_x$, and take the ordinary potential on anchor x to be a linear combination of

$$u_x^4, \qquad b^{-s}\phi_x^2, \qquad b^{2-s}k_x(\phi), \qquad b^{2-s}R^{-2}(u_x^2 - \phi_x^2), \qquad k_x(\phi) = \tfrac{1}{4}\sum_{|e|_1=1}(\phi_{x+e} - \phi_x)^2. \tag{3.54}$$

Write G for the sum of the absolute coefficients, with the coefficient of the last density weighted by $1 + \log(2 + b/R)$. In the oriented covariance row of Lemma 3.4, put the third density in its second type and the other three in the first type. Define Loc *to keep the vacuum, the constant-field quartic, and the mass and kinetic moments of the bulk quadratic, represented on the same anchors by the displayed vacuum, filtered quartic, mass and nearest-neighbour kinetic densities. Then*

$$\big\|(1 - \mathrm{Loc})X_S(U)\big\|_{\ell_b,0;P} \le C_L(\lambda + \lambda^2 + \lambda^3)G^2 \tag{3.55}$$

on each ordinary tile P. The bound is bilinear, including first variations with this strengthened coefficient norm. The remaining polynomial has degree at most six.

Proof. Set $P_1 = TS$ and $D = TST^*$. They obey the same bounds and range as S, by the absolute exponential filter moments. The shell sums give

$$\|S\|_1 + \|P_1\|_1 + \|D\|_1 \le C_L\lambda b^s, \qquad \|\nabla P_1\|_2^2 \le C_L\lambda^2/R^2. \tag{3.56}$$

When ϕ is tested on this stencil, a spatial difference of order k costs $C_L b^{-k}\ell_b$. An even filter of mass one thus has $(T - I)\phi = O_L((R/b)^2\ell_b)$, also after any fixed tested derivative of order k, with another factor b^{-k}. This follows by subtracting the constant and linear terms before using the second absolute moment of the filter. Convolutions commute with these differences in the ordinary chart.

The Gaussian Wick product identity for P_{-S} in the oriented row leaves precisely the contractions with at least one cross pairing between the two vertices. It follows, for instance, by applying P_S to two separately Wick-ordered polynomial generating exponentials and expanding the cross exponential. For two quartics the one-cross term has norm $C_L\lambda G^2$, using $b^d\ell_b^6 = b^{-s}$. In the two-cross term, subtract the constant-field representative u_x^4 at the anchor x. The remaining external-leg factor is bounded by $C_L\ell_b^4 \min(1, |r|/b)$ in multilinear tests. Summing over the internal displacement r gives $C_L\lambda^2$ after the factor $b^d\ell_b^4 = 1$:

$$\sum_r |D(r)|^2 \min(1, |r|/b) \le C_L\lambda^2. \tag{3.57}$$

The three-cross term at x contains $D(r)^3 u_x u_{x+r}$. Replace u_{x+r} in the even row by the average of u_{x+r} and u_{x-r}. On subtracting u_x^2, the external-leg bound is now $C_L \ell_b^2 |r/b|^2$. Since

$$\sum_r |D(r)|^3 |r|^2 \le C_L \lambda^3 b^{2-s} \qquad (s < 2), \tag{3.58}$$

multiplication by $b^d \ell_b^2 / b^2 = b^{s-2}$ gives $C_L \lambda^3$. For replacing the filtered on-site square by its mass representative, use $\sum_r |D(r)|^3 \le C_L \lambda^3 R^{-s}$ and the filter difference $(R/b)^2$. The normalized kinetic moment of the quadratic has the same bound, since it is bounded by

$$C_L b^{s-2} \sum_r |D(r)|^3 (|r|^2 + R^2) \le C_L \lambda^3. \tag{3.59}$$

The R^2 term accounts for the external filter legs. The symbol computation of these moments uses translation invariance and cubic symmetry. Their nearest-neighbour representatives have the same local norm bound.

For clarity, let Q denote the Hessian of any sum of the displayed quadratic densities with its stated coefficients. The mass kernel is of size Gb^{-s}, and the other kernels are even zero-row kernels with second absolute moment at most CGb^{2-s}. Hence $Q\phi$ costs $C_L G b^{-s} \ell_b$ on the ordinary stencils, by subtracting the constant and linear terms. Its fixed differences of order k satisfy the same estimate times $C_{L,k} b^{-k}$. For a one-cross quartic–quadratic contraction rooted at the quartic, the contracted direction is $TSQ\phi$, bounded by $C_L \lambda G \ell_b$. At a mass root use P_1 on the smooth external cubic from the quartic. At a structured head root the expression with that smooth field v has the form $u_x (TSv)_x - \phi_x (Sv)_x$. Its decomposition

$$(u_x - \phi_x)(TSv)_x + \phi_x ((T - I)Sv)_x \tag{3.60}$$

recovers the factor $(R/b)^2$ on the head coefficient. The difference-square root in a mixed quartic–kinetic contraction does not occur, by orientation. These bounds, with $b^d \ell_b^4 = 1$, control the quartic and its constant-field projection.

For the two-cross quartic–mass contraction the row is a multiple of $P_1(r)^2$ times the filtered square at the quartic anchor. When the root is the mass anchor, moving that square to its root and replacing it by the mass density costs $C_L \ell_b^2 (\min(1, |r|/b) + (R/b)^2)$. The first factor eliminates the marginal logarithm of the absolute row, and the second times $1 + \log(2 + b/R)$ is bounded. The difference-square root is again avoided by orientation in the kinetic contraction. At the quartic root the row is exactly the filtered square there times a coefficient bounded by $C_L G^2 b^{2-s} \sum_r |\nabla P_1(r)|^2$. Subtracting the mass and kinetic representatives costs $C_L (R/b)^2$ in its normalized coefficient bound, which cancels the factors from the kinetic input and the derivative row. These statements also control the kinetic subtraction: with both surviving legs at a quartic row, the relative second moment of the external kernel is $O(R^2)$ times its scalar row coefficient, independently of the displacement of that row from the selected anchor.

For the head–quartic two-cross term put $F(r) = D(r)^2 - P_1(r)^2$. Even filtering and the second difference estimate give

$$|D(r) - P_1(r)| \le C_L \lambda R^2 (R + |r|)^{-s-2}. \tag{3.61}$$

One obtains this by using the second difference for shifts of size at most $(R + |r|)/2$ and the exponentially weighted absolute row bound on the other shifts. Thus $\sum |F| \le C_L \lambda^2$ and $\sum |r|^2 |F(r)| \le C_L \lambda^2 R^2 [1 + \log(2 + b/R)]$. A head root requires moving the square to its root. Pairing r with $-r$ makes this cost $C_L \ell_b^2 |r/b|^2$. The filter and kinetic comparison cost $C_L \ell_b^2 (R/b)^2$ against the unweighted row. The resulting normalized bounds use precisely the strengthened head coefficient. At a quartic root there is no need to move the square.

For the quadratic–quadratic one-cross term the summed partner contributes $SQ\phi$, with test norm $C_L\lambda G\ell_b$. A mass root uses its coefficient Gb^{-s}. At a difference-square root the external field and $SQ\phi$ each supply a difference at scale b^{-1}. At a head root use the preceding $(T-I)$ decomposition with $v = Q\phi$. All three cases thus give $C_L\lambda G^2$ after multiplication by $b^d\ell_b^2$. The mass and the relative second moment of their spatial sum are bounded by differentiating $\widehat{Q}_1\widehat{S}\widehat{Q}_2$ at zero, using the row and second moment estimates for S, Q_1, Q_2. There is only a vacuum from all higher cross orders of these contractions. This exhausts the retained degrees and proves the stated bound, with bilinearity giving the same estimates on polarized coefficients. □

For this ordinary coefficient space, Lemma 3.4 gives the literal heat split on each anchor. If X_i^+ is the anchored $X_{S+\Gamma,i}(P_\Gamma U)$, set

$$D_i = (\mathrm{Loc}X^+)_i - P_\Gamma(\mathrm{Loc}X_S(U))_i. \tag{3.62}$$

Then $R_i - D_i = (1-\mathrm{Loc})X_i^+$ for the incoming polynomial $W_i = (1-\mathrm{Loc})X_{S,i}(U)$. The heat of the quartic contributes the fixed filtered square, which is the sum of the chosen mass and head densities on that anchor. The heat of a quadratic contributes only a vacuum. These identities make the split use the same anchored mass, quartic, kinetic and head densities on the two scales. In particular $D_i = O_L(g^2)$ wherever the strengthened coefficients of $P_\Gamma U$ are $O_L(g)$, since R_i has that bound and the outgoing polynomial is bounded by the lemma. If the literal outgoing potential differs from $P_\Gamma U$ by $O_L(g^2)$ in the strengthened norm, recomputation of its ordinary polynomial costs $O_L(g^3)$, with the corresponding bilinear variation bound. This proves the ordinary comparison for changes within this coefficient space. Additional localized divergences or changes of boundary representatives must still be accounted for by their actual allocations.

Exact removal of extracted quadratic divergences. Ordinary localization produces second-moment densities that are lattice divergences. To remove them without changing the linearized activity map, choose the ordinary kinetic representative proportional to

$$k_x^{\mathrm{c}}(\phi) = \tfrac{1}{4}\sum_i(\phi_{x+e_i} - \phi_{x-e_i})^2. \tag{3.63}$$

The preceding proof applies to this representative as well. Its difference-square span is two, and its bulk quadratic symbol has leading Taylor term $|k|^2$, up to the fixed normalization of the coefficient.

In a translation-invariant cubically symmetric ordinary region, the extracted coefficients of $J_{P,i}$ and all off-diagonal $Z_{P,ij}, H_{P,ij}$ vanish. The diagonal Z and H coefficients are independent of i. Their contribution can therefore be written as a multiple of the central kinetic density and a multiple of $\sum_i(Z_{P,ii} + H_{P,ii})$. The latter can be carried by the boundary collar. In fact the commutator identity in Proposition 3.2 writes its sum over ordinary sites as a sum of local interface tensors. Each interface tensor has normalized fine coefficient at most $C\tau$, by the interface-bond count in that proof. The same reasoning applies to a first-moment divergence on any symmetrically assigned exceptional region. Keep the nonuniform extracted coefficients in the fine halo literally.

To implement this change at the level of all circle products, write the literal extracted and the new reference factors as $I_i = e^{-\Delta_i}\widehat{I}_i$, with the same polynomial factors, where Δ_i is the quadratic density removed on the ordinary labels and its negative is inserted on the halo as given by the commutator identity. Then $\sum_i \Delta_i = 0$ as a polynomial. For every occupied set with components X,

$$\prod_{i\ \text{uncovered}} I_i \prod_X K(X) = \prod_{i\ \text{uncovered}} \widehat{I}_i \prod_X \left(e^{\sum_{i\in X}\Delta_i} K(X)\right). \tag{3.64}$$

This holds on the actual fine labels with their stencils, without dropping local divergences from any single label. On each label the newly introduced quadratic has small-radius size $C\tau$. In the fused reblocking expansion it uses the same quadratic reserve as in Lemma 3.10, including when the label belongs to an originally occupied component. In particular this step need not map an unrestricted fine activity into the identical fine regulator before Gaussian integration.

Write the correction on each occupied component by integrating its exponential difference:

$$e^{\sum_{i\in X}\Delta_i}-1=\int_0^1 e^{t\sum_{i\in X}\Delta_i}\sum_{i\in X}\Delta_i\,dt. \tag{3.65}$$

There is one additional local insertion marked at i. Its costs are $C\tau$ and $C\tau g^{-1/2}$ at the two radii. The choice of i adds only a factor $|X|$, paid by the stronger size count on large X, or by a fixed constant on small X. The factor it multiplies already contains an independent activity or a local extraction mark. Thus this exact removal changes no first variation at $K=0$. Its connected bound is included in the enlarged-disc estimate. The ordinary bilinear polynomial can then be computed in the kinetic representative k^{c}, because the uniform divergence has been removed before that polynomial is formed.

3.8 The inverse bulk mass chart on the cube

The mass and kinetic terms in the second-order covariance polynomial can be included in the definitions of the bulk couplings. At the terminal scale we vary the resulting mass coordinate and need its derivative in the original local potential.

Let W_j^{cov} be the cumulative bulk covariance before step j. Define its normalized row sum and second moment by

$$a_j=b_j^{-s}\sum_x W_j^{\mathrm{cov}}(x),\qquad m_j^{(2)}=b_j^{-s-2}\sum_x x_1^2 W_j^{\mathrm{cov}}(x). \tag{3.66}$$

Summing the slice ℓ^1 bounds as a geometric series gives $|a_j|\le C\lambda$. Their range $O(b_j)$ gives $|m_j^{(2)}|\le C\lambda$. At the terminal scale the cumulative covariance excludes the remaining tail and satisfies the same bounds.

Lemma 3.14 (Compensated bulk coordinates and their inverse)**.** *Suppose the bulk covariance has cubic symmetry and the quadratic Taylor coefficient of the local kinetic symbol is $|k|^2$. Set*

$$\begin{aligned}\nu_j&=\nu_j^{\mathrm{loc}}+a_j(\nu_j^{\mathrm{loc}})^2,\\ g_j&=g_j^{\mathrm{loc}}(1+4a_j\nu_j^{\mathrm{loc}}),\\ z_j&=(1+2a_j\nu_j^{\mathrm{loc}})z_j^{\mathrm{loc}}-\tfrac12 m_j^{(2)}(\nu_j^{\mathrm{loc}})^2.\end{aligned} \tag{3.67}$$

These changes compensate the second-order cumulative mass–mass, mass–kinetic and mass–quartic projections in the uniform bulk coordinates. On the complex box $|\nu_j^{\mathrm{loc}}|\lesssim|g_j|$ with $\lambda|g_j|$ sufficiently small, the branch of the inverse through zero local mass is holomorphic. Holding g_j,z_j, the covariance and the separate tensor coefficients fixed gives

$$\partial_{\nu_j}\nu_j^{\mathrm{loc}}=1+O(\lambda|g_j|),\qquad |\partial_{\nu_j}g_j^{\mathrm{loc}}|\le C\lambda|g_j|,\qquad |\partial_{\nu_j}z_j^{\mathrm{loc}}|\le C\lambda(|\nu_j^{\mathrm{loc}}|+|z_j^{\mathrm{loc}}|). \tag{3.68}$$

Proof. For a quadratic Hessian Q the covariance polynomial has quadratic part $\frac12\langle Q\phi,W_j^{\mathrm{cov}}Q\phi\rangle$. Its constant and quadratic Fourier coefficients are those of

$$\big(b_j^{-s}\nu_j^{\mathrm{loc}}+b_j^{2-s}z_j^{\mathrm{loc}}|k|^2+O(|k|^4)\big)^2\widehat{W}_j^{\mathrm{cov}}(k). \tag{3.69}$$

Cubic symmetry gives the coefficient of $|k|^2$ in $\widehat{W}_j^{\mathrm{cov}}$ as $-\frac{1}{2}\sum_x x_1^2 W_j^{\mathrm{cov}}(x)$. In mass and kinetic units the two coefficients are therefore $a_j(\nu_j^{\mathrm{loc}})^2$ and $2a_j\nu_j^{\mathrm{loc}}z_j^{\mathrm{loc}} - \frac{1}{2}m_j^{(2)}(\nu_j^{\mathrm{loc}})^2$. The one-line mass–quartic contraction gives $4a_j\nu_j^{\mathrm{loc}}g_j^{\mathrm{loc}}$, as in Section 2.4. Heating the extracted uniform quadratic changes only its constant term. Thus the changes absorb the differences of the cumulative projections to second order. The separate zero-row terms retain their own polynomial increments.

Write $u = \nu_j^{\mathrm{loc}}$. The inverse through zero for $\nu_j = u + a_j u^2$ has derivative $\partial_{\nu_j} u = (1 + 2a_j u)^{-1}$. The remaining two equations are linear in g_j^{loc} and z_j^{loc}, respectively. Differentiating them gives

$$\partial_{\nu_j} g_j^{\mathrm{loc}} = -\frac{4a_j g_j^{\mathrm{loc}}}{1 + 4a_j u}\,\partial_{\nu_j} u, \qquad \partial_{\nu_j} z_j^{\mathrm{loc}} = \frac{m_j^{(2)} u - 2a_j z_j^{\mathrm{loc}}}{1 + 2a_j u}\,\partial_{\nu_j} u. \tag{3.70}$$

The denominators are uniformly bounded away from zero by the hypothesis on the mass box. The moment bounds then give the three derivative estimates. □

Thus the inverse mass variation reproduces the uniform mass to leading order, also in $d = 4$ with a small marginal kinetic coefficient. The finite-volume boundary tensors are assigned separately in the outgoing reference. Their bounds and those of the activity enter the terminal large-field comparison.

3.9 An invariant retained-coefficient chart in dimensions one through three

To apply the preceding activity estimate in $d = 1, 2, 3$, we now specify the outgoing coefficients on the cube. An ordinary site retains its bulk density even when it lies in the outer part of a cell meeting the boundary. In the inner collar the full nonuniform quadratic is retained. When a bulk difference-square density enters this collar, its contribution is thus measured in the boundary tensor norm. The same assignment enters the coefficient recursion and the change of reference in the activity map.

On an interval, a layer at a fixed distance from the boundary has at most two sites. Each coarse interval meets $O(1)$ fine intervals in the old collar and contains $O(L)$ fine intervals altogether. Take, for example, $p = 1/8$ at $s = 1/2$. The old quadratic collar tensor then has transport factor $CL^{-3/4}$ by Lemma 3.6, while the bulk kinetic coefficient transports with $L^{-3/2}$. A newly generated tadpole may occupy $O(Lb)$ sites. Its norm is bounded by $C|G|\lambda_j^{\mathrm{tad}} b^{-1/2}(Lb)\ell_{Lb}^2 = CL^{1/2}|G|\lambda_j^{\mathrm{tad}}$ by (3.31). On the next step it belongs to the old collar. The initial boundary head costs $C\lambda G_0 M_R^{-3/4}$ by Proposition 2.16. For the ordinary quadratic projection there is just one coordinate z_1, so the four moments are $A(1,1), A(1,z_1), A(z_1,z_1), A(1,z_1^2)$. The centered difference and its square give the same invertible moment matrix as in Proposition 3.2. Each interface contributes $O(1)$ crossing bonds to its commutator. These counts also apply to the terminal interval when a leftover piece is merged with the last full cell.

Proposition 3.15 (Retained coefficient estimates below the local endpoint)**.** *Let $d \in \{1, 2, 3\}$ and $\max\{0, (s-1)/2\} < p < s/2$, with $p \neq 1/2$. Use the frozen finite-range covariance, the initially truncated filter, the joint initialization, and the comparison sequence of Sections 2.3, 5.1 and 6. Suppose the activity agrees with its bulk comparison off the grouped boundary labels, and norm the finite-volume and bulk activities by the maximum of their mixed norms. There is a single outgoing chart with ordinary mass, quartic, local kinetic, and fixed-shape quadratic coefficients, full collar tensors of degrees two, four and six, and ordinary margin factors. In this chart the ordinary and collar reference comparisons required in Proposition 3.12 hold on its enlarged extraction disc. On the shooting box, the retained coefficient map obeys the quartic, mass and stable rows of Assumption 6.1,*

including their normalized first variations. Boundary comparisons preserve equality with the bulk off the new grouped boundary labels.

Proof. Use $b = b_j$, $B = Lb$. Away from the physical boundary use the site densities of Lemma 3.13, with the central kinetic representative in the subsequent divergence identity. Include also the fixed initial zero-row head H_0 of Proposition 2.16, in difference-square form. Use the notation $Z_T = T^*T - I$ for the structured head, with T normalized by its bulk row sum. Its anchored density is $(T\phi)_x^2 - \phi_x^2$. The actual filtered rows, and the uniform mass density, are shared by the ordinary and the collar factors. The other quadratic terms inside the collar are represented by a full quadratic $Q_{\partial,2}$. The additional terms of degrees four and six there are stored in $1 + B_{\partial,4} + B_{\partial,6}$.

The inner collar at scale B consists of the fine-layer anchors within a fixed multiple of B of the physical boundary. This accommodates the covariance range $O(B)$: in the scalar polynomial construction the largest time in an ordinary slice is at most a constant times the coarse length in units of R. Choose the collar width with a fixed margin over this range and the fine-scale localization and support enlargements. The incoming collar is correspondingly within a fixed multiple of b, so only $O(L^{d-1})$ old collar labels are reblocked in a coarse collar tile. New promotions may use all $O(L^d)$ fine labels of a coarse tile. They are sources from the incoming ordinary coefficients, with constants allowed to depend on L. The stencils used by every new inner factor lie within a fixed enlargement of its coarse cell and within depth $O(B)$. All remaining anchors in that cell use the ordinary factors, including their polynomial corrections. This set of ordinary anchors is the margin of the cell. Their polynomial corrections can be added on the cell, with the cost of the omitted products as in Section 3.7. At initialization the collar contains the finite boundary head and the bulk difference squares with anchors inside it. The ordinary second-order polynomial at that scale is zero. The cutoff defining which fine anchors enter the outgoing collar can be rounded to fine layers with no change in these bounds. For sides not divisible by the block length, start the grid at the same corner at every scale. On each coordinate axis merge the leftover interval with its last full interval and use the Cartesian products. The merged interval at the fine scale lies within the merged interval at the coarse scale, since the earlier coordinate of the coarse merged interval is a fine grid coordinate still to its left. This gives nested labels whose volumes and stencil enlargements are comparable to those of ordinary cubes.

We first check the ordinary comparison in this chart. The bulk Fourier estimate in Section 2.4 with an extra multiplier bounded by $|k|^r$ gives $\|\nabla^r \Gamma_i^{\mathrm{bulk}}\|_\infty \le C_r \lambda b_i^{-s-r}$ for each fixed order needed in the comparison. On the low-frequency ball the extra multiplier is handled by scaling $Rk/\sqrt{r_i}$, where $r_i = (R/b_i)^2$ is the spectral scale in that section. On the complement, the fixed gap gives the factor r_i^{2m-1} of the higher-resolvent estimate. The observation weight has integral with the extra factor $|k|^r$ bounded by $C_r R^{-d-r}$. Indeed the product formula for the spread-out box gives $|\widehat{D}_R(k)| \le C(1 + R|k|_\infty)^{-1}$; the excluded zero site contributes only $O(R^{-d})$. Choose the fixed smoothing power larger than $d + r$. The white channel contributes $O(e^{-R})$, which obeys the same bound. Increasing the fixed higher-resolvent order makes the large-frequency part no greater than the claimed slice bound. Finite propagation then gives the cumulative bounds in Lemma 3.13, by summing scales of size at least a constant multiple of the observed displacement. This includes the unmerged pieces of the first covariance.

The ordinary bilinear estimate allows H_0 in the second part of the orientation in Lemma 3.4. For this extension use the size

$$z_{H_0}(b) = b^{s-2} \sum_y |y|^2 |H_0(y)| \tag{3.71}$$

in the coefficient norm. At each difference-square root, the external and contracted legs contribute a difference over the same span y, so the product costs $C|y|^2/b^2$ in the ordinary field norm. Summing

with $|H_0(y)|$ gives exactly $z_{H_0}(b)$. For the mixed quartic–head term with two cross pairings, the orientation charges it to the quartic root. Both exterior legs there form the same on-site filtered square. The coefficient is bounded by

$$C \sum_y |y|^2 |H_0(y)| \, \|\nabla(TS)\|_2^2, \tag{3.72}$$

where S is the cumulative covariance. This follows by applying the two endpoint differences to the two contracted legs and using $\|f(\cdot + y) - f\|_2 \le C|y| \|\nabla f\|_2$. The subtraction of the on-site mass and kinetic moments uses $(R/b)^2$, as in Lemma 3.13. These are all changes involving the fixed head. They also apply to its variation by polarization.

Write G for the raw quartic, $u = b^{-s}\nu^{\mathrm{loc}}$ for the raw mass row, $S_j^0 = \sum_y S_j(y)$ for the row of the cumulative covariance, S_j for its kernel, and $\gamma_j = (T\Gamma_j T^*)(0)$ in the bulk. In particular $3B^s|\gamma_j| \le C_L \lambda_j^{\mathrm{tad}}$, where $\lambda_j^{\mathrm{tad}} \le C_L \Lambda_R$. The first heat gives the raw mass row $u_1 = u + 3\gamma_j G$ and the structured term with Hessian $3\gamma_j G Z_T$, on the same anchors. Use

$$D_i = (\mathrm{Loc}\, X_{S_j+\Gamma_j}(P_{\Gamma_j} U))_i - P_{\Gamma_j}(\mathrm{Loc}\, X_{S_j}(U))_i \tag{3.73}$$

on the outgoing ordinary anchors. This is the polynomial subtracted from the first-heat potential. The constant terms are recorded in the vacuum. The ordinary projection keeps scalar quartic, mass and kinetic coefficients. There is one further structured density in D_i because the localized quartic in $\mathrm{Loc}\, X_{S_j}(U)$ heats to a filtered square. If $a_{4,j}$ is the coefficient of that localized quartic in the normalization $G(T\phi)^4/4$, the extra contribution to the outgoing quadratic has Hessian $3\gamma_j a_{4,j}(Z_T - \kappa_T K_{\mathrm{loc}})$. Here K_{loc} is the kinetic representative with second Taylor coefficient $|k|^2$ and κ_T is that coefficient for Z_T in the same units. Its mass and second Taylor coefficients vanish. This formula follows by subtracting the scalar mass and kinetic projections from the heat of $a_{4,j}(T\phi)^4/4$. In particular the actual outgoing structured coefficient is updated by first heat and this explicit correction alone.

Put $\theta = L^{s-2} < 1$ and choose a_1 with $\theta < a_1 < 1$, also $L^{-2} < a_1$. Give the normalized structured head the weight $h_j = \Lambda_R x_0 a_1^j$; the fixed initial head can use the same weight. Give the local kinetic and full collar quadratic the weight $\Lambda_R x_j$, and $B_{\partial,4}, B_{\partial,6}$ the weights x_j^2. On fixed boxes in these weights,

$$h_j[1 + \log(2 + b_j/R)] \le C_L x_j. \tag{3.74}$$

To check this, $b_j/R = M_R L^j$, the reciprocal sequence estimate implies $x_0/x_j \le C(1 + \lambda^2 x_0 j)$, and $\Lambda_R(1 + \log(2 + M_R)) = o(1)$. The factor a_1^j absorbs any fixed power of $1 + j$. The ordinary bilinear estimate is therefore $O_L(x_j^2)$ on the common enlarged box, including the variation in these coordinates.

For the structured source itself the constant-field quartic projection gives, by the two-line quartic and the one-line quartic–mass contractions,

$$a_{4,j} = 9G^2 \sum_y (TS_j T^*)(y)^2 + 4GuS_j^0. \tag{3.75}$$

Every other quadratic vertex has zero row and contributes zero to this quartic projection. The cumulative bubble bound gives $\sum_y (TS_j T^*)(y)^2 \le C_L \lambda^2 (j + \log(2 + M_R))$, by summing its increments. With $|S_j^0| \le C_L \lambda b^s$, this gives $|a_{4,j}| \le C_L x_j$ on the enlarged box, and the same bound on first variations in a unit normalized quartic direction. In a unit normalized mass direction the bound is $C_L \lambda \Lambda_R x_j^2$, by differentiation of this formula and the inverse mass chart. The old structured density transports with θ. Both its first-heat and corrected source are bounded in outgoing units by

$$C_L \lambda_j^{\mathrm{tad}} (R/b_j)^2 (|G| + |a_{4,j}|). \tag{3.76}$$

Dividing by h_{j+1} works uniformly since

$$\frac{\lambda_j^{\mathrm{tad}}(R/b_j)^2}{\Lambda_R a_1^{j+1}} \le C_L M_R^{-2}(L^{-2}/a_1)^j. \tag{3.77}$$

The normalized stable diagonal from the fixed-head transports is $\theta/a_1 < 1$. The outgoing structured change from first heat is included in $P_{\Gamma_j}U$ when comparing its polynomial with the outgoing ordinary polynomial. The additional D_i correction has size $O_L(x_j^2)$ even in the log-strengthened norm: before the logarithm is applied it is bounded by

$$C_L\lambda_j^{\mathrm{tad}}(R/b_j)^2[\lambda^2(j+\log(2+M_R))|G|^2 + \lambda|G\nu^{\mathrm{loc}}|]. \tag{3.78}$$

The geometric decay of $(R/b_j)^2$ absorbs the additional logarithm. All other changes of ordinary representatives from $P_{\Gamma_j}U$ are $O_L(x_j^2)$ in the strengthened norm by the bounded projection and covariance estimate. Recomputing the ordinary second-order factor therefore costs $O_L(x_j^3)$.

For completeness the scalar quartic projection gives the raw update

$$G_+ - G = -\mathfrak{b}_j G^2 - 4G(u_1 S_{j+1}^0 - uS_j^0). \tag{3.79}$$

Indeed the quartic part of the localized row of $X_{S_j+\Gamma_j}(P_{\Gamma_j}U)$ is the formula for $a_{4,j}$ at the new covariance and with u_1. Heat preserves the quartic part of the old localized row, and the difference of the two bubbles is $\mathfrak{b}_j/9$. In compensated coordinates $g = G(1+4uS_j^0)$, $g_+ = G_+(1+4u_+S_{j+1}^0)$ this implies

$$g_+ - g + \mathfrak{b}_j G^2 = 4S_{j+1}^0\big[(G_+ - G)u_1 + G_+(u_+ - u_1)\big]. \tag{3.80}$$

There is no estimate here on the unprojected size of the cumulative bubble; only its increment is used. The right side is $O_L(x_j^3)$ since $B^s|u_+ - u_1| + |G_+ - G| \le C_L x_j^2$. Replacing G^2 by g^2 costs $O_L(x_j^3)$. The same bounds give

$$|\nu_+ - L^s\nu - 3B^s\gamma_j g| + |z_+ - L^{s-2}z| \le C_L x_j^2 \tag{3.81}$$

in the compensated mass and local kinetic chart. Its quadratic compensation has uniformly bounded coefficients by Section 3.8. Since the first filtered tadpole was assigned to Z_T, it does not contribute to z_+. All these error bounds hold on the enlarged boxes, so Cauchy's inequality gives the stated orders in a unit normalized potential direction.

We next describe the outgoing split on the inner collar. Use the same oriented covariance row and the sum of the actual incoming polynomials on each of its site or tensor anchors. A fine correction may be assigned to the collar whenever one of its stencils meets the region of disagreement. By finite propagation all such anchors are in the chosen new inner collar. In D_i keep the constant and full quadratic part of R_i there, and the unique primary quartic change giving the common outgoing scalar quartic on those filter rows. In forming the full outgoing collar quadratic, subtract the common outgoing mass on the same anchors. The nonprimary parts of degrees four and six of $R_i - D_i$ are grouped into $B_{\partial,4}^+, B_{\partial,6}^+$; the quartic discrepancy from changing the coefficient on an anchor is included in its fourth-degree polynomial. This gives the literal exponent with an outgoing ordinary margin and a full inner quadratic. The uniformly bounded $O_L(x_j^2)$ polynomials on the newly promoted fine anchors are moved into this same collar polynomial. Products discarded in merging them are exactly the products of at least two perturbative factors estimated in Section 3.7.

On the retained old collar tensors, linear heat and reblocking are those of Lemma 3.6. A newly promoted fine anchor carries an ordinary kinetic or head term in the linear quadratic part. Its old stencil is at depth at least cb_j, where the fine tested norm of these terms is at most C_L times their

ordinary coefficient norms. Restrict the coarse tests to the fine stencils and sum the $O(L^d)$ fine labels per coarse tile; this bounds the contribution by $C_L\Lambda_R x_j$. It is a triangular linear arrow from the stable ordinary coefficients to the collar. The common raw mass has no such arrow: its linear uncontracted square is the same before and after engineering rescaling at each anchor. The primary first heat on the $O(B)$ collar satisfies (3.31). At a newly promoted anchor, subtracting the common first-heat mass leaves the zero-row filtered square and the covariance discrepancy in that estimate. Both are retained in the outgoing quadratic, so the connected remainder has no first-order term from this promotion.

The nonlinear collar terms have norm $C_L x_j^2$ in the small-radius units, by the finite-stencil covariance bound. In detail, subtract the independent-copy Gaussian expectation from the covariance. Each factor then has at least one fluctuation direction, and the Gaussian moment and polarization bound that bilinear map by $C_L\sigma_j$, where the raw jet amplitude $\sigma_j = o(1)$ also for the merged step. Each normalized local coefficient direction is bounded by its own weight in raw units. An ordinary margin polynomial is a bounded bilinear function of the outgoing ordinary coordinates and has no collar derivative. Its derivative in a raw mass or stable quadratic direction is at most $C_L x_j$ times that direction in the strengthened norm. These bounds also treat the old ordinary polynomials promoted to the new collar. Thus mass-to-collar quadratic derivatives are $O_L(x_j)$ after normalization, mass-to-higher tensors are $O_L(\Lambda_R)$, and quadratic-to-higher nonlinear stable derivatives are $O_L(\Lambda_R)$. For the linear heat from degree six to degree four the extra factor is $O_L(\sigma_j)$. From degrees four and six to degree two it is $O_L((\sigma_j + \sigma_j^2)x_j/\Lambda_R)$. The diagonal collar factors are at most $CL^{d-1-r(s/2+p)}$ for $r = 2, 4, 6$, times ratios of successive weights tending to one. The normalized quartic-to-collar derivatives are bounded, by the $C_L\lambda_j^{\mathrm{tad}}|G|$ bound for the quadratic source and the $O_L(x_j^2)$ higher sources. The mass row receives only the bulk derivatives. In particular its mass derivative after removal of $L^s\nu$ is $o(1)$ in normalized units.

Finally, the extraction of a small mark changes the ordinary and collar coefficients by $O_L(\tau)$, where $\tau = \|K\|$. Use the exact divergence reassignment preceding Section 3.8. On ordinary anchors this is a common scalar displacement; the nonuniform part and the actual interface terms are assigned to the widened collar. Its quadratic terms there use the full tensors. The corresponding difference of ordinary margin factors is $O_L(x_j\tau)$ by the bilinear comparison on the same stencils. The structured head has no independent extracted coefficient. Its extraction dependence is through the displayed two source formulas. A change of size $O_L(\tau)$ in the extracted coefficients changes $G + a_{4,j}$ by at most $C_L\tau$ on this box: the factor $\lambda^2(j + \log(2 + M_R))|G|$ is uniformly bounded. The normalized head variation is therefore at most $C_L M_R^{-2}\tau/x_0$. The connected expansion only needs the bounded coefficient changes on the enlarged disc $\tau \le cg^{9/4}$. In the shooting norm its potential derivatives and its variation in a unit activity direction are bounded by the same estimates with τ replaced by Cf_j, using Cauchy's inequality on the enlarged boxes. The relevant uniform small factors are

$$x_j/\Lambda_R, \qquad f_j/(\Lambda_R x_j), \qquad f_j/x_j^2. \tag{3.82}$$

For the transient part use $x_j^{-1} \le C(x_0^{-1} + \lambda^2 j)$ and the geometric factor d_0^j. The initial ratios U/G_0^2 and $U/(\Lambda_R G_0)$ tend to zero by Section 5.1; the remaining factors follow from the microscopic scalings.

The only linear stable coupling with a bound not tending to zero besides the diagonals is from the ordinary coordinates to the collar. Use an equivalent fixed product norm giving the collar a sufficiently small relative weight. Choose L large for the diagonals first, and then R large for the small reverse and off-diagonal terms, with the radii of the boxes fixed in this order. All scalar and tensor formulas use identical ordinary stencils away from the halo. Finite propagation of the reblocking and the same choice of localization on the common grid preserve bulk agreement there, by

grouping every disagreeing label into the new boundary family. This proves the coefficient bounds and compatibility with the outgoing split. □

The preceding proposition controls the retained coefficients. The following corollary combines their transport with the contraction of the activity and identifies the data passed to the terminal integration.

Corollary 3.16 (Coefficient and activity bounds along the fractional cube path)**.** *In dimensions $d = 1, 2, 3$, the chart in Proposition 3.15 satisfies the one-step estimates in Assumption 6.1 on the shooting boxes. At its terminal scale, its coefficient bounds and inverse-mass variation are those required in Propositions 7.2 and 7.3.*

Proof. Use the paired bulk and finite-volume activity charts throughout Proposition 3.12. The bulk version has the same range and product estimates, with no collar. Proposition 3.3 bounds the leading linear map by a quantity tending to zero as $L \to \infty$ and then $R \to \infty$. Choose L, the size weight, and then R so that this bound and the errors in Proposition 3.12 are less than $d_0/4$ in total. The same choice applies in the bulk. The uniform disc bound there and Cauchy's inequality give

$$\|D_K \Psi_j(V, K) - D_K \Psi_j(V, 0)\| \le C\|K\|/g_j^{9/4} \tag{3.83}$$

on a smaller concentric disc. The shooting box is contained in this smaller disc and its relative radius tends to zero, since $U/G_0^{9/4} = o(1)$ and $f_j = x_j^3 + U d_0^j$. Since $f_{j+1} \ge d_0 f_j$, the normalized activity diagonal has a strict contraction. The bound on the source term at zero is Cx_j^3, and the potential derivatives on the enlarged shooting boxes are $C(x_j^3 + \|K\|)$ by Cauchy's inequality. This gives the full activity row with values and derivatives in units f_{j+1}. The other rows are supplied by Proposition 3.15.

At the terminal scale the ordinary mass, local kinetic and collar quadratic sizes are $O(\Lambda_R x_N)$. The fixed heads satisfy $h_N \le C\Lambda_R x_N$. The margin and the collar polynomial factors have sizes $O(x_N^2)$ and fixed field degrees. The full coefficient box is holomorphic, including its terminal inverse mass chart. The full collar tensors, fixed heads, and the activity are independent raw coordinates in that inverse. The raw quartic and local kinetic coefficients have derivative in ν_N of size $O(\lambda x_N)$ by Lemma 3.14. In the total quadratic, the raw mass contributes the specified uniform square with multiplicative derivative $1 + O(\lambda x_N)$. The difference from the target uniform mass therefore has derivative $o(1)$ at the small radius. The remaining hypotheses of the two terminal propositions are the raw Gaussian moments in Section 2.3, the fixed-filter bounds, and the already normalized activity estimate. This proves the asserted terminal bounds. □

4 The marginal kinetic coefficient in dimension four

At the local endpoint a quadratic term with two spatial derivatives has zero scaling dimension. We first isolate its coefficient in the bulk quadratic interaction. We then estimate its increments and its transport together with the terms supported near the boundary.

4.1 Local representatives for quadratic terms

An anchorwise cumulative bound in dimension four. The sunset kernel has a logarithmic cumulative second moment at $d = 4$. We now treat this endpoint. Removing the second moment at the original anchor gives a uniform ordinary second-order remainder on a coefficient box where the retained kinetic coefficient is $O(g)$.

Lemma 4.1 (Anchored ordinary reference in $d = 4$)**.** *Use the bulk covariance and filter assumptions of Lemma 3.13 with $s = 2$, and use the same degrees and oriented rows. Represent the ordinary kinetic term on anchor x by*

$$k_x^{\mathrm{a}}(\phi) = -\phi_x \Delta_{\mathrm{c}} \phi_x, \qquad \Delta_{\mathrm{c}} \phi_x = \tfrac{1}{4} \sum_{i=1}^{4} (\phi_{x+2e_i} - 2\phi_x + \phi_{x-2e_i}). \tag{4.1}$$

The quadratic Fourier coefficient of its spatial sum is $|k|^2$. Keep the mass and kinetic moments in the local representatives $\phi_x^2, k_x^{\mathrm{a}}$, with the prescribed quartic and vacuum also kept. Then the ordinary remainder $(1 - \mathrm{Loc})X_S(U)$ has the bilinear bound of Lemma 3.13, with the same strengthened filtered-head coefficient. The bound includes first variations in those coordinates. It also holds after the exact removal of ordinary extraction divergences, with the H-representative kept instead of Z.

Proof. The sums of k^{a} and the central difference squares k^{c} coincide. Indeed, for a fixed coordinate their sums both equal $\frac{1}{4} \sum_x (\phi_{x+2e_i} - \phi_x)^2$. The ordinary $Z + H$ divergence has the same commutator identity as before. Thus the exact exponential change uses the ordinary H-representative with its sign and the interface tensor from that identity.

In the sunset row put $D = TST^*$ and $u = T\phi$. Expand u_{x+r} to spatial order two at x. The discrete Taylor polynomial can use central first and second differences, since these are exact on polynomials of degree two. Its tested error is $C_L \ell_b |r/b|^3$ for $r \neq 0$, using three bounded differences on the ordinary stencil. The odd linear term has zero row because $D(r)^3$ is even. Cubic symmetry makes the quadratic term a scalar multiple of $\Delta_{\mathrm{c}} u_x$. The normalized test bound after subtracting these terms is

$$C_L b^{s-3} \sum_r |D(r)|^3 |r|^3 \leq C_L \lambda^3, \qquad s = 2. \tag{4.2}$$

To replace $u_x \Delta_{\mathrm{c}} u_x$ by $\phi_x \Delta_{\mathrm{c}} \phi_x$, each filtered leg adds $(R/b)^2$ in the corresponding tested difference norm. The logarithm in the removed coefficient therefore costs only $C_L \lambda^3 (R/b)^2 [1 + \log(2 + b/R)]$. The constant term of this spatial expansion is u_x^2 times an absolute row sum $O(\lambda^3 R^{-2})$. The difference from its mass representative has the same $(R/b)^2$ filter gain and the kinetic moment is $O(R^2)$ times this row sum. Thus both give bounded normalizations. These replacements agree exactly with the quadratic Taylor moment of the summed anchor polynomial: the relative k^2 moment of the external filter legs is the one removed from the constant row, and their row sum is one on the spatial quadratic term.

The mixed rows admit the bounds of Lemma 3.13 with this kinetic density. In the two-cross mixed row with a kinetic vertex, the chosen anchor is the quartic and the internal sum is a multiple of $\sum_r (TS)(r)(-\Delta_{\mathrm{c}})(TS)(r)$. By summation by parts over the bulk partner it is the same difference-square sum as before. In a quadratic–quadratic one-cross row rooted at k^{a}, the two contracted possibilities leave Δ_{c} either on the external field or on the smooth summed partner $SQ\phi$. Both give b^{-2}. The Fourier moments of the sum and all the remaining shell estimates are those in the ordinary proof. The finite-jet and outgoing-reference comparison therefore holds on the $O(g)$ kinetic box. □

A fixed short head of vanishing second moment. The local power admits a microscopic quadratic remainder whose coefficients need not decay in the original core units. Its normalized effect decays in the ordinary polynomial after removing its constant and second moments on the actual anchors. The following bound also applies to the bulk head produced by a short inverse in Proposition 4.4.

Lemma 4.2 (Ordinary estimate for a higher-order short head)**.** *In the setting of Lemma 4.1, add to the second type of the incoming quadratic potential the fixed ordinary density*

$$H_x(\phi) = \tfrac{1}{2}\sum_v a_v(\phi_{x+v} - \phi_x)^2, \qquad a_{-v} = a_v, \qquad \sum_v a_v v_i v_k = 0 \quad (1 \le i, k \le 4). \tag{4.3}$$

Assume it has finite range $O(b_0)$ and

$$\sum_v |a_v|\,|v|^2 \le M, \qquad \sum_v |a_v|\,|v|^4 \le CMR^2. \tag{4.4}$$

Then in the bilinear ordinary bound at scale $b \ge b_0$ this head contributes as a coefficient of size $CM(R/b)^2$. In the ordinary kinetic increment, the quartic–head two-cross term from the change of the covariance at an ordinary scale has magnitude at most $C_L\lambda^2 gM(R/b)^2$. For a merged covariance with cumulative gradient bound $\|\nabla TS\|_2^2 \le C\lambda^2 R^{-2}$, it has size $C_L\lambda^2 gM$. These statements are bilinear, with the corresponding first variations.

Proof. In the homogeneous bilinear form of H_x, pair v and $-v$ and Taylor-expand the two differences to degree two. Their cubic contributions cancel. The two linear differences give the tensor $\sum_v a_v v_i v_k$, which vanishes. The remaining tests cost $C_L\ell_b^2|v/b|^4$, by the same discrete Taylor interpolation with bounded differences through order three. Consequently H_x has the bound of a quadratic coefficient $CM(R/b)^2$ on the tile, since $b^d\ell_b^2 = b^2$. The Hessian of the sum has bounded absolute fourth moment CMR^2, zero row and zero second moment. As a convolution on ordinary fields it costs at most $C_L MR^2 b^{-4}\ell_b$ and the corresponding bound with any fixed extra spatial differences. This follows by pairing $v, -v$ and subtracting the Taylor polynomial through degree three.

For a one-cross quartic–head term the chosen anchor is the quartic, and the contracted direction is the convolution of the cumulative TS with this Hessian applied to the field. The bound gains $M(R/b)^2$. In a one-cross quadratic–quadratic term the same convolution argument applies if the head is the summed partner. If the head is the root, the contraction has one difference of the external field and one difference of the smooth summed field. The paired local estimate just proved applies to these two fields, giving the same gain. The summed mass and second coefficients of this term vanish by the zero row and zero second moment of the head Hessian.

In the two-cross quartic–head term the quartic is again the root. Write $P_1 = TS$. The scalar multiplying its external square u_x^2 is bounded by

$$Cg\sum_v |a_v| \sum_r |P_1(r+v) - P_1(r)|^2 \le C_L gM\lambda^2 R^{-2}, \tag{4.5}$$

using telescoping and the cumulative $\|\nabla P_1\|_2^2$ bound. The full ordinary projection of the external square is its mass plus its kinetic moment on that anchor. In fact evenness of the fixed filter, cubic symmetry and its fourth absolute moment imply

$$u_x^2 = \phi_x^2 + c_T\phi_x\Delta_c\phi_x + O_L(\ell_b^2(R/b)^4) \tag{4.6}$$

as a tested polynomial, with $c_T = O(R^2)$. This follows by Taylor-expanding the filter to degree three, with cancellation of both odd degrees; the square of the second-degree correction also has order $(R/b)^4$. The second-degree term has exactly the relative moment of the external square. Thus its unprojected remainder improves on the mere $(R/b)^2$ filter estimate. Multiplication by $b^2(R/b)^4$ in the preceding scalar row bound gives $C_L g\lambda^2 M(R/b)^2$. There is no other nonconstant degree from two quadratics, so the claimed strengthened estimate follows.

Finally write $P_C = TC$ for the filtered ordinary slice. In the kinetic increment the quartic and this fixed short input coefficient do not change under first heat. The scalar difference of the two

covariance rows uses two products of differences where at least one difference is on P_C. By the ordinary slice and cumulative bounds,

$$\|\nabla P_C\|_\infty \le C_L \lambda b^{-3}, \qquad \|\nabla P_1\|_1 + \|\nabla P_C\|_1 \le C_L \lambda b, \qquad s = 2. \tag{4.7}$$

Every shifted v-difference in a product is therefore bounded in the sum over r by $C_L|v|^2\lambda^2 b^{-2}$, using telescoping, one supremum and one absolute row sum. Multiply by gM and by the $O(R^2)$ kinetic moment of the filtered external square. The merged case uses the earlier cumulative gradient estimate. This is the asserted increment and derivative bound. □

A kinetic representative bounded up to the physical faces. A common choice of kinetic density up to the physical faces is also possible. In the following estimate take $J_R = R^{-2} P_\Lambda L_R P_\Lambda^*$ and normalize J_R by a bounded positive factor so that its bulk symbol has leading term $|k|^2$, and use the anchored density $\phi_x (J_R \phi)_x$ with zero extension. Its sum over an enlarged fine tile has bounded bilinear F_b-norm at radius ℓ_b, uniformly up to the faces, edges and corners. For each tested direction v, at distance more than CR from the exterior, the symmetry of the walk gives

$$|(J_R v)_x| \le C \ell_b (\rho_b(x)/b)^p \rho_b(x)^{-2}, \tag{4.8}$$

by summing centered second differences over displacements of size at most R. At smaller distances the same estimate comes directly from the absolute jump-row bound and $\rho_b(x) \asymp R$. Multiply by the field bound on the second direction and use

$$\sum_{x \in P^*} (\rho_b(x)/b)^{2p} \rho_b(x)^{-2} \le C b^2, \tag{4.9}$$

which follows by summing the layers of at most Cb^3 sites and using $2p > 1$. With $\ell_b^2 = b^{-2}$ this proves the uniform norm. Freezing one normalized kinetic coordinate therefore gives the same local bound for its first variation.

In the ordinary region the alternative representative differs from k_x^{a} in Lemma 4.1 by a tested coefficient of size $O((R/b)^2)$ after summing over a tile. The difference kernel is even and has zero constant and second moments, and fourth absolute moment $O(R^2)$, so the fourth-order discrete Taylor bound gives this estimate. The same argument gives the corresponding bound on the difference of smooth contracted legs in the ordinary bilinear estimate. The spread-out representative thus keeps the ordinary kinetic moment and admits uniform tests on an anchor newly assigned to the collar.

4.2 Summable perturbative sources for the local kinetic coordinate

The cumulative sunset has a second moment of logarithmic size in dimension four. The shooting argument of Corollary 6.4 uses its successive increments. At least one covariance line in each increment belongs to the current slice, and this gives a scale-uniform bound.

Lemma 4.3 (Kinetic moment of the sunset increment in $d = 4$)**.** *Write W for a cumulative filtered bulk covariance and C for the next ordinary filtered slice, at side length $b \ge R$. Assume they are supported in $|x| \le C_1 b$ and obey the bounds from Section 2.4,*

$$|W(x)| \le C_2 \lambda (R + |x|)^{-2}, \qquad |C(x)| \le C_2 \lambda b^{-2}. \tag{4.10}$$

The kernels are taken even. The displayed cumulative envelope also holds after filtering: the finite Neumann filter has bounded absolute exponential row and column norms at scale R, and

$(R+|x-y|)^{-2} \le (1+|y|/R)^2 (R+|x|)^{-2}$. *The zero-row kinetic second moment of the three-line quadratic increment, before multiplication by its quartic factors, is bounded by*

$$\sum_x |x|^2 \big|(W+C)(x)^3 - W(x)^3\big| \le C\lambda^3. \tag{4.11}$$

For a merged ultraviolet slice with $W=0$, the bound is $C\lambda^3 \log(2+b_0/R)$.

Proof. Expand the difference as $3W^2C + 3WC^2 + C^3$. By counting lattice shells the first term has moment at most

$$C\lambda^3 b^{-2} \sum_{1\le r\le Cb} r^5 (R+r)^{-4} \le C\lambda^3. \tag{4.12}$$

For the second term the factor b^{-4} multiplies the shell sum with exponent -2. For the third the factor b^{-6} multiplies the sum with exponent zero. These sums are $O(b^4)$ and $O(b^6)$, respectively. Making the kernel zero-row subtracts a multiple of δ_0, which does not affect its second moment. If the two external legs carry the normalized filter, the kinetic coefficient of an even zero-row kernel is unchanged: multiply its symbol by $|\widehat{\overline{T}}|^2 = 1 + O(R^2|k|^2)$ and use that the kernel symbol vanishes to second order at zero. The merged envelope $C\lambda(R+|x|)^{-2}$ gives the logarithm on taking its cube and the second moment. □

The quartic factors in the polynomial increment are bounded by Cx_j^2, and have the same first-variation bound in a unit normalized quartic direction. Thus the sunset contribution has summable size and variation, at most

$$C\lambda^3 \left(x_0^2 \log(2+b_0/R) + \sum_{j<N} x_j^2 \right) \le C\lambda^3 x_0^2 \log(2+b_0/R) + C\lambda x_0. \tag{4.13}$$

Here we used the reciprocal-increment comparison of Section 6. The first-heat structured head has an additional geometric gain: its raw coefficient is bounded by $C\lambda_j^{\mathrm{tad}} g_j b_j^{-2}$ since the filtered slice diagonal gives the tadpole. The normalized zero-row second moment of its fixed shape is $O(R^2)$ at $s=2$. Consequently the incremental bound in that marginal coordinate is $C\lambda_j^{\mathrm{tad}} g_j (R/b_j)^2$, with the same bound for its normalized quartic derivative. Its sum is finite and tends to zero on the microscopic initial scales. These are the polynomial sources in the marginal flow. The connected remainder and its variation require an activity estimate on the chosen kinetic domain. Proposition 4.7 supplies it on the shrinking domain used for shooting.

Appendix C gives the corresponding identities when the kinetic coefficient is held in a fixed disc. For the four-dimensional trajectory we now choose a family of finite-range covariances and bound the residual kinetic coefficient scale by scale.

4.3 A finite-range family of kinetic Gaussian references

For the four-dimensional shooting argument we vary the Gaussian reference before decomposing its covariance. The local split allows a kinetic parameter in the physical block of the auxiliary precision, so that each ordinary slice still has finite propagation. The compensating quadratic density is explicit and the physical spin measure and its pressure are unchanged. The kinetic moment of this density will supply the second adjustable bare coordinate alongside the temperature.

Proposition 4.4 (A tunable finite-propagation Gaussian family)**.** *Work in $d = 4$ with the local split of Lemma 2.1. Let $\zeta = (1 + t_R)^{-1}$, $J_\Lambda = P_\Lambda L_R P_\Lambda^*$, and fix a sufficiently small real δ in an interval independent of R, n. Set*

$$B^\delta = B + \delta L_R, \qquad \mathcal{Q}_{c,\delta} = \begin{pmatrix} I & -\mathcal{L}_\Lambda \\ -\mathcal{L}_\Lambda^* & B_\Lambda^\delta/c \end{pmatrix}, \qquad V_{c,\delta} = \big(A_\Lambda^{-1} - c(B_\Lambda^\delta)^{-1}\big)^{-1}. \tag{4.14}$$

Use the original observation $\mathcal{O} = (\mathcal{L}_\Lambda^*, 0)$. For c below the principal pole $c_{*,\Lambda}^\delta$,

$$V_{c,\delta} = \mathcal{O}\mathcal{Q}_{c,\delta}^{-1}\mathcal{O}^*. \tag{4.15}$$

Define

$$\xi_i = \text{the } i\text{-th increasing eigenvalue of } A_\Lambda^{-1/2} J_\Lambda A_\Lambda^{-1/2}. \tag{4.16}$$

Then $c_{*,\Lambda}^\delta = 1+(\zeta+\delta)\xi_1$. The pole observation direction is independent of δ up to scalar normalization. Set

$$c^\delta = c_{*,\Lambda}^\delta - \theta_0 \lambda^2 \bar{g}_N (R/n)^2. \tag{4.17}$$

The polynomial finite-range decomposition of $C^\delta = a_R V_{c^\delta,\delta}$ has the annular Gaussian moments and bulk filtered bubble bounds of Sections 2.3 and 2.4, uniformly in this real interval. The physical Gaussian density change from C^0 to C^δ inserts the quadratic Hessian

$$D^\delta = (C^0)^{-1} - (C^\delta)^{-1} = a_R^{-1}[c^\delta (B_\Lambda^\delta)^{-1} - c^0 B_\Lambda^{-1}] \tag{4.18}$$

in the residual potential. Its bulk kernel and its boundary defect, in the exponentially weighted absolute row and column norms at length R, are bounded by $C a_R^{-1}|\delta|$. The boundary defect has an additional factor $e^{-a\,\mathrm{dist}(x,\Lambda^c)/R}$ in row x. These estimates also hold for one derivative in δ, with $|\delta|$ replaced by one. All coefficients and finite-horizon Gaussian data are continuous in the real parameter.

Proof. The physical Schur complement at shift $r \geq 0$ is

$$H_{r,\delta} = B_\Lambda^\delta/c + rI - A_\Lambda/(1 + r). \tag{4.19}$$

For $r = 0$ the block inverse tested through the original observation is $A_\Lambda + A_\Lambda H_{0,\delta}^{-1} A_\Lambda$. The matrix inverse identity gives $V_{c,\delta}$. The inverse exists when $H_{0,\delta}$ is positive, which is equivalent to $c < 1 + (\zeta + \delta)\xi_1$, since

$$B_\Lambda^\delta = A_\Lambda + (\zeta + \delta) J_\Lambda. \tag{4.20}$$

The generalized eigenvalues of A_Λ against B_Λ^δ are $[1 + (\zeta + \delta)\xi_i]^{-1}$. If v satisfies $J_\Lambda v = \xi_1 A_\Lambda v$, then $[A_\Lambda^{-1} - c_{*,\Lambda}^\delta (B_\Lambda^\delta)^{-1}] A_\Lambda v = 0$. So the pole direction in the observation space is $A_\Lambda v$. The established two lowest generalized levels at $\delta = 0$ give $\xi_1 \asymp (R/n)^2$ and $\xi_2 - \xi_1 \asymp (R/n)^2$, with the upper bound on ξ_2 from the same min–max comparison. The formula for the generalized eigenvalues therefore preserves the gap uniformly. The exact displacement in c^δ gives the same principal variance in the adapted source coordinate, namely $n^2/(\theta_0 \lambda \bar{g}_N)$.

To verify finite propagation and the local estimates, put $Q^\delta = Q + \delta D_R$, so that $B^\delta = (1+\delta)I - Q^\delta$. The choice of the smoothing amplitude in the proof of Lemma 2.1 gives $Q(x) \geq a D_R(x)$ off the diagonal for a fixed $a > 0$. Thus Q^δ is nonnegative with zero diagonal, range $O(R)$, and supremum $O(R^{-4})$ for sufficiently small signed δ. Its row sum is $q + \delta$, and $(q + \delta)/(1 + \delta) \leq 1 - c_1$ uniformly. The block operator has the original finite propagation and uniform spectral upper bound. In the

polynomial decomposition the degree bounds are unchanged, so each covariance slice has finite range in the same scale units.

Here are the specific uniform inputs for its boundary estimates. In the physical Schur complement the bulk convolution can be written as

$$H_{r,\delta} = ((1+\delta)/c + r)I - Q^{\delta}/c - A/(1+r) = \frac{\zeta+\delta}{c} L_R + rI + \left(c^{-1} - (1+r)^{-1}\right) A. \tag{4.21}$$

The off-diagonal jump kernel in the first expression has the same nonnegative projected spread-out component used in the barrier proof. Thus the normal-coordinate barrier and the loss on jumps killed outside the cube give the same $r^{-1}w(x)$ resolvent row bound, for $w(x) = (\min(b, R+\operatorname{dist}(x, \Lambda^c))/b)^{p_0}$, $p_0 < 1$, with $b = Rr^{-1/2}$. The bounded correction multiplying A is absorbed by choosing the spectral shift as in the unperturbed proof. In the auxiliary Duhamel comparison the diagonal entries are $1+r$ and $(1+\delta)/c+r$; all the remaining couplings are nonnegative. Its no-coupling subtraction and the bulk domination are unchanged. The bulk Schur symbol in the second expression is comparable to $r + \omega(k)$ after the same enlarged regularization. The high-frequency gap, the $O(R^{-4})$ squared integral of the physical off-diagonal multiplier, and the observation smoothing are uniform, since $Q^{\delta} = Q + \delta D_R$ and A is fixed. The interior block equations and cutoff derivative estimates in Lemma 2.11 use a finite-range even physical convolution with bounded row norms, second moment $O(R^2)$, and these low- and high-frequency bounds. All hold with Q^{δ}, including the Fourier derivative bounds on the scaled torus because D_R has fixed scaled support. This verifies the same annular variance and exponential moment proof, including the microscopic collar and merged scales. The tail comparison, principal source bounds and bubble bounds then follow by the same spectral polynomial estimates: the lower bubble band uses the Schur comparison at two fixed regularizations and the observation lower bound at low frequency, and its upper bound uses the just established bulk diagonal bound. These arguments use δ only in fixed compact bounds.

For the local density change use the original and modified short inverse series

$$(B_{\Lambda}^{\delta})^{-1} = \frac{1}{1+\delta} \sum_{m \geq 0} \left(\frac{Q_{\Lambda}^{\delta}}{1+\delta} \right)^{m}. \tag{4.22}$$

Their weighted absolute row sums converge geometrically with a fixed exponential distance weight at length R. A killed-path defect at row x requires at least $c \operatorname{dist}(x, \Lambda^c)/R$ steps. Reserving part of the geometric decay for those steps proves the stated boundary bound. A derivative in δ marks one step or differentiates the scalar normalization and changes the per-length bound by at most a fixed polynomial in m. The path bound with a slightly weakened exponential is uniform for the derivative chain. Since $c^{\delta} - c^0 = \delta\xi_1$, the mean value formula using these estimates gives the additional $|\delta|$ for D^{δ}. Symmetry gives the column estimates. Finally, substituting the two precision matrices into the density identity gives the displayed sign of D^{δ}; the scalar Gaussian normalization is a vacuum. □

The initial kinetic and mass coordinates. Let $p = \widehat{B}(0)$, $B_2 = \frac{1}{2}\sum_x x_1^2 Q(x)$, and $j_2 = \frac{1}{2}\sum_x x_1^2 D_R(x) \asymp R^2$. The bulk two-jet of this density change is explicit:

$$\widehat{D}^{\delta}(k) = \frac{\delta\xi_1}{a_R p} - \frac{\delta\xi_1 B_2 + \delta c^{\delta} j_2}{a_R p^2} |k|^2 + O_{R,\delta}(|k|^4). \tag{4.23}$$

Indeed both short inverse denominators have row sum p and $\widehat{B}^{\delta}(k) = p + (B_2 + \delta j_2)|k|^2 + O_R(|k|^4)$. At fixed c, the mass term is exactly zero. With c^{δ} chosen above its magnitude is $|\delta|\xi_1/(a_R p)$. In the scale-b_0 mass coordinate this is $O(|\delta|\lambda^{-1}(b_0/n)^2)$. The derivative of the kinetic coefficient is

$$-\frac{\xi_1 B_2 + (c^{\delta} + \delta\xi_1) j_2}{a_R p^2}, \tag{4.24}$$

negative and of order λ^{-1}, uniformly at large n/R. In particular a change $|\delta| = O(\lambda^2 G_0)$ changes the initial kinetic coefficient by $O(\lambda G_0)$. Before canonical field conversion the added Hessian has size $O(\lambda^2 G_0) = O(g_{\mathrm{mic}})$ in the exponentially weighted row norm. Its bulk mass may be assigned to the scalar mass chart, its zero-row remainder to the bulk quadratic heads, and its clipping defect to the boundary quadratic heads. These assignments retain the same short-tail path estimates as physical initialization. Thus the initial kinetic coefficient can be varied while keeping the ordinary fluctuations within the finite-propagation construction.

Exact transfer from the unperturbed scale decomposition. For comparison of the terminal distributions across this family, it is convenient to express the perturbed covariance directly in terms of the unperturbed one. This also gives control of the ordinary slices and their first variations. Write $d_\delta = (\zeta + \delta)/\zeta$ and $c = 1 + d_\delta(c' - 1)$, where $c' < c^0_{*,\Lambda}$. For the zero-shift Schur complements one has

$$H_{0,\delta}(c) = \frac{d_\delta c'}{c} H_{0,0}(c'), \qquad V_{c,\delta} = t_\delta V_{c',0} + (1 - t_\delta) A_\Lambda, \qquad t_\delta = \frac{c}{d_\delta c'}. \tag{4.25}$$

Indeed both the coefficient of L_R and the coefficient of A in the first identity have the displayed common factor. Substitution in $V = A_\Lambda + A_\Lambda H_0^{-1} A_\Lambda$ gives the second identity. In the new reference the distance to its pole is d_δ times the original pole displacement. In the bulk two-jet formulas for this choice of c, replace ξ_1 by $(c' - 1)/\zeta = O((R/n)^2)$. The microscopic tail bounds, strict kinetic slope and bound on the scalar perturbation are unchanged in size.

To obtain a positive finite-range decomposition for both signs of δ, fix $\eta > 0$ with $\eta \| \mathcal{Q}_{c',0} \| \le 1/2$. Multiply each original auxiliary slice polynomial by $1 - \eta t$ in its spectral variable t, including the functional-calculus terminal remainder. Functional calculus preserves positivity and bounds each new slice between one-half and one times the original slice. The ordinary propagation range increases by at most the microscopic range of the auxiliary precision. The sums of the resulting slices and remainder give $\mathcal{Q}^{-1}_{c',0} - \eta I$. On applying the observation, multiplying by t_δ, and adding $[1 - t_\delta + t_\delta \eta] A_\Lambda$ to the merged slice, one obtains $V_{c,\delta}$ exactly. The coefficient of A_Λ is positive on a fixed neighbourhood of $\delta = 0$. The positive remainder of the merged slice is comparable to its unperturbed counterpart, and the added A_Λ has fixed microscopic range.

The ordinary annular variance bounds transfer by positive operator comparison, since a stencil variance is a quadratic form in the covariance. The ordinary first δ variation is the same scale kernel times $\partial_\delta t_\delta$, which is uniformly bounded when c' is fixed. The bubble lower bound transfers by the positive slice comparison on the bulk Fourier fibers. The upper bound transfers by the absolute row and diagonal estimates. In the merged slice the additional A_Λ may be bounded directly. Its canonical diagonal is $O(\lambda R^{-2})$ and its supported test differences of order k have standard deviation $O(\sqrt{\lambda} R^{-1-k})$, using the smoothing power in Lemma 2.1 chosen larger than the tested order and the exponentially small white channel. For the supremum of the weighted stencil tests one may even use a union bound on a merged tile: the extra fluctuation parameter is at most $C \lambda M_R^{2P+2p+2} \log(2 + b_0)$, which still tends to zero. This gives the growing fluctuation reserve at the merged scale. Its canonical diagonal alone yields the tadpole bound with λM_R^2. These statements also hold for one variation of the scalar multiplier of this added covariance. The resulting decomposition varies the terminal reference and is available uniformly on the fixed external kinetic domain $|\delta|/\lambda \le c$.

The kinetic source on the shrinking ordinary box. Backward shooting for the residual kinetic coefficient requires a summable bound on its increment and its first variation. It is important here to retain the mass projection. Its cancellation removes the tadpole multiplied by the cumulative bubble, which would otherwise exceed the cubic error after quadratic compensation.

Proposition 4.5 (Bulk kinetic increment in the shrinking chart)**.** *Assume the exact ordinary split of Lemma 3.4 and Lemma 4.1, with the mass and kinetic being the complete quadratic Fourier two-jet. Use the kinetic representative* k^{a}*. Let the bulk quadratic Hessian before the step have symbol*

$$Q(k) = u + z^{\mathrm{loc}}|k|^2 + O(|k|^4). \tag{4.26}$$

Assume $|u| \le Cb^{-2}x_j$*,* $|z^{\mathrm{loc}}| \le Cx_j$ *and* $G \asymp x_j$ *on the enlarged potential box. Include the zero-two-jet short heads as in Lemma 4.2. Let* M_j *bound their total second absolute moment and* CR^2M_j *their total fourth absolute moment. Require* $M_jR/b_j \le Cx_j$ *and* $M_0 \le Cx_0$*. This includes the fixed microscopic remainder and the zero-two-jet part of the filtered square. The comparison sequence and the activity satisfy the bounds of Section 6. For the enlarged extraction chart, assume the retained scalar coefficient and its first variations change by* $O_L(\tau)$ *and the corresponding bounds in Proposition 3.12, with* $\tau \le Cf_j$*. The Gaussian covariance is held fixed in these variations.*

Then for z *in the compensated bulk quadratic chart of Section 3.8,*

$$|z_+ - z| + \|D(z_+ - z)\| \le C_L[\lambda^3 x_j^2 + f_j + \lambda x_j(R/b_j)^2]. \tag{4.27}$$

Here D *measures directions in the enlarged normalized potential and activity boxes, including* z/x_j*. At the merged step the right side may be replaced by*

$$C_L[\lambda^3 x_0^2 \log(2 + M_R) + f_0 + \lambda x_0]. \tag{4.28}$$

This proposition estimates the ordinary bulk source. It assumes availability of the stated local outgoing split and does not construct the coupled boundary chart.

Proof. Use $S = S_j$ for the cumulative bulk covariance and C for the incoming slice. Put $D = TST^*$, $E = TCT^*$, $\gamma = E(0)$, $P = ST^*\delta_0$, $P_C = CT^*\delta_0$. A superscript $+$ on a kernel means that S is replaced by $S + C$. First heat replaces the quadratic Hessian by $Q_1 = Q + 3\gamma G\, T^*T$. The kinetic coefficient of this heat is $O(R^2|\gamma G|)$.

The quadratic part of the Wick covariance projection has three kinds of terms. The three-cross quartic pair has Hessian obtained from $6G^2D(r)^3$ between the two external filtered legs. In the two-cross quartic–quadratic pair both external legs are at the quartic anchor. Its Hessian on that anchor is

$$3G\,\langle P, QP\rangle\,(T^*\delta_x)(T^*\delta_x)^*. \tag{4.29}$$

Indeed the second derivative of $G(T\phi)_x^4/4$ along two filtered directions has scalar $3G(T\phi)_x^2$, and the Hessian of the quadratic is Q. The Wick factor $1/2!$ cancels on passing to the Hessian. This is independent of the anchor assignment when summed in the bulk. Lastly, the one-cross quadratic pair has summed Hessian Q^2S. All other nonconstant degrees are quartic or sextic. The quartic projection before the step has coefficient

$$a_4 = 9G^2 \sum_r D(r)^2 + 4Gu \sum_r S(r). \tag{4.30}$$

The old localized quartic is heated in the anchor subtraction, supplying the same filtered square with Hessian $3\gamma a_4 T^*T$. These formulas use $G/4$ as the quartic normalization throughout.

For the mixed pair, the difference of scalar rows is exactly

$$\langle P^+, Q_1P^+\rangle - \langle P, QP\rangle = \langle P^+, (Q_1 - Q)P^+\rangle + 2\langle P_C, QP\rangle + \langle P_C, QP_C\rangle. \tag{4.31}$$

The mass part uI of the last two summands is bounded by $C_L|u|\lambda^2$, since $\|P\|_1 \le C_L\lambda b^2$ and $\|P_C\|_\infty \le C_L\lambda b^{-2}$, and the same bounds hold for the corresponding factors with P_C. For an even

zero-row kernel Z of second absolute moment M_Z, use $\langle f, Zh\rangle = -\frac{1}{2}\sum_v Z(v)\langle f(\cdot+v)-f, h(\cdot+v)-h\rangle$. The bounds $\|\nabla P_C\|_\infty \le C_L\lambda b^{-3}$ and $\|\nabla P\|_1 + \|\nabla P_C\|_1 \le C_L\lambda b$ give $C_L M_Z \lambda^2 b^{-2}$ for the scalar row. This applies to the local kinetic and to any zero-row head. For a fixed zero-second-moment head one can improve it: put Z on the fresh smooth factor and subtract its Taylor polynomial through order three. With fourth absolute moment $M_{Z,4}$ the bound becomes $C_L M_{Z,4}\lambda^2 b^{-4}$, since $\|\nabla^4 P_C\|_\infty \le C_L\lambda b^{-6}$. These estimates apply to both mixed terms in the display. The first summand in that display is $3\gamma G\sum_r (D^+(r))^2$, whose absolute value is at most $C_L|\gamma G|\lambda^2[1+\log(2+b/R)]$.

Every mixed scalar row is converted to its kinetic coefficient by multiplying by the relative quadratic moment of the two external filter legs, of size $O(R^2)$. Including $3G$, its increment is bounded by

$$C_L\lambda^2 x_j (R/b)^2 (x_j + M_j(R/b)^2 + \lambda x_j[1+\log(2+b/R)]) \tag{4.32}$$

at an ordinary step. This is $O_L(\lambda x_j(R/b)^2)$. Indeed the comparison sequence and $b/R = M_R L^j$ give $\lambda^2 x_j[1+\log(2+b/R)] \le C_L$. At the merged step $S=0$. The bounds $\|P^+\|_2^2 \le C_L\lambda^2[1+\log(2+M_R)]$ and $\|\nabla P^+\|_2^2 \le C_L\lambda^2 R^{-2}$ give the asserted merged envelope for the mixed increment. The kinetic coefficient of the quartic pair is $O_L(\lambda^3 x_j^2)$, by Lemma 4.3. The external filter moment contributes only $C_L R^2 G^2 \sum_r |(D+E)(r)^3 - D(r)^3|$; that sum is at most $C_L\lambda^3 b^{-2}[1+\log(2+b/R)]$ at an ordinary step, which is also sufficient. The merged bound is the stated logarithmic second moment. Finally, $R^2|\gamma|(|G|+|a_4|)$ is at most $C_L\lambda x_j(R/b)^2$, using $|a_4| \le C_L x_j$, and at most $C_L\lambda x_0$ at the merged step.

We next give the raw mass bound required for the quadratic compensation. In its projection the sunset increment has potentially singular term $18G^2\gamma\sum_r D(r)^2$. The difference of its D^2E row from that singular part is $O_L(\lambda^3 b^{-2}G^2)$: use $|E(r)-\gamma| \le C_L\lambda b^{-2}|r|/b$ and $\sum_r |D(r)|^2|r|/b \le C_L\lambda^2$ on the covariance support. The terms with at least two E have the same bound. The mixed row from the first-heat change $Q_1 - Q$ has singular mass $9G^2\gamma\sum_r D(r)^2$. Its difference from the full mixed contribution with $Q_1 - Q$ is bounded by $C_L G^2|\gamma|\lambda^2$, since $\sum_r |(D^+(r))^2 - D(r)^2| \le C_L\lambda^2$. The old localized quartic heat subtracts precisely $27G^2\gamma\sum_r D(r)^2$ in the potential correction. Thus all three singular coefficients cancel. The remaining part of $3\gamma a_4$, involving $Gu\sum S$, has size $C_L b^{-2}x_j^2$. The previously estimated non-heat mixed rows also have this bound in raw mass units: for the fixed zero-second-moment head use $b^{-4}M_{Z,4}$ and $M_j(R/b)^2 \le C_L x_j$. For the merged step the uncorrected cubic row bound $C\lambda^3 R^{-2}$ replaces this cancellation. After multiplication by b_0^2 its additional factor $\lambda^3 M_R^2$ is bounded. The merged non-heat and heat mixed rows similarly use $\lambda^2 M_R^2$ and $\lambda^3 M_R^2[1+\log(2+M_R)]$ times bounded normalized inputs, so give the same conclusion.

For the pure quadratic part and its compensation, work in the algebra of Fourier two-jets. The reference update subtracts $Q_1^2(S+C) - Q^2 S$ as well as the non-quadratic projected correction just bounded. Each quadratic quantity here is in dimensionful mass and kinetic units. The pure projected kinetic correction is $O_L(x_j^2)$, using $S(0), (S+C)(0) = O_L(\lambda b^2)$ and their second moments $O_L(\lambda b^4)$. Its mass correction is $O_L(b^{-2}x_j^2)$ by the same calculation. Together with the preceding mixed, sunset and tadpole estimates this proves

$$|u_+ - u_1| \le C_L b^{-2} x_j^2, \qquad |z_+^{\mathrm{loc}} - \mathrm{kin}\, Q_1| \le C_L x_j^2. \tag{4.33}$$

In the kinetic estimate the old localized quartic heat is $O_L(x_j^2)$ even though its rough bound above was linear in x_j: retain $|a_4| \le C_L[\lambda^2(1+j+\log(2+M_R))+\lambda]x_j^2$ and use the geometric factor $R^2|\gamma|$. Here kin denotes the coefficient of $|k|^2$. The merged bound holds since $a_4 = 0$ there.

The compensated coordinate is the kinetic two-jet of $Q + Q^2 S$. Substituting the reference update into this expression cancels the pure quadratic projection exactly to second order. Its only remaining term in this calculation is the corresponding two-jet of

$$(Q_+^2 - Q_1^2)(S+C) = (Q_+ - Q_1)(Q_+ + Q_1)(S+C). \tag{4.34}$$

Its kinetic coefficient is $O_L(\lambda x_j^3)$ by the last displayed bounds. It fits inside f_j. Using the exact rational quadratic two-jet instead changes these estimates only by a cubic expression on this shrinking domain. The extraction changes cost $O_L(\tau)$ in the retained coefficient and have the stated variation bound. All the polynomial estimates hold also on the complex potential box. Applying Cauchy's inequality on a fixed smaller normalized box gives their derivative bounds with the same envelopes. This proves the assertion for the bulk step. □

Orientation when a filtered quadratic is split into two jets. The first heat of the quartic produces a kinetic two-jet and a short quadratic remainder. Both pieces can participate in a two-cross contraction with a quartic. Assigning that contraction to the quartic anchor requires a redistribution of the individual covariance rows. The next lemma bounds this redistribution and uses its antisymmetry to preserve the sum for each pair.

Lemma 4.6 (Reorientation of the ordinary two-cross rows)**.** *Use the covariance, test kernels and ordinary enlarged shrinking box of Proposition 4.5. In addition to the fixed difference-square head in Lemma 4.2, allow a short filtered quadratic of the form*

$$q_T\big[(T\phi)_x^2 - \phi_x^2 - c_T\phi_x\Delta_{\mathrm{c}}\phi_x\big], \qquad |q_T|R^2 \le M_j, \qquad M_jR/b_j \le Cx_j, \qquad M_0 \le Cx_0, \tag{4.35}$$

where c_T is the quadratic coefficient of the summed external filtered square in this representation. Use the symmetric covariance row for the terms other than the mixed two-cross term of a quartic and a zero-row quadratic. Assign the latter term to the quartic anchor.

This ordinary representative has the bilinear bounds of Lemma 4.1, with an effective strength $CM_j(R/b)^2$ for the zero-two-jet heads, including the fixed microscopic head if its moments are included in M_j. In the anchorwise heat identity it suffices to replace the symmetric fresh covariance row by a row with the same pair totals and an $O_L(x_j^2)$ bound per fine label after summing its neighbors. The row is supported on pairs at distance $O_L(b)$, and the bounds hold uniformly in the enlarged complex potential domains.

Proof. Let $P = TS$, $D = TST^*$, and $A_T = T - I$ as bulk convolution kernels. If $T = I + (c_T/2)\Delta_{\mathrm{c}}$ to quadratic order, $A_T - (c_T/2)\Delta_{\mathrm{c}}$ has zero row and second moment. Its absolute moments of orders two and four are $O(R^2)$ and $O(R^4)$. For the mixed two-cross term of this filtered head and a quartic the scalar radial row is, up to fixed Wick factors,

$$F_S = D^2 - P^2 - c_TP\Delta_{\mathrm{c}}P = 2P(A_T - (c_T/2)\Delta_{\mathrm{c}})P + (A_TP)^2, \tag{4.36}$$

where the square and the first multiplication are pointwise. By fourth-order filter cancellation and the cumulative kernel difference bounds,

$$|F_S(r)| \le C_L\lambda^2R^4(R + |r|)^{-8}. \tag{4.37}$$

For this estimate use paired Taylor subtraction for shifts at most $(R + |r|)/2$, and the exponentially weighted row bound for the complementary shifts. The anchored external filtered square has a fourth-order remainder after mass and kinetic subtraction, as in Lemma 4.2. The two-cross term assigned to the quartic anchor therefore costs $C_Lg\lambda^2M_j(R/b)^2$. A one-cross term rooted at the head is the polarized difference of the filter product and its degree-two spatial Taylor polynomial, evaluated on the external field and a smooth summed partner. Its fourth-order bound gives the same geometric gain. If the head is the summed partner the bulk convolution Hessian has zero row and second moment, so the gain follows by putting its four differences on that smooth partner. These also handle quadratic–quadratic rows. A local kinetic one-cross term may be kept symmetric:

at its root its two alternatives put the Laplacian on the external field or on the summed smooth partner. Both cost b^{-2}. This proves the ordinary reference bounds with the chosen convention.

Denote the correction of the ordinary row on the pair (i,k) by $J^S_{ik}(U)$. For each mixed term it is half the corresponding scalar times the difference of the two external filtered squares; thus $J^S_{ki} = -J^S_{ik}$. The exact heat row is the symmetric fresh row plus

$$J^{S+C}_{ik}(P_C U) - P_C J^S_{ik}(U). \tag{4.38}$$

This follows by applying the heat identity of Lemma 3.4 to the original symmetric row and adding the correction on both sides. Its pair totals are unchanged. In the bulk P_C of the difference of the two external filtered squares adds zero because the diagonal heat is uniform on the anchors.

We give estimates for this difference. Write $P_C = TC$ for the fresh filtered kernel. For the filtered head above the change of F_S is a sum of the cross products in its last formula with one or both kernels replaced by P_C. With $v(r) = \min(1, |r|/b)$, its weighted absolute row sum is

$$\sum_r v(r)|F_{S+C}(r) - F_S(r)| \le C_L \lambda^2 R^3 b^{-3}. \tag{4.39}$$

Indeed the term with P_C multiplying $(A_T - (c_T/2)\Delta_{\mathrm{c}})P$ uses $\|P_C\|_\infty \le C_L \lambda b^{-2}$ and $\sum |r||((A_T - (c_T/2)\Delta_{\mathrm{c}})P)(r)| \le C_L \lambda R^3$. The latter follows from the pointwise bound $C_L \lambda R^4 (R + |r|)^{-6}$. For the oppositely placed fourth difference use $C_L \lambda R^4 b^{-6}$ on P_C and $\|P\|_1 \le C_L \lambda b^2$. In the cross product with A_T on each factor use $\|A_T P\|_1 \le C_L \lambda R^2 [1 + \log(2 + b/R)]$ and $\|A_T P_C\|_\infty \le C_L \lambda R^2 b^{-4}$. The terms with two fresh factors satisfy the same bound. Use $(R/b)[1 + \log(2 + b/R)] \le C$ to conclude. The factor $v(r)$ is supplied in multilinear tests by the difference of the external filtered squares. Multiplication by $g|q_T|b^2$ gives $C_L g \lambda^2 M_j R/b = O_L(g^2)$.

For a fixed difference-square head of zero second moment, the scalar cross row is

$$\sum_v a_v (\Delta_v P)(\Delta_v P_C), \qquad \Delta_v f(r) = f(r+v) - f(r). \tag{4.40}$$

Pair opposite v and Taylor-expand both differences to order two. The degree-two product cancels by the moment condition and the cubic product by symmetry. The absolute summed remainder is bounded by

$$C_L M_j R^2 (\|\nabla^3 P\|_1 \|\nabla P_C\|_\infty + \|\nabla^2 P\|_1 \|\nabla^2 P_C\|_\infty + \|\nabla P\|_1 \|\nabla^3 P_C\|_\infty). \tag{4.41}$$

To see this without a pointwise restriction on the span, subtract the linear and quadratic products in an order that pairs a third-order remainder with the full first difference, a second-order remainder with the second-order term, and a third-order remainder with the first-order term. The telescoping remainder bounds in ℓ^1 and ℓ^∞ give $|v|^4$ in each case. The cumulative derivative row norms of orders one, two and three are respectively $C_L \lambda b$, $C_L \lambda [1 + \log(2 + b/R)]$ and $C_L \lambda / R$. The fresh derivative bounds then give $C_L \lambda^2 M_j R b^{-3}$. The same calculation applies when both factors are fresh, and multiplication by gb^2 again yields $O_L(g^2)$.

For the local kinetic with coefficient $O(g)$, there is no need for a zero second moment. One supremum of a fresh first difference and one cumulative ℓ^1 first difference give $C_L \lambda^2 |z| b^{-2}$, as in the kinetic source estimate. Its row correction is therefore $O_L(g^2)$ in the fine jets. The remaining change of the quadratic coefficients under the first heat has zero-row part $3\gamma G(T^*T - I)$. Here the scalar row is $(D^+)^2 - (P^+)^2$, whose absolute row norm is $O_L(\lambda^2)$ by even-filter second differences. The coefficient is $O_L(\lambda g b^{-2})$, so its changed orientation also costs $O_L(g^2)$. At the merged step use $S = 0$, the cumulative kernel bounds from length R, and the factor b_0^2/R^2 in raw mass normalization. The

resulting coefficients are bounded by $C_L\lambda^2 M_R^2 g^2$, including the filtered head and first heat after their factors of λ are included. This preserves the bound.

All these transfer rows have the finite support of the cumulative and fresh kernels and their stencils. The finite cube subtraction may use their bulk versions on the pairs read by any new ordinary anchor, and zero or their antisymmetric completion on the other pairs inside the boundary chart. A fixed coarse collar contains all nonordinary neighbors. Antisymmetry leaves the pair covariance sums unchanged, so the same connected row cancellation is valid; no equality of bulk and clipped coefficients at an individual collar anchor is needed. □

4.4 The coupled cube chart for a residual kinetic coefficient

Fix a member of the Gaussian reference family during the RG iteration. The bulk source estimate controls the small residual kinetic coefficient in this reference. It remains to use the same quadratic coefficients in the ordinary and collar parts of the activity map. The ordinary reference retains the full kinetic two-jet. When a fine anchor enters the coarse collar, its actual fine quadratic is assigned to the collar tensor. In particular, the central difference stencil is used only on the ordinary side of the physical boundary.

Proposition 4.7 (Coupled anchored split on the shrinking kinetic box)**.** *Let $d = 4$, $1/2 < p < 1$, and use the canonical finite-range covariance, including any fixed member of Proposition 4.4. Use the comparison sequence of Section 6, and take $\Lambda_R \asymp \lambda M_R^2$. Assume the initialized free and ordinary activities agree off the grouped boundary family and have the bounds of Proposition 5.1 with the chosen quadratic inputs. Keep the Gaussian reference fixed when defining derivatives in the shooting variables.*

There is one outgoing anchored split as required in Proposition 3.12, on the enlarged complex domains

$$|\nu_j| + |z_j| + \|Q_{\partial,2}\|_{\ell_j,0} \le C\Lambda_R x_j, \qquad \|B_{\partial,4}\|_{\ell_j,0} + \|B_{\partial,6}\|_{\ell_j,0} \le C x_j^2, \tag{4.42}$$

where the displayed output tensor bounds describe their coordinate units. The incoming structure also allows the fixed short zero-two-jet head from physical initialization and a filtered zero-two-jet head in the geometric units specified below. The stable-coordinate derivative is strictly contracting in an equivalent product norm, with the marginal kinetic coordinate held fixed. The normalized mass and quartic rows obey Assumption 6.1. The additional kinetic row obeys Proposition 4.5, with bounded normalized arrows from $z_j/(\Lambda_R x_j)$ to the other rows. The output retains agreement with the bulk off the new boundary family.

Proof. Representatives and weights. Use ordinary anchored mass, the common positive filtered primary quartic, and the kinetic density k_x^{a} of Lemma 4.1. This density is used only on stencils a fixed positive fine depth from the physical boundary. The ordinary quadratic Hessian is then its full quadratic Fourier two-jet plus the two short heads of zero two-jet. The fixed microscopic head is put in difference-square form, with its mass and kinetic moments removed before it is assigned to ordinary anchors. The finite truncation can be performed first and its second moment then removed exactly. Subtract a central difference square with the same cubic second moment, and transfer that square to k^{a} by the exact divergence identity in Section 3.7. Its interface piece belongs to the initialized boundary tensor. The exponentially small discarded quadratic tail is assigned as in physical initialization. The retained fixed head has bounded span $O(b_0)$, second absolute moment $O(\lambda x_0)$, and fourth absolute moment $O(\lambda x_0 R^2)$.

For the initial off-shell kinetic variable, multiply the reference on each ordinary initial anchor by the exponential of its assigned kinetic change, and multiply an occupied component by the product

of these exponentials on its actual labels. Keep the full boundary quadratic fixed in this operation. On the real physical parameter the kinetic change is zero relative to the bare coefficient, so this is the original initialized interaction. For all off-shell values the additive distribution makes the circle product identity exact, just as for initial mass reweighting. In the initial ordinary region all central stencils have depth comparable to b_0, so a change of size $O(x_0)$ has small-radius quadratic cost $O(x_0)$ per label and large-radius cost $O(\sqrt{x_0})$. The uniform initial field reserve absorbs this exponent on the enlarged disc, with an additional fixed fraction if the mass is varied simultaneously. Thus the new initial activity has size $O(U)$, and a unit normalized kinetic variation has the same bound by Cauchy's inequality. Ordinary anchors outside the boundary family are assigned in the same way in the two initializations, preserving their agreement.

Use the filtered quadratic density

$$H_x^T(\phi) = (T\phi)_x^2 - \phi_x^2 - c_T\phi_x\Delta_{\mathrm{c}}\phi_x \tag{4.43}$$

from Lemma 4.6. Write its raw Hessian coefficient as $q_{T,j}$, so its contribution to the potential has coefficient $q_{T,j}/2$. Its second absolute-moment parameter is at most $CR^2|q_{T,j}|$. At scale $b = b_j$ define the effective coefficient $m_{T,j} = (R/b_j)^2R^2q_{T,j}$. Its unit is

$$h_j = \Lambda_R x_0 a_1^j, \qquad L^{-2} < a_1 < L^{-1}. \tag{4.44}$$

On the enlarged boxes this gives

$$\frac{R}{b_j}R^2|q_{T,j}| \le C\Lambda_R M_R x_0 (La_1)^j \le C_L x_j. \tag{4.45}$$

The bound uses the reciprocal comparison for x_j and $\Lambda_R M_R = o(1)$. At the merged step $R^2|q_{T,0}| \le C\Lambda_R M_R^2 x_0 \le Cx_0$. These are the strengthened estimates for the effective orientation row and the mixed kinetic source. In particular the enlarged domains allow $R^2|q_{T,j}|$ to be larger than its value on the physical trajectory. The effective quartic coefficient norm of Lemma 4.6 is $O(x_j)$ in these domains.

Ordinary reassembly. Write G for the raw primary coupling, $u = b^{-2}\nu^{\mathrm{loc}}$, S_j for the cumulative bulk covariance and $\gamma_j = (T\Gamma_j T^*)(0)$ in the bulk. First heat on the ordinary quartic provides the full filtered square $3\gamma_j G(T\phi)_x^2$ in Hessian units. The old localized quartic provides another such square with coefficient $3\gamma_j a_{4,j}$, where

$$a_{4,j} = 9G^2\sum_r (TS_jT^*)(r)^2 + 4Gu\sum_r S_j(r). \tag{4.46}$$

The first term follows from the two-cross quartic pair, the second from the one-cross mass pair. A zero-row quadratic has zero constant-field quartic projection. The first-heat square splits literally into its mass, k^{a} part, and H_x^T; the same split is used for the square from the old localized quartic. The ordinary outgoing Hessian two-jet is the projection of the full $P_{\Gamma_j}U - D$, with the total kinetic coefficient including those two filtered squares. Thus the remaining head coefficient has the exact scalar equation

$$q_{T,j+1} = q_{T,j} + 3\gamma_j(G + a_{4,j}). \tag{4.47}$$

This is the equation before extraction when the input is U, and the same formula with the extracted ordinary coefficients is used on the enlarged activity disc. The effective transport of the old term is L^{-2}. The absolute effective forcing is at most

$$C_L\lambda_j^{\mathrm{tad}}(R/b_j)^4 x_j, \qquad \lambda_j^{\mathrm{tad}} = \begin{cases} C\lambda M_R^2, & j = 0, \\ C\lambda, & j > 0. \end{cases} \tag{4.48}$$

Here $3|\gamma_j| \le C_L \lambda_j^{\mathrm{tad}} b_j^{-2}$ and $|a_{4,j}| \le C_L x_j$ by the cumulative bubble sum and the reciprocal comparison. Dividing by h_{j+1} gives a uniformly small bound, since $L^{-4} < a_1$. The same estimate holds for a unit normalized potential derivative by the explicit formula and Cauchy's inequality. Thus this head has a strict stable diagonal L^{-2}/a_1, with small extra normalized entries.

Let $\widetilde{X}_{S,i}$ be the anchored ordinary covariance polynomial using the antisymmetric reorientation of Lemma 4.6, and let $W_i = (1 - \mathrm{Loc})\widetilde{X}_{S_j,i}(U)$. In the ordinary region let

$$D_i = \big(\mathrm{Loc}\, \widetilde{X}_{S_j+\Gamma_j}(P_{\Gamma_j} U)\big)_i - P_{\Gamma_j}\big(\mathrm{Loc}\, \widetilde{X}_{S_j}(U)\big)_i. \tag{4.49}$$

Use the effective fresh row of Lemma 4.6 in the covariance reapportioning. The anchorwise identity gives $R_i - D_i = (1-\mathrm{Loc})\widetilde{X}_{S_j+\Gamma_j,i}(P_{\Gamma_j}U)$. Here all ordinary coefficient formulas are evaluated on the bulk comparison of the extracted input where appropriate. These identities use the actual anchored representatives, including the split of the first-heat square. If a fine label contains both site densities and a previously grouped ordinary polynomial, use the latter's sum of anchored densities in this linear calculation. The product errors from grouping are counted only after the two orders have been matched.

Both R_i and the displayed remainder have $O_L(x_j^2)$ fine jet bounds. For R_i, use the fixed-stencil Gaussian moment for $P_{\Gamma_j} W_i$, and the bounded effective fresh row. For the remainder, use Lemmas 4.1, 4.2 and 4.6. Hence $D_i = O_L(x_j^2)$. Its quartic is along the primary density. As a consequence the ordinary positive sector persists throughout the real coefficient segments used for the product estimates. The change of q_T in the ordinary first-heat input and in the literal outgoing potential differs only by $3\gamma_j a_{4,j}$. In the strengthened units for the row reorientation its size is at most $C_L \lambda_j^{\mathrm{tad}} (R/b_j)^3 |a_{4,j}|$. Using

$$|a_{4,j}| \le C_L[\lambda^2(1 + j + \log(2 + M_R)) + \lambda]x_j^2 \tag{4.50}$$

makes this $O_L(x_j^2)$, including the merged input. The other ordinary coefficients in the literal outgoing potential differ from $P_{\Gamma_j} U$ by $O_L(x_j^2)$, by the raw mass and kinetic estimates in Proposition 4.5 and the quartic projection. Thus replacing the ordinary bilinear factor evaluated at $P_{\Gamma_j} U$ by the one evaluated at the literal outgoing coefficients has size $O_L(x_j^3)$ on the same anchors. The local comparison of Section 3.7 applies.

The common collar and ordinary margin. At each scale mark a fixed number of cell layers along the physical boundary. For a side length not divisible by the grid size, start from the same corner and merge the last incomplete interval into the previous full interval in each coordinate. The fine merged interval starts within the last coarse interval, and all other fine boundaries respect the integer ratio of the two grids. The Cartesian products thus give nested rectangular labels. The clipped labels belong to the boundary family. Keep the disjoint nested cell labels and, if grouping at a face or an edge is used, group only a bounded number of adjacent cells per label. Thus each label has volume comparable to b_j^4. Inside that boundary family keep a collar of thickness $C_* b_j$ at scale b_j; the remaining anchors there are its ordinary margin. Choose C_* larger than the range and support constants in coarse units, and leave a fixed additional buffer of coarse labels for ordinary suffixes of the enlarged stencils. Thus at the step to $B = Lb$ the new collar can have thickness $C_* B$. It contains the old boundary stencils, their localization halo, and every anchor whose incoming covariance row or polynomial support reads nonbulk data. This uses finite propagation at range $O(B)$; it does not require the range constant in fine units to be independent of L. The incoming boundary tensors still occupy only a fixed width in fine units. The ordinary expression applies to every margin anchor outside the buffer. One can implement the buffer by using the slightly enlarged collar in the boundary labels and reducing the margin accordingly; the width and stencil enlargement constants remain fixed in coarse units.

Use the actual finite-cube covariance row on these supports. Add the ordinary antisymmetric correction on pairs for which it is required by a new ordinary anchor; complete it antisymmetrically on the adjoining collar anchor. Elsewhere in the collar one may use the symmetric row. Both pair totals and support bounds are unchanged. All added rows have fine jet norm $O_L(x_j^2)$ by Lemma 4.6; they need agree with the actual finite covariance only in their symmetric pair totals. The definition of R_i then uses the whole $P_{\Gamma_j}W_i$ on each old support. On each new inner anchor take the complete constant and quadratic part of R_i into D_i. The quartic part put in D_i is just that needed to change the local primary coefficient on its actual filter row to the outgoing common value. The difference in these scalar quartic coefficients is $O_L(x_j^2+\tau)$. Put the remaining homogeneous degree-four and degree-six coefficients in the corresponding inner polynomial factor, and record constants separately. Define the new inner quadratic by summing the literal $P_i - D_i$ there and subtracting only the common mass on those same anchors. All kinetic and short-shape quadratics inside the collar are therefore in its full tensor. Exponents match exactly, and the linear polynomial factors match through second order. Grouping their higher products has the mixed cubic cost proved in Section 3.7. No summation-by-parts equality of distinct finite-volume cell factors is used.

For the linear coefficient bounds on this chart, inherited boundary tensors are supported within $O(b)$ of the physical boundary on disjoint fine labels. They therefore have the gains of the full collar heat estimate, namely CL^{1-2p}, CL^{-1-4p}, and CL^{-3-6p} for the matched degrees two, four, and six. The constants are independent of L. They follow by counting at most CL^3 fine boundary labels in one coarse enlargement, and restricting each external direction with factor CL^{-1-p}. The Gaussian heat off-diagonal terms with $r-m$ contracted directions add the factor $C_L\sigma_j^{(r-m)/2}$, where $\sigma_j = o(1)$ also at the merged step. These estimates apply equally to face, edge, and corner stencils since the union of the face layers overcounts their intersections.

There are also genuine linear terms on newly promoted ordinary anchors. For those stencils the smallest physical depth is a fixed positive multiple of b; shallower stencils were already stored in the old tensor or its fixed halo. A coarse test direction on a stencil at depth t has value bounded by $C\ell_B(t/B)^p$ and a difference of order a bounded by the same value times t^{-a}, with t capped at B. For a kinetic density the two legs thus cost $C\ell_B^2(t/B)^{2p}t^{-2}$. There are at most CB^3 sites in a unit depth layer. The kinetic contribution in normalized coarse units is at most

$$C|z^{\mathrm{loc}}|B^{1-2p}\left(\sum_{cb\le t\le B} t^{2p-2} + B^{2p-1}\right) \le C|z^{\mathrm{loc}}|. \tag{4.51}$$

The term at depth comparable to or larger than B uses the capped bound. The analogous polarized bound for H^T or the fixed short head has t^{-4} in place of t^{-2} and coefficient CM_jR^2, where $M_j = R^2|q_{T,j}|$ for H^T and M_j bounds the second absolute moment for the fixed head. The polarized fourth-order angular cancellation is the same as in Lemma 4.2. It gives

$$CM_jR^2B^{1-2p}\left(\sum_{cb\le t\le B} t^{2p-4} + B^{2p-3}\right) \le CM_j(R/b)^2L^{1-2p}. \tag{4.52}$$

At the lowest scale choose the finite-head cutoff inside the fixed initial ordinary stencil buffer, as permitted by the exponentially small tail bound. Thus the line segments used in the Taylor estimate lie at comparable depth even at that scale. If a stencil in the fixed halo reads a nonordinary anchor, the whole halo is kept in a full tensor instead. These linear arrows are bounded in the quoted units, because $|z^{\mathrm{loc}}| = O(\Lambda_R x_j)$ and the effective head weights are $h_j = O(\Lambda_R x_j)$. They are arrows into the collar. There is no reverse linear arrow into an ordinary coefficient.

The uncontracted bare mass density and its engineering rescaling are identical in the two charts on every anchor. Its linear contribution cancels exactly upon subtracting the common outgoing mass. The first-heat quadratic compares $3G(T\Gamma_j T^*)_{xx}$ with its bulk coefficient on the corresponding filter rows. Applying (3.31) with $s = 2$ on the entire $O(B)$ collar gives $CL^2\lambda_j^{\mathrm{tad}}|G|$ per coarse tile, including the difference of the filtered and bare squares at promoted anchors. Its normalized quartic derivative is at most $C_L\lambda_j^{\mathrm{tad}}/\Lambda_R$. In the inverse mass chart G depends only on g_j, ν_j. The derivative of this source with respect to the residual kinetic coordinate at fixed covariance is therefore zero.

All remaining nonlinear sources and the old ordinary polynomial factors promoted into the collar have fine coefficient size $C_L x_j^2$. This follows from polarization and the ordinary comparison on the same anchors. In the genuinely clipped covariance rows use the two centered Gaussian increments, which together cost $C_L\sigma_j$ times the product of the fixed-stencil potential norms. The effective antisymmetric rows are bounded by their ordinary analytic estimates. The resulting coarse tensor bound is $C_L x_j^2$ since the number of fine labels inside the new fixed coarse collar is $O_L(1)$ per coarse label. The same statement holds for the residual mass after its common linear density has been removed. Cauchy's inequality in the normalized boxes gives $O_L(x_j/\Lambda_R)$ for its contribution to the quadratic row after division by $\Lambda_R x_{j+1}$. For derivatives from a quadratic coefficient into a quartic or sextic factor, bilinearity gives $O_L(\Lambda_R)$ after division by x_{j+1}^2. The bound uses the effective norm of a head on its ordinary polynomial stencils. The ordinary margin factors have no boundary derivative and are subject to the same ordinary bilinear variation bound. The normalized heat 6 to 4 is $O_L(\sigma_j)$; the heat $4, 6$ to 2 is $O_L((\sigma_j + \sigma_j^2)x_j/\Lambda_R)$. Thus all reverse and nonlinear variations in the stable block are small, apart from the stated triangular linear arrows. Choose L large for the diagonal gains, and then give the collar a smaller relative weight than the ordinary stable coordinates. This absorbs the bounded one-way stable arrows. The marginal input is kept independent in this stable-block estimate.

Scalar rows and the activity domain. The raw ordinary scalar quartic obeys

$$G_+ - G = -\mathfrak{b}_j G^2 - 4G(u_1\widehat{S}_{j+1}(0) - u\widehat{S}_j(0)), \qquad u_1 = u + 3\gamma_j G. \tag{4.53}$$

The proof is the constant-field projection used for $a_{4,j}$, since the kinetic and both short heads have zero row. Multiplying by the compensated factor $1 + 4u\widehat{S}_j(0)$ cancels the second term to cubic order, by the raw coefficient bounds in Proposition 4.5. The mass equation after subtracting its engineering term has forcing $3B^2\gamma_j G + O_L(x_j^2)$ in outgoing units. The only potentially unbounded cumulative bubble term in that error cancels exactly as proved in the same proposition. The compensated quadratic projection removes the leading pure-quadratic variation. Its inverse changes a unit raw mass variation by $1 + O(\lambda x_j)$ and the remaining raw coefficients by $O(\lambda x_j)$, as in Section 3.8 on this shrinking box. The normalized mass entries from the residual mass error are small since $x_j/\Lambda_R = o(1)$. The kinetic source estimate itself is Proposition 4.5. All these scalar maps are holomorphic on the larger coefficient box, with the asserted bounds for their first variations by Cauchy's inequality.

On the extraction disc the uniform ordinary displacement of mass, quartic and kinetic is $O_L(\tau)$ in normalized raw fine-scale units. Implement the extra spatial-moment densities by the exact exponential divergence reassignment of Section 3.7, using k^{a} instead of the central difference square. The scalar kinetic span is unchanged. The nonuniform displacement stays local in the fine halo and is assigned to the new collar. No independent filtered head displacement is extracted. Its formula above changes by $C_L\lambda_j^{\mathrm{tad}}(R/b_j)^4\tau$ in its effective units, because $\lambda^2[1 + j + \log(2 + M_R)]|G| \le C_L$. The new collar extraction difference is $O_L(\tau)$ and the ordinary margin difference is $O_L(x_j\tau)$, by the local matching estimates and their bilinear variation. These statements hold on the entire enlarged disc, and their first variations follow as in Proposition 3.12. On the shooting boxes the

additional normalized factors are $f_j/(\Lambda_R x_j)$ and f_j/x_j^2, both $o(1)$ by the microscopic comparison. The same bounded matching proves the connected disc bound and the strictly contracting activity derivative by that proposition. Thus extraction and transport use the same outgoing reference on each anchor. □

The number of cells on a face grows with its area. To obtain a uniform bound when an ordinary label enters the boundary family, we sum over connected covers containing one fixed cell.

Lemma 4.8 (Anchored boundary promotion on a cube)**.** *Use the bounded-degree graph of cube cells of scale b_0. Let $K^0(X)$ be its connected activity at zero boundary defects, with $\sup_D \sum_{X\ni D} e^{A_0|X|}\|K^0(X)\|_{r,*} \le C\tau_r$. Assume the single-label reference factors have uniformly bounded norm in the same reserve. Suppose the first marked terms obtained by Cauchy extraction on the boundary-link discs have bounds*

$$\begin{aligned}\|D_{X,e}\|_{r,*} &\le CR^{-4}\tau_r e^{-A_0|X|}e^{-ad_\partial(B,C)/R}, \qquad e=\{B,C\}\subset X,\\ &\text{where} \qquad d_\partial(B,C)=\operatorname{dist}(B,C)+\min(\delta(B),\delta(C)).\end{aligned} \tag{4.54}$$

The norm includes the spare field reserve, X is graph connected, and $\delta(B)$ is its distance to the physical boundary. For every fixed output exponent A_1, take A_0 and b_0/R sufficiently large and τ_r sufficiently small. There are connected covers $X\cup J_e$ reaching the grouped boundary layer with

$$\sup_D \sum_{X,e:\ D\in X\cup J_e} e^{A_1|X\cup J_e|}\|D_{X,e}\|_{r,*} \le C_{A_1}R^{-4}\tau_r \tag{4.55}$$

uniformly in the size of the cube. Regrouping the circle product on these covers preserves the field density. The activity of components with at least one mark has anchored norm at most $C_{A_1}R^{-4}\tau_r$. On components avoiding the boundary labels the activity equals K^0.

Proof. If $K_X(\mathbf{z})$ is the interpolated activity with physical defects $\mathbf{z}$, integration along $t\mathbf{z}$ gives $K(X)=K^0(X)+\sum_e D_{X,e}$, where $D_{X,e}=\int_0^1 z_e\,\partial_e K_X(t\mathbf{z})\,dt$. Attach to e a path from one of its cells to a nearest face and, if needed, a path between its endpoints. These paths form J_e. They use at most $C(1+d_\partial(B,C)/b_0)$ cells, and the distance in the cell graph from either endpoint to any cell of J_e obeys the same bound.

For a root $D\in X$, choose X rooted at D and sum over at most $|X|^2$ marks, dropping the remaining defect weight. The connected set count through D is at most a fixed exponential in $|X|$. For $D\in J_e$, both endpoints lie within $C(1+k)$ grid units of D if $k=\lceil d_\partial(B,C)/b_0\rceil$. There are at most $C(1+k)^{2d}$ choices. For each pair sum over connected X rooted at B. The cover size weight uses at most $CA_1(1+k)$ from the exponential decay in the distance, which leaves a convergent sum in k.

Expand the original circle product into packings of its nonadjacent connected actual supports. Each occupied support is now assigned either to $K^0(X)$ with cover X or to its selected marked term with cover $X\cup J_e$. Group the covers that overlap or touch and insert the old reference factors on their unused labels. This is an exact regrouping because the actual occupied supports are disjoint, independently of their covers, and the filler labels are exactly the unused labels consumed by grouping. The field reserve multiplies on the disjoint actual labels, even if their tested neighborhoods intersect. A grouped component having no mark is just one unmarked activity, since the original actual supports are mutually nonadjacent.

For completeness, the local norm bound for the regrouping follows from a rooted tree sum. Absorb filler constants and the desired output size weight into the spare exponents of the cover weights. A parent cover Y can attach each child at one of $O(|Y|)$ possible labels, the labels in its

closed graph neighborhood. A connected family containing a given root has a spanning tree of touching covers, and may be overcounted by specifying an ordered list of children at each such label. If u bounds the resulting total weight of descendant trees rooted at any label, the children's lists cost at most $(1-u)^{-C|Y|}$. A fixed spare size exponent pays this factor for sufficiently small u. Induction on the tree height bounds u by $C\tau_r$ from the cover-anchored input. To bound the subseries with at least one mark, multiply all marked weights by R^4 in this positive majorant. Its anchored input is still $O(\tau_r)$, so the same induction applies. Returning to the original weights gives the gain $O(R^{-4}\tau_r)$ for the marked subseries. All its output covers reach the boundary, so components avoiding the boundary labels retain their unmarked activity. □

5 Physical initial data and terminal covariance

To apply the one-step estimates, the effective interaction at the initial scale must represent the spin measure at the chosen temperature. We construct the initial polymer activity with the mass left free, then identify the mass at which the representation is physical. In dimension four the same construction leaves the kinetic coefficient free as well.

5.1 Joint initialization in the cube chart

To initialize the bulk and boundary charts in the same activity norm, we must account for the two truncations in the microscopic density. The pressure terms use the full inverse of the short filter, whereas the reference on unused labels uses its finite Neumann truncation. We keep this reference fixed and expand the filter tail together with the quadratic tail. Recording both cell labels of each filter entry, including the column label, permits an exact comparison with the bulk off the grouped boundary labels.

Proposition 5.1 (Joint initial activity and its mass derivative)**.** *Use initial cells of side comparable to $b_0 = RM_R$, with M_R a sufficiently large fixed multiple of* $\log R$*. Use fixed enlarged cell stencils and the raw field tests F_{b_0} of Section 3. On the boundary use the same tests on the clipped stencils. Clipped cells can be joined to adjacent cells so that the resulting labels have volume comparable to b_0^d. Let T be a finite Neumann truncation of the frozen primary filter $H_F = B_\Lambda^{-1}$. Decompose the primary-field Hessian into its bulk scalar row, a finite head, and the remainder, as specified below. There is an exact circle-product initialization with the quartic reference at T and the finite quadratic head. Its activity $K_{\mathrm{init}}(\beta,\nu)$ is equal to the corresponding bulk activity on every polymer avoiding the grouped boundary labels.*

Set $g = g_{\mathrm{mic}}$, $G_0 \asymp A_{\mathrm{sc}}^4 g$, and

$$U = A_{\mathrm{sc}}^6 g^2 + G_0^{9/4} g^{1/2} + G_0^3. \tag{5.1}$$

At the two initial radii, use the growing regulator and the decaying and quotient regulator, respectively. With any fixed exponential polymer weight and a fixed reserve in both regulators, the mixed activity norm obeys

$$\sup_{|\nu| \le cG_0} \|K_{\mathrm{init}}(\beta,\nu)\|_{\mathrm{mix},0} \le CU. \tag{5.2}$$

The activity is holomorphic in this disc, and on a fixed smaller disc its mass derivative is at most CU/G_0 in the same norm. The estimates are uniform in the volume and in a fixed physical window $|\beta - 1| \le A_\beta g$ for $n \ge n_0(R)$.

Proof. Filter and Hessian tails. All calculations of filter kernels below are in the physical field $h = A_{\mathrm{sc}} R^{-d/4}\phi$. Write $\delta(x) = \operatorname{dist}(x, \Lambda^c)$. The Neumann expansion, the short-range row gap and the weighted Hessian estimate of Lemma 2.2 give

$$H_F = T + E^0 + E^\partial, \tag{5.3}$$

$$\sup_x \sum_y e^{a|x-y|/R}(|E^0(x,y)| + e^{a\delta(x)/R}|E^\partial(x,y)|) \le \epsilon_R. \tag{5.4}$$

The same estimate holds with rows and columns exchanged. Here E^0 is the restriction of the bulk Neumann tail. In the difference E^∂ every walk has length exceeding the truncation order and visits the exterior. The geometric row bound therefore pays both the truncation and the extra boundary weight after reducing $a > 0$. The truncation order and M_R can be increased by fixed multiples of $\log R$ so that $\epsilon_R \le R^{-M}$ for any fixed M, with all finite stencils still in a fixed number of cell collars.

Remove the bulk scalar row from the primary-field Hessian. Write its bulk zero-row kernel Z^0 in difference-square form $-\frac{1}{4}\sum_{x,y} Z^0(x,y)(h_x - h_y)^2$, summing over the actual cube sites. The correction to the diagonal is then the missing row $-\sum_{y\notin\Lambda} Z^0(x,y)$. Include it with the finite-volume Hessian discrepancy. The bulk kernel has exponentially weighted absolute rows and columns of size Cg. The boundary discrepancy, including the missing row, has the same bound with the extra factor $e^{a\delta(x)/R}$ in (5.4), with ϵ_R replaced by Cg. For the Hessian pullback this follows also by expanding the two conjugating factors in short walks. In a discrepancy one factor is a Hessian discrepancy or a walk reaches an exterior summation site. In either case the weighted bound and $\delta(x) \le \delta(y) + |x-y|$ retain the boundary weight at the initial site. Take the bulk head for $|x-y| \le c_1 b_0$ and the boundary head for $|x-y| + \min(\delta(x), \delta(y)) \le c_1 b_0$, for a fixed $c_1 > 0$. Assign these heads by fixed symmetric local weights to their incident labels, using grouped boundary labels for the boundary head. The remainder has the same estimates as the original kernels with bound $Cge^{-c_2 M_R}$ after reducing a. The diagonal correction belongs entirely to the boundary part. These conventions leave the bulk pair density unchanged when both endpoints are present.

Interpolation of cell pairs. For a cell pair $e = (B, D)$, let $d_e = \operatorname{dist}(B, D)$. For boundary terms use $d_e^\partial = d_e + \min(\delta(B), \delta(D))$, where $\delta(B) = \min_{x\in B}\delta(x)$. Choose a tail variable for each ordered block of E^0 or E^∂, and independently for each unordered cell pair of the bulk or boundary quadratic remainder, including the diagonal pairs. Each variable multiplies its block or its exact signed pair contribution. Write $D_e = d_e$ or $D_e = d_e^\partial$ according to the type of variable. For a small fixed $c_3 > 0$ use two polydiscs, with radii

$$\rho_e^+ = \epsilon^{-1/4} e^{c_3 D_e/R}, \qquad \rho_e^- = \epsilon^{-1/8} e^{c_3 D_e/R}, \tag{5.5}$$

where $\epsilon = R^{-M}$. By increasing M_R and the Neumann cutoff before fixing the cell collars, we can take both weighted tail bounds at most ϵ. Hence on the larger polydisc the common absolute majorant of the amplified tails has weighted row and column bounds $C\epsilon^{3/4}$. This applies also when several types of variables have the same pair of cells. For a quadratic difference-square term the estimates control the absolute pair coefficient, and hence the diagonal contribution as well as the off-diagonal contribution.

Apply the cell BKAR formula to the Hessian-subtracted short pressure. On each nonsingleton tree component C, integrate the tree derivative on its induced short graph, evaluated on the filter rows of C. On a singleton B, subtract the fixed quartic filler $I_{4,T}(B) = \exp[-V_4((Th)_B)]$ from the exact singleton pressure weight. Denote these two types of pressure atom by p_C. Since the Hessian-subtracted pressure weight is a quadratic exponential times a finite spin partition function divided by its zero-field value, it is entire in the external field. Thus the pressure formula holds for complex tail variables, with p_C holomorphic. In the pressure packing the marked atoms are disjoint,

and all its other labels carry the quartic fillers. The finite quadratic head factor on a label is present whether or not the label is marked.

We first estimate the product for a fixed packing, restricted to a set of labels W that contains all its marked atoms and both endpoints of every selected tail variable. Other tail variables are zero. Insert the filler, with its finite head, at every unused label of W. Let F_W be this product times the exponential of the selected quadratic tails. Denote the two physical radii by

$$e = A_{\mathrm{sc}} R^{-d/4} b_0^{-d/4}, \qquad a = b_0^{-d/4}. \tag{5.6}$$

In the remainder of this proof r denotes either physical radius. Define $\tau_r = b_0^d(r^6 + gr^4)$. We use normalized squares

$$Z_B^2 = \frac{\|h\|_{2,B}^2}{b_0^d r^2}, \qquad Q_B^2 = \frac{q_{B^*}(h)^2}{r^2}, \tag{5.7}$$

where q_{B^*} is the adapted affine quotient on the fixed stencil B^*. Its defining affine decomposition and the finite-dimensional affine norm comparison on B give

$$\frac{\sup_{B^*} |h|^2}{r^2} \le C(Z_B^2 + Q_B^2). \tag{5.8}$$

Indeed the residual is controlled in supremum by q_{B^*}, and the L^2 norm of the affine part on B controls its supremum on B^*. These comparisons are uniform for a clipped label joined to a neighbour because it contains a comparable cube and is contained in a fixed enlargement of that cube.

For real tail variables in the larger polydisc all marked rows use T plus an error with absolute row and column bounds $C\epsilon^{3/4}$, and the fillers use T. The same is true of any convex mixture of those rows, with the same absolute majorant. For x in W at distance more than the range of B_Λ from the artificial interfaces of W, the identity $B_\Lambda T = I - Q_\Lambda^{m+1}$ recovers h_x using only those rows that are available in W. The operator norm error is $C\epsilon^{3/4}$ by the row and column bounds. Its input consists of the fixed stencils and the selected column labels. The omitted strip in any cell has volume $O((R/b_0)b_0^d)$. Squaring, summing and using (5.8) in these strips and for the stencil input gives, for the available filtered rows u,

$$\sum_{B \in W} Z_B^2 \le C \frac{\|u\|_{2,W}^2}{b_0^d r^2} + C(R/b_0 + \epsilon^{3/2}) \sum_{B \in W} (Z_B^2 + Q_B^2). \tag{5.9}$$

Here m is the order of the truncation. Physical-face rows are handled by the restricted short inverse and need no extension of h through the physical face.

Lemma 2.5 and Proposition 2.6 give, at a real filter, damping by $e^{-c\sum \Psi(u_x)}$ on all the marked rows with a factor τ_r per marked atom. One may take an arbitrarily small additional loss in the normalized cell supremum squares, and any fixed exponential size margin on the nonsingleton atoms. To use the fixed singleton filler, split the singleton difference at $e^{-V_4(u)}$. The difference of the remaining quartic exponentials, by integration along their segment, costs at most $C\epsilon^{3/4} b_0^d r^4$ times a fixed polynomial in the normalized cell suprema. This is smaller than τ_r. The exponential in the segment formula damps a convex combination of the two row fields, by convexity of the quartic. The estimates through a fixed number of field jets hold in the same way because filtered test directions have supremum at most Cr. When the derivative falls on the outer finite head or on the selected quadratic exponential, its polynomial factors have quadratic coefficient bounded by $C(g + \epsilon^{3/4})$. At either radius

$$b_0^d r^2 (g + \epsilon^{3/4}) = o(1). \tag{5.10}$$

Both the finite heads and the tail exponents therefore consume only an arbitrarily small fixed multiple of the sum in (5.8). The cell supremum penalties of the pressure estimates have the same property, by the ℓ^2 bound on the matrix of cell suprema of the common filter majorant. That matrix has bounded row and column sums: the finite head has fixed cell range, and the tails have an exponentially weighted row and column bound.

At the large radius $r = a$, $b_0^d r^4 = 1$. For every fixed $K > 0$ and all sufficiently large R, $c\Psi(t) \geq Kt^2/(b_0^d r^2) - C_K/b_0^d$. For $|t| \leq 1$ this is quartic Young's inequality, and for $|t| > 1$ it follows from $\Psi(t) \geq t^2/2$. Take K to leave a strict reserve in the required negative Z_B^2 exponents. In (5.9) the coefficient of Q_B^2 then tends to zero and fits in any prescribed positive quotient allowance. At the small radius one keeps a fraction of the growing regulator directly, using (5.8) and discarding the real damping. The finite jet Leibniz coefficients, the singleton polynomial factors and the fixed losses per cell are bounded by a fixed exponential in $|W|$. Thus, in either fixed weighted jet norm with a reserve, denoted by $\|\cdot\|_{r,*}$, the real bound takes the form

$$\|F_W\|_{r,*} \leq C^{|W|} \prod_{C \text{ marked}} (\tau_r e^{-A_1|C|}). \tag{5.11}$$

Here A_1 is arbitrary. On singleton atoms its fixed cost can be included in $C^{|W|}$, or retained as a separate fixed factor per singleton.

For use in Cauchy's formula we also need a bound on the large complex polydisc. For a pressure graph U_C, with J the interpolated short coupling, set $A_J = L_J^{\mathsf{T}} S_J L_J$. The proof of the weighted Hessian estimate applies to small fixed complex perturbations of the coupling variables around the tree interpolants: the absolute row gap and the bound on the largest coefficient are preserved. Indeed the proof of that estimate only uses the absolute weighted gap of the complex hyperbolic-tangent matrix and the small maximum entry. It gives $A_J = I - J + O(g)$ in absolute rows and columns. Integrating the zero-field coupling derivative from zero along the segment tJ gives $|\log(Z_J(0)/2^{|U_C|})| \leq C|U_C|g$, since the derivative uses only off-diagonal spin covariances and each is $O(g)$. For $u = v + iw$, positivity of the real spin weight gives $|Z_J(L_J u)| \leq Z_{\operatorname{Re} J}(\operatorname{Re} L_J u)$. Compare its field to $L_{\operatorname{Re} J} v$, using the bound one on every coordinate of the real pressure gradient. The difference costs $C\|w\|_1$. The quadratic subtractions differ in absolute value by $C(g\|v\|_2^2 + \|w\|_2^2 + \|v\|_2\|w\|_2)$ using $\operatorname{Re}(I - J) = I - \operatorname{Re} J$. The complex pressure weight is consequently bounded by the real Ψ damping at v times the exponential of these costs and $C|U_C|g$. For the quartic filler use $\operatorname{Re}(x + iy)^4 \geq \frac{1}{2}x^4 - Cy^4$ and $|y| \leq Cr$. For a real displacement v_1, the inequality $\Psi(v_0 + v_1) \geq c\Psi(v_0) - C\Psi(v_1)$ follows from the triangle inequality and the bounded ratio $\Psi(2t)/\Psi(t)$. Thus displacing the real field by at most Cr consumes only a fixed multiple of $b_0^d r^4$ per label.

For the finitely many tested field jets, evaluate every factor on $h + \sum_i z_i v_i$ with $|z_i| \leq 1$ and $\|v_i\|_\infty \leq Cr$ on the stencils. The imaginary part has ℓ^2 norm at most $C\epsilon^{3/4}$ times the norm of h on the stencils and selected columns, plus $Cr\sqrt{b_0^d|W|}$; the same bound controls the displacement between v and the real row substitution at h, with the tail term omitted for that displacement. Sum over the atoms before using Young's inequality, as their row labels are disjoint. The part quadratic in the background uses a coefficient $C(g + \epsilon^{3/4})$ or an arbitrarily small fixed multiple of the normalized squares in (5.8). Young's inequality for the remaining terms costs at most $Cb_0^{2d}|W|$. Thus the real recovery bound (5.9) at the real row substitution still leaves the prescribed reserves. Cauchy's formula on fixed coupling circles costs at most a constant per tree edge, and there are at most a fixed exponential number of trees. Cauchy's formula on the field-jet circles costs a fixed factor. We obtain the uniform bound $e^{Cb_0^{2d}|W|}$ in $\|\cdot\|_{r,*}$ on the large polydisc, with constants independent of the number of tail variables.

For a point $\boldsymbol{z}$ of the small polydisc consider the common scalar line $\operatorname{Re}\boldsymbol{z} + w \operatorname{Im}\boldsymbol{z}$. It stays in the larger polydisc for $|w| \leq c\epsilon^{-1/8}$ and has real entries on the diameter. At $w = i$ the semicircle

has harmonic measure $O(\epsilon^{1/8})$. Apply the two-constants theorem to each fixed normalized field-jet evaluation and take the supremum. For M sufficiently large, $\epsilon^{1/8} b_0^{2d} = o(1)$ and $\epsilon^{1/8} |\log \tau_r| = o(1)$. The bound (5.11) follows on the small polydisc with a different C and half the size margin. Thus taking any fixed desired size margin is possible also there.

We give the exact support expansion and its count. For a fixed packing and for a subset S of the tail variables, let $\Delta_S F = \sum_{D \subset S} (-1)^{|S|-|D|} F(\mathbf{1}_D)$. The evaluation at all variables one is $\sum_S \Delta_S F$, including $S = \varnothing$. In a finite volume this is an identity of finite sums. The term uses precisely the Taylor monomials with positive exponent in each variable of S and zero exponents outside S. By Cauchy's formula and the small polydisc bound each selected variable contributes a factor at most

$$\sum_{k \geq 1} (\rho_e^-)^{-k} \leq 2\epsilon^{1/8} e^{-c_3 D_e / R}. \tag{5.12}$$

Associate a connected grid path to the pair of each selected variable, and in a boundary variable also include a path to its selected physical face. The supports used to form connected components are the actual marked atom sets and these paths. Group touching as well as overlapping supports, and include a filler for every otherwise unused label in the resulting component W. An atom reads only the selected filter columns connected to it by these paths and its fixed finite head stencil. The latter need not be added to the polymer labels, since field dependence on these stencils is allowed for the reference and for every activity. As functions of the selected variables, different components factor exactly. This is also true of Δ_S, since the variable sets in different components are disjoint. An unused label outside the components keeps its original reference. Conversely all pairwise nonadjacent choices of output components and their internal data give precisely the corresponding terms of the product. This constructs the circle product on the enlarged support graph without any division by a reference factor. In particular, at zero adjustable mass,

$$e^{-Q^{\mathrm{head}} - Q^{\mathrm{tail}}} e^{-\mathcal{U}_\Lambda(H_F h)} = (I_T \circ K_{\mathrm{init}}(\beta, 0))(\Lambda). \tag{5.13}$$

Sum over connected components. Regard each marked atom or selected edge with its path as an object with a

connected cover Y. Include the fixed $C^{|W|}$ and any desired output polymer size weight in the per-object exponent $e^{c|Y|}$, using $|W| \leq \sum |Y|$. The sum of atom weights through one root label is $O(\tau_r)$, by choosing its size margin larger than the exponential number of connected sets through that label. For each edge type the anchored sum through one root is $O(\epsilon^{1/8})$ after any fixed size margin, by the connected-cover count of Section 3. For completeness, convergence of the connected-object sum follows by choosing a rooted spanning tree of touching covers. Root at an object covering the given label, and assign each child to the first label in the closed neighbourhood of its parent cover that belongs to the child. There are at most $C|Y|$ possible attachment labels at a parent of cover Y. Allow an ordered list of children at each attachment label. If the anchored sum of all trees of height at most $h-1$ is T_{h-1}, the extra factor at this parent is bounded by $(1 - T_{h-1})^{-C|Y|}$. Reserve an additional fixed $e^{|Y|}$ in the input weights and take the anchored sum v of the inflated object weights small. Induction gives $T_h \leq 2v$ for every h, since $(1 - 2v)^{-C|Y|} \leq e^{|Y|}$. Every connected family is counted at least once. The resulting rooted activity bound is $C(\tau_r + \epsilon^{1/8})$, with the same field reserve. The bound is uniform in the cube size. Norm tests on a component restrict to all the stencils used above, and the estimate retained its reserves on the disjoint actual labels. In particular the counting step incurs no extra field cost.

A component of the exact activity that avoids boundary labels contains no boundary variable, because such a selected variable was assigned a face-reaching path. On interior rows its finite Neumann head is the bulk head. Its pressure atoms use the same induced short graph as the

corresponding bulk atoms, and its bulk difference-square tails are exactly the same as in the bulk calculation. Use the same path assignments for bulk pairs and average over tied paths or faces when enforcing lattice symmetry. The ordinary paths can be taken coordinate-monotone, so a path between two cells of the cube stays in the cube. No bulk coefficient on a boundary-free polymer changes due to clipping, because its local stencils and all retained labels lie outside the clipping collar. Thus the boundary-free coefficients agree term by term. The bulk activity on a fixed finite polymer is defined by these same finite component formulas.

Finally give each label the factor $e^{-\nu m(B)}$, where $m(B) = \frac{1}{2} b_0^{-s} \sum_{x \in B} \phi_x^2$ is the normalized scalar mass density. The product over an occupied component is exactly $e^{-\nu m(W)}$, by additivity, including the fillers in that component. The reweighting preserves the circle product and the bulk agreement. Its quadratic cost in the two radii, for $|\nu| \le c G_0$, is $O(G_0)$ and $O(G_0^{1/2})$, respectively, and fits inside the retained reserve. It therefore defines a holomorphic family in the fixed norm. Shrinking the mass disc gives the claimed derivative by Cauchy's formula. Since

$$\tau_e \le C A_{\mathrm{sc}}^6 g^2, \qquad \tau_a \le C g^{1/2}, \tag{5.14}$$

and $\epsilon^{1/8}$ can be made smaller than G_0^3, the weighted sum of the two estimates proves (5.2). All bounds used here persist on the real thermal window. For fixed finite volume the same family is continuous in β in the activity norm. The finite coefficient formulas are continuous on bounded field sets, uniformly over the tested directions. Outside a bounded field set the strict regulator reserve makes the norm contribution uniformly small, by (5.8). As in the pressure estimates, for each fixed window width the quartic defect proportional to that width is absorbed by the sextic seed when R is sufficiently large. □

Corollary 5.2 (Physical entry into the mass chart)**.** *For $d = 1, 2, 3, 4$, take the finite heads of Proposition 2.16 and the joint activity of Proposition 5.1. Fix A_0 and $|\beta - 1| \le A_0 g_{\mathrm{mic}}$. For sufficiently large R and $n \ge n_0(R)$, their circle product equals the physical density at the scalar mass $\nu_0^{\mathrm{bare}} = \lambda^{-1} M_R^s r_h(\beta)$. This mass is $O_{A_0}(\Lambda_R G_0)$ and lies on the smaller mass disc of the proposition. In the units of Assumption 6.1, $U/(\Lambda_R G_0) = o(1)$ and $U/G_0^2 = o(1)$ as $R \to \infty$.*

Proof. At the stated mass, the reference factor and the activity reweighting together restore the scalar term subtracted from the physical Hessian. At other ν the same circle product gives the off-shell family. The crossing estimate in Section 2.1 gives the bound on ν_0^{bare}, and $\Lambda_R = o(1)$ places it inside the smaller disc. Substituting the microscopic scales into U gives

$$\frac{U}{\Lambda_R G_0} \le \frac{C}{M_R^s} \left(g_{\mathrm{mic}}^{1-4a_*} + g_{\mathrm{mic}}^{7/4-7a_*} + g_{\mathrm{mic}}^{2-10a_*} \right) = o(1), \qquad \frac{U}{G_0^2} \le C \left(g_{\mathrm{mic}}^{2a_*} + g_{\mathrm{mic}}^{3/4-a_*} + G_0 \right) = o(1). \tag{5.15}$$

Together with (5.2) and the derivative bound of Proposition 5.1, these ratios give the initial activity and mass-derivative estimates in the shooting chart. □

5.2 Initial and terminal data for the kinetic reference family

The shooting construction uses two endpoint properties of the kinetic reference family. At the initial end, matching both bare coordinates must recover the microscopic density. At the terminal end, changing the reference must preserve the Gaussian estimates for curvature and magnetic response. The two lemmas verify these properties before the nonlinear matching.

Exact physical entry with a free initial kinetic coordinate. We first extract the kinetic moment from the bulk quadratic and keep the full boundary quadratic fixed. This determines the variation of the bare kinetic coordinate while confining the central difference stencil to the ordinary region.

Lemma 5.3 (Initial kinetic extraction and off-shell reweighting)**.** *Let $d = 4$, $1/2 < p < 1$, and use Proposition 4.4 with $|\delta| \le C\lambda^2 G_0$. After the bulk row is removed, let z_0^{bare} be the k^2 coefficient of the retained bulk quadratic head in Hessian units. Then $|z_0^{\mathrm{bare}}| \le C\lambda G_0$. The remainder of its ordinary bulk head has zero row and zero second moments and absolute moments of degrees two and four bounded respectively by $C\lambda G_0$ and $C\lambda G_0 R^2$. The full physical collar quadratic, including the interface terms from this extraction, has norm $C\lambda G_0$ at radius ℓ_{b_0}.*

There is an exact initial activity in the outgoing chart with independent mass and kinetic coordinates ν, z, holomorphic on $|\nu| \le cG_0$, $|z - z_0^{\mathrm{bare}}| \le cG_0$, with norm CU. On a fixed smaller disc both partial derivatives are at most CU/G_0. It equals the physical interaction at the bare mass and kinetic coordinates and agrees with the bulk activity off the grouped boundary labels. The initial kinetic coefficient has uniformly negative derivative in the scaled external parameter $a = \delta/\lambda$, bounded away from zero for sufficiently large R.

Proof. The canonical weighted row of the zero-row bulk head is bounded by $C\lambda G_0 R^{-2}$: combine Proposition 2.16 with the density-change bound $Ca_R^{-1}|\delta|$ in Proposition 4.4. Its second and fourth absolute moments are consequently bounded by $C\lambda G_0$ and $C\lambda G_0 R^2$. Keep the finite head in anchored difference-square form. Subtract its scalar second moment times a unit-curvature difference-square representative with spans $2e_i$. The cut head is cubically symmetric, so this removes all its second moments. The subtraction preserves zero row and satisfies the same absolute moment bounds. The starting coefficient has size $C\lambda G_0$.

For exact placement, separate the ordinary sites S from the initial collar so that their stencils are in the bulk-comparison region and the interface is at depth at least a fixed positive multiple of b_0. At those sites only, replace the kinetic difference-square density by k_x^{a} of Lemma 4.1, using the same normalization for the potential and its Hessian. The difference of the two unit-curvature representatives is a fixed multiple of

$$\sum_{i=1}^{4} (\phi_{x+2e_i}^2 - 2\phi_x^2 + \phi_{x-2e_i}^2). \tag{5.16}$$

For example the representative $\frac{1}{8}\sum_i[(\phi_{x+2e_i} - \phi_x)^2 + (\phi_{x-2e_i} - \phi_x)^2]$ differs from $-\phi_x \Delta_{\mathrm{c}} \phi_x$ by $1/8$ times this display. Summing over S cancels every internal bond and leaves the crossing-bond differences of ϕ^2. On the interface the raw first difference bound gives $|\phi_y^2 - \phi_x^2| \le C\ell_{b_0}^2/b_0$ for tested directions at radius ℓ_{b_0}; the bilinear bound follows by polarization. At most Cb_0^3 bonds contribute per interface label. Since $\ell_{b_0}^2 = b_0^{-2}$, the complete interface tensor has normalized size $C|z_0^{\mathrm{bare}}|$. Assign its actual crossing stencils to the collar. The entire original head in the inner collar is retained there, including its clipping defect. Its norm has the claimed bound by Proposition 2.16 for $p > 1/2$. No ordinary nearest-neighbour density is tested at the physical face.

Denote by Δ_i the changes in local quadratic exponents from this rearrangement, including the interface density. They have zero total as a polynomial. On each affected label their jet norm at ℓ_{b_0} is at most $C\lambda G_0$. On an originally occupied component X, multiply its activity by $e^{\sum_{i \in X} \Delta_i}$ with the convention that the old uncovered exponent is the new one minus Δ_i. This is precisely the finite circle-product identity used for the extraction divergences. The same rearrangement on the bulk comparison uses the translation-invariant divergence with no interface density. Hence the transformed activities still agree on all polymers avoiding the boundary labels.

Now allow z to differ from z_0^{bare}. Multiply on each ordinary label by the exponential of the ordinary kinetic density with coefficient $-(z - z_0^{\mathrm{bare}})$, in the potential normalization. Put this same factor on every site label in an occupied component, as well as on each uncovered label. This defines the off-shell family exactly, and keeps the raw collar tensor fixed. The off-shell mass is introduced by the uniform scalar multiplication in Proposition 5.1. Any of these multipliers of quadratic coefficient size A at the small radius has large-radius exponent bounded by

$$CAG_0^{-1/2} \sum_i \left(Z_i^2 + Q_i^2 \right), \tag{5.17}$$

where Z_i, Q_i are the local normalized L^2 and quotient quantities. This follows by testing the quadratic on its fixed enlarged stencil and using $F_{b_0}(\phi)/h_{b_0} \le C(Z_i + Q_i)$. The same bound at the small radius uses A instead of $AG_0^{-1/2}$. For the rearrangement take $A = C\lambda G_0$, and for the off-shell multipliers take $A = CG_0$. These costs and the fixed number of tested jets use only a fixed fraction of the spare regulator and size margins of the joint initialization. Thus the norm remains CU on the two complex discs. Cauchy's inequality gives the two derivative bounds. The equality at the bare coordinates is exact.

Finally the derivative in $a = \delta/\lambda$ of the bulk kinetic coefficient of the full density difference is negative and bounded away from zero by Proposition 4.4. Its discarded tail has derivative in the same kinetic moment at most $Ce^{-cb_0/R}$, by the weighted short inverse estimate for one δ derivative and multiplication by λ. The old microscopic quadratic is independent of a. Hence the same sign and lower bound hold for the bare coordinate of the cut head. All exact rearrangements and initial norms are continuous on the physical temperature and counterterm windows. □

For the terminal comparison we must control both the remaining covariance in the principal coordinate and the magnetic response of the slices already integrated. The transfer formula gives both covariances in terms of the same unperturbed decomposition.

Lemma 5.4 (Terminal Gaussian bounds under covariance transfer)**.** *Use the transfer decomposition of Section 4.3, with the same scale horizon as its unperturbed decomposition. Denote the unperturbed terminal and integrated covariances by $C_T^0, C_<^0$. There is a common choice of the dual principal direction ℓ, up to scalar normalization, for all the references. Let T_δ, P_δ denote their terminal and integrated covariances. Uniformly in the physical counterterm window, the terminal variance, conditioned Gaussian exponential moment, magnetic profile, and integrated magnetic direction satisfy the bounds used in Propositions 7.3 and 7.4.*

Proof. Write q_T for the unperturbed auxiliary terminal kernel in spectral variables and

$$M = a_R \mathcal{O} \mathcal{Q}_{c',0} q_T(\mathcal{Q}_{c',0}) \mathcal{O}^*. \tag{5.18}$$

Since $0 \le q_T(t) \le t^{-1}$ on the positive spectrum, one has $0 \le M \le a_R A_\Lambda$ and $M \le \|\mathcal{Q}_{c',0}\| C_T^0$. The transferred covariances are exactly

$$T_\delta = t_\delta(C_T^0 - \eta M), \qquad P_\delta = t_\delta C_<^0 + (1 - t_\delta) a_R A_\Lambda + t_\delta \eta M. \tag{5.19}$$

In particular at $\delta = 0$ the transfer moves ηM from the prescribed terminal covariance into the integrated covariance. The comparison $\frac{1}{2} t_\delta C_T^0 \le T_\delta \le t_\delta C_T^0$ follows from the multiplier bound. The generalized principal vector solves $J_\Lambda \ell = \xi_1 A_\Lambda \ell$, independently of δ. Its finite-volume normalizations are uniformly equivalent since $B_\Lambda^\delta \asymp I$. Use one and the same vector in the comparisons below.

For $\sigma_0^2 = \ell^{\mathsf{T}} C_T^0 \ell$ and $u_0 = C_T^0 \ell / \sigma_0^2$ the old bounds give $\sigma_0^2 \asymp n^2/(\lambda x_N)$ and $\|u_0\|_{F_{n,p}} \le C n^{-2}$. If D is a weighted stencil test in $F_{n,p}$, the unperturbed hard covariance bound and the decomposition along ℓ imply

$$DC_T^0 D^* \le C(\lambda n^{-2} + \sigma_0^2 n^{-4}). \tag{5.20}$$

By covariance Cauchy–Schwarz and $M \le a_R A_\Lambda$,

$$|DM\ell| \le C\sqrt{\lambda} R(\sqrt{\lambda} n^{-1} + \sigma_0 n^{-2}). \tag{5.21}$$

It follows that the variance of T_δ along ℓ has the same order as σ_0^2. Dividing the last display by σ_0^2 and accounting for the change in the denominator shows $\|u_\delta - u_0\|_{F_{n,p}} = o(n^{-2})$ as $n \to \infty$ at fixed R and terminal aspect. In particular the magnetic overlap has the same lower bound. The conditioned covariance is monotone under positive covariance comparison, by the infimum formula in the proof of Lemma C.9. Its Gaussian norm moment follows by adding an independent Gaussian to attain a fixed multiple of the old conditioned covariance and applying conditional Jensen's inequality.

For the integrated magnetic direction, $\|A_\Lambda \mathbf{1}\|_{F_{n,p}} \le C_R n^p$. Outside a strip of microscopic width the field is constant, and inside the strip the local tests have only the boundary weight loss and microscopic differences. Also

$$\|M\mathbf{1}\|_{F_{n,p}} \le C\sqrt{\lambda} R n^2 (\sqrt{\lambda} n^{-1} + \sigma_0 n^{-2}) \le C_R n^2. \tag{5.22}$$

For the last inequality use $x_N^{-1} = O_R(1 + \log n)$. Together with the old bound for $C_<^0 \mathbf{1}$, the exact formula for P_δ proves the integrated-source estimate. All these covariances and normalizations are continuous in the real external parameter. □

6 A dimension-independent shooting theorem

The linearized bulk mass grows by a factor L^s at each step, whereas the quartic coupling decreases only logarithmically. We therefore prescribe the quartic coupling and the stable coordinates at the microscopic scale and determine the mass from the zero-curvature condition at scale N. Theorem 6.2 implements these two boundary conditions by a contraction on paths. Its proof allows the spaces of boundary tensors and polymer activities to vary with the scale.

Let $b_j = b_0 L^j$ be the step scales, $N \asymp \log n$, $x_0 = G_0 \asymp A_{\mathrm{sc}}^4 g_{\mathrm{mic}}$, and set

$$x_{j+1} = x_j - \mathfrak{b}_j x_j^2, \qquad f_j = x_j^3 + U d_0^j, \qquad 0 < d_0 \le \tfrac{1}{4}. \tag{6.1}$$

Assume $0 \le \mathfrak{b}_j \le C\lambda^2$ and $\mathfrak{b}_j \ge c\lambda^2$ except on a bounded number of initial or terminal steps. The merged first step may instead have $\mathfrak{b}_0 \le C\lambda^2 \log(2 + M_R)$, where $M_R = b_0/R \asymp \log R$. Assume $\sup_j \mathfrak{b}_j x_j = o(1)$ and

$$\Lambda_R = C_0 \lambda M_R^s, \qquad x_0/\lambda^2 + U/x_0 + \Lambda_R = o(1). \tag{6.2}$$

For the physical initialization of Proposition 5.1, $x_0/\lambda^3 = O(g_{\mathrm{mic}}^{1-10a_*}) \to 0$ and

$$U = A_{\mathrm{sc}}^6 g_{\mathrm{mic}}^2 + G_0^{9/4} g_{\mathrm{mic}}^{1/2} + G_0^3 = o(G_0^{9/4}). \tag{6.3}$$

The latter follows since the first term divided by $G_0^{9/4}$ is $O(g_{\mathrm{mic}}^{3a_* - 1/4})$. Thus the microscopic estimate has the smallness required at entry into the map.

Write $\gamma_j = (g_j - x_j)/x_j$, $\mu_j = \nu_j/(\Lambda_R x_j)$, and $k_j = K_j/f_j$, and let w_j denote W_j expressed in its scale-dependent stable weights. The one-step map must contract the stable coordinates in these

weights. We state its estimates on fixed enlarged convex boxes $|\gamma| \le \varepsilon$ and with fixed sufficiently large radii in w, μ, k. All derivatives below are in these normalized input coordinates, holding the frozen covariance fixed.

Assumption 6.1 (Weighted RG and endpoint data)**.** The step equations have the form

$$g^+ = g - \mathfrak{b}_j g^2 + e_j, \qquad |e_j| + \|De_j\| \le Cf_j, \tag{6.4}$$

$$\nu^+ = L^s \nu + \Lambda_R x_j M_j(\gamma, w, \mu, k), \tag{6.5}$$

$$w^+ = \mathcal{W}_j(\gamma, w, \mu, k), \qquad k^+ = \mathcal{K}_j(\gamma, w, \mu, k). \tag{6.6}$$

The maps $M_j, \mathcal{W}_j, \mathcal{K}_j$ are bounded at zero in the normalized input and output coordinates. For $\kappa < 1$ and $\delta_R = o(1)$ they have the derivative bounds

	γ	w	μ	k
$\mathcal{W}_j$	C	κ	δ_R	δ_R
M_j	C	C	δ_R	δ_R
$\mathcal{K}_j$	C	C	C	κ.

(6.7)

All constants are uniform in j, N. The fixed initial w_0 is bounded and $g_0 = x_0$. The initial family has $\|K_{\text{init}}\| \le CU$ and $\|D_{\nu_0} K_{\text{init}}\| \le CU/x_0$. The exact terminal curvature equation has a graph $\mu_N = B_N(\gamma_N, w_N, k_N)$ on these boxes, bounded at zero, with derivative bounds (C, C, δ_R). The maps and data are real on real arguments and continuous in the external data in the same charts. Uniform constants and the estimates hold on the boxes with the successive radii used in the proof below.

The term M_j includes the mass generated by Gaussian integration. With a fixed bulk filter of row sum one its actual first-order coefficient is $3b^s_{j+1}(\overline{T}\Gamma_j^{\text{bulk}}\overline{T}^*)(0)$. Centering at the actual slice keeps the Gaussian reference independent of the mass variable. At the terminal scale the natural profile in the curvature calculation is $\rho_N = n^{s/2} u_T$; the derivative of the uniform quadratic variation is $b_N^{-s}\|\rho_N\|_2^2 \asymp 1$ if $\|u_T\|_2 \asymp 1$. Proposition 7.3 derives the graph and its first variations from the terminal coefficient and activity bounds and a small inverse-chart derivative. At the other end of the path, the reference and the polymer packing must be evaluated at the same trial value of ν_0.

Theorem 6.2 (A trajectory with prescribed terminal curvature)**.** *For $d \in \{1, 2, 3, 4\}$, under Assumption 6.1 and the stated comparison-sequence bounds, there is a unique trajectory in the controlled box, for all sufficiently large R and every horizon with these bounds. It has $g_0 = x_0$, the prescribed W_0, $K_0 = K_{\text{init}}(\beta, \nu_0)$, zero terminal curvature, and*

$$|g_j - x_j| \le \varepsilon x_j, \qquad |\nu_j| \le C\Lambda_R x_j, \qquad \|w_j\| \le C, \qquad \|K_j\| \le Cf_j. \tag{6.8}$$

The constants and contraction factor are uniform in the horizon, and ν_0^* is continuous in the external data. For the canonical scales at large n,

$$g_N \asymp (G_0^{-1} + \lambda^2 N)^{-1} \asymp_{d,R} (\log V)^{-1}, \qquad \|K_N\| \le C(g_N^3 + U d_0^N). \tag{6.9}$$

Proof. Put $a_j = 1 - 2\mathfrak{b}_j x_j$. Since $0 \le a_j \le (x_{j+1}/x_j)^2$, the forward Green function of the linearized quartic equation is bounded by $(x_j/x_{i+1})^2$ for a source at step $i < j$. Reciprocal increments satisfy

$$x_{i+1}^{-1} - x_i^{-1} = \frac{\mathfrak{b}_i}{1 - \mathfrak{b}_i x_i}. \tag{6.10}$$

Thus $x_j^{-1} \asymp x_0^{-1} + \lambda^2 j$, with constants uniform under the small first-step correction. If $a_0' \in (0,1)$ is fixed sufficiently close to one, then $x_{i+1} \geq a_0' x_i$ and $d_0 < (a_0')^3$. It follows that

$$\sup_j x_j \sum_{i<j} \frac{f_i}{x_{i+1}^2} \leq C(x_0/\lambda^2 + U/x_0) = o(1). \tag{6.11}$$

For the cubic part the sum is bounded by $Cx_j \sum_{i<j} x_i$; the reciprocal comparison and $\log(1+t)/(1+t) \leq 1$ give the bound. For the transient use $x_{i+1} \geq (a_0')^{i+1} x_0$ and $x_j \leq x_0$. The residual $-\mathfrak{b}_i(g_i - x_i)^2$ has normalized size $O(\varepsilon^2)$ and Lipschitz norm $O(\varepsilon)$ after the same Green function. Indeed

$$x_j \sum_{i<j} \frac{\mathfrak{b}_i x_i^2}{x_{i+1}^2} \leq Cx_j \sum_{i<j} (x_{i+1}^{-1} - x_i^{-1}) \leq C. \tag{6.12}$$

The forward quartic equation thus maps an error of size f_i at scale i to a uniformly small relative error at every later scale.

Define the mass component by solving the linear recurrence backward with the terminal graph evaluated at the input path. The normalized forcing at step $i \geq j$ is multiplied by $L^{-s(i-j+1)} x_i/x_j$, and the terminal graph by $L^{-s(N-j)} x_N/x_j$. Their sums are uniformly bounded. Use one-step forward maps for w, k, with the prescribed w_0 and $k_0 = K_{\mathrm{init}}(\beta, \nu_0)/f_0$. The latter has normalized mass derivative at most $CU\Lambda_R/f_0 \leq C\Lambda_R$.

Choose ε small for the quartic self-coupling. Fix successively large radii for w, μ, k, in this order. Each radius absorbs its bounded input from the preceding coordinates and its strict diagonal transport. The backward mass component has only a small derivative in μ. Terms from the later coordinates into earlier ones are absorbed by increasing R after the radii are fixed. For a contraction, use the supremum norm on paths with coordinate weights $1, t, t^2, t^3$ in order (γ, w, μ, k). For $t > 0$ sufficiently small, the strictly lower triangular bounded derivatives are small in this norm. The diagonal bounds are strictly below one and the upper triangular bounds tend to zero by the Volterra estimate and $\delta_R \to 0$. This gives a uniform strict contraction of the closed box. Its fixed point satisfies the exact recurrences. The contraction estimate in common charts yields continuity, and the reciprocal increment comparison gives the bounds at N. □

Corollary 6.3 (Matching the bare mass by varying the temperature)**.** *Suppose the hypotheses of Theorem 6.2 hold for the physical initial family of Proposition 5.1 on the real window $|\beta - 1| \leq A_0 g_{\mathrm{mic}}$, with the constants in the bound for the target mass fixed before A_0. If A_0 is sufficiently large, then for all sufficiently large R and $n \geq n_0(R)$ there is a β in the window such that $\nu_0^{\mathrm{bare}}(\beta) = \nu_0^*(\beta)$. At this temperature the physical initial family follows the trajectory with the prescribed terminal curvature.*

Proof. By Corollary 5.2, the bare scalar is $\nu_0^{\mathrm{bare}} = \lambda^{-1}(b_0/R)^s r_h(\beta)$, where r_h is the bulk row of the interaction Hessian in the primary field. By Corollary 2.4, $r_h = h_0^2 r$, with $h_0 = p^{-1}$ bounded above and below, $r_h(1) = O(g_{\mathrm{mic}})$, and $r_h'(\beta) = -h_0^2(1 + O(g_{\mathrm{mic}}))$. The fixed box and target constants must be chosen before A_0. The nonscalar variations are controlled by Proposition 2.16 and the estimates following it, with the threshold on R allowed to depend on A_0. The target from Theorem 6.2 is

$$r_h^*(\beta) = \lambda(R/b_0)^s \nu_0^*(\beta), \qquad |r_h^*(\beta)| \leq C\lambda^2 x_0 \leq C' g_{\mathrm{mic}}. \tag{6.13}$$

Choosing A_0 large gives opposite signs of $r_h - r_h^*$ at the endpoints of the window. At a zero of this continuous function, the full initialization, including its reference and polymer activity, is physical and has the prescribed terminal curvature. □

A marginal coefficient can also be prescribed at the microscopic scale if the sum of its increments and first variations is small. The next corollary treats that case, including the dependence of the bubble coefficient on the additional coordinate.

Corollary 6.4 (A summable marginal coordinate)**.** *Theorem 6.2 permits an additional finite-dimensional coordinate with initial value z_0 and step equation $z^+ = z + \zeta_j$ in fixed normalized units. Suppose the weighted step and terminal bounds are uniform on an enlarged box for z, with bounded derivatives in z of the normalized forcing in the stable and mass rows and the terminal graph, and with $\|D_z e_j\| \le C f_j$. The quartic equation may use a coefficient $\mathfrak{b}_j(z)$. In that case use $\mathfrak{b}_j = \mathfrak{b}_j(z_0)$ for the comparison sequence and assume $\|D_z \mathfrak{b}_j(z)\| \le C\mathfrak{b}_j$ on the enlarged box. If $\sum_{j<N}(\sup|\zeta_j| + \sup\|D\zeta_j\|) = o(1)$, with the derivative in all normalized coordinates on the box, the same conclusion holds with $\sup_j |z_j - z_0| = o(1)$.*

Proof. Use the cumulative equation $z_j = z_0 + \sum_{i<j} \zeta_i$ in the path map. Its size and all its derivatives tend to zero uniformly in the horizon, so place this coordinate first in the weighted order in the proof of Theorem 6.2. For the variable bubble, the additional quartic source at step i is $-[\mathfrak{b}_i(z_i) - \mathfrak{b}_i(z_0)]g_i^2$. Its normalized forward Green bound is at most

$$C\|z - z_0\|_\infty \sup_j x_j \sum_{i<j} \frac{\mathfrak{b}_i x_i^2}{x_{i+1}^2} \le C'\|z - z_0\|_\infty. \tag{6.14}$$

The derivative in z has a bounded forward norm, and the derivative in the normalized quartic path is $O(\|z - z_0\|_\infty)$ by the same sum. Use a ball $\|z - z_0\|_\infty \le r_R = o(1)$ containing the image of the cumulative equation. The quartic source preserves the quartic box there. The other rows have bounded derivatives in this earlier coordinate, so the same weighted norm gives the contraction. For example, unnormalized step bounds of order $x_j^2 + U d_0^j$ have summable total $O(x_0/\lambda^2 + U)$ under the reciprocal comparison. □

For a kinetic coefficient in dimension four, the corollary requires uniform control of the remainder and its first variation on an enlarged domain of fixed radius. The coupled chart of Section 4.4 controls the residual kinetic coefficient on a domain that shrinks with g_j.

6.1 Backward shooting for a small residual kinetic coefficient

To use this shrinking domain, prescribe zero residual kinetic coefficient at the terminal scale and sum its increments backward. The sum from scale j to the endpoint must then be small in units of x_j. An external parameter in the Gaussian reference will allow the resulting initial value to be matched along with the mass.

Proposition 6.5 (Backward kinetic shooting on the shrinking box)**.** *Use the comparison sequence and the normalized coordinates of Theorem 6.2. In dimension four suppose the maps of Assumption 6.1 are available on $\sup_j |z_j|/(\omega_R x_j) \le c_z$, with $0 < \omega_R \le 1$ and an additional residual kinetic equation $z_{j+1} = z_j + \xi_j$ and terminal condition $z_N = 0$. Assume the unaltered bounds of Assumption 6.1 are uniform for these kinetic inputs. Derivatives of the normalized stable, mass and activity maps, and of the terminal graph, with respect to $z_j/(\omega_R x_j)$ are bounded. In the quartic row suppose the added derivative is bounded by $C(f_j + \mathfrak{b}_j x_j^2)$ in the same direction. The initial activity may depend on $z_0/(\omega_R x_0)$ with bounded derivative after division by f_0, and the prescribed other initial coefficients retain their bounds. If*

$$\sup_{j \le N} \frac{1}{x_j} \sum_{i=j}^{N-1} \left(\sup|\xi_i| + \sup\|D\xi_i\|\right) \le \vartheta_R = o(\omega_R), \tag{6.15}$$

where D uses all the normalized coordinates, there is a unique trajectory in the controlled box satisfying both terminal conditions. It satisfies $\sup_j |z_j|/x_j \le \vartheta_R$. The bare target values z_0^*, ν_0^* depend continuously on external data when the stated maps and data do.

Proof. In the path map of Theorem 6.2 insert the component

$$(z_j/(\omega_R x_j))^{\mathrm{out}} = -(\omega_R x_j)^{-1} \sum_{i=j}^{N-1} \xi_i. \tag{6.16}$$

On the full controlled box it has supremum bound and Lipschitz constant at most ϑ_R/ω_R. Place this component first in the weighted path norm. In the quartic component the new derivative is bounded uniformly, since its forward Green bound in normalized coordinates is

$$C \sup_j x_j \sum_{i<j} \frac{f_i + \mathfrak{b}_i x_i^2}{x_{i+1}^2} \le C'. \tag{6.17}$$

With $z = 0$ the original self-map bounds apply. The value of the other arrows from this component on the smaller ball $\sup |z_j|/x_j \le 2\vartheta_R$ is $O(\vartheta_R/\omega_R)$. Fix the other box sizes with the strict margins in the proof of Theorem 6.2, and then take R large. This gives a self-map. All arrows from the kinetic component are bounded, so weighted contraction in the other rows is proved in the same order as before. All arrows into the kinetic component are small after that choice of weights, by the bound ϑ_R/ω_R. The fixed point is the desired trajectory. Every controlled solution belongs to the smaller ball by the cumulative kinetic equation, which gives uniqueness in the controlled box. Continuity follows from the same uniform contraction in common charts. □

The one-step bounds in the coupled chart contain a sunset term, a cubic remainder and a contribution from the short-range quadratic kernel. The following estimate sums these three sources in the units required by Proposition 6.5.

Lemma 6.6 (A summable source in the backward kinetic norm)**.** *Under the comparison-sequence bounds of Theorem 6.2, suppose the ordinary step bounds for $\sup|\xi_j| + \sup\|D\xi_j\|$ are*

$$C[\lambda^3 x_j^2 + f_j + \lambda x_j (R/b_j)^2]. \tag{6.18}$$

At the merged step one may replace the first coefficient by $\lambda^3 \log(2 + M_R)$ and the last by $\lambda x_0 M_R^2 (R/b_0)^2$. These bounds imply the cumulative hypothesis of Proposition 6.5, uniformly in the horizon, with $\omega_R = \Lambda_R$ in the microscopic scalings of this paper.

Proof. On all ordinary scales with the bubble lower bound, $\lambda^2 x_i^2 \le C(x_i - x_{i+1})$. The bounded number of other scales contribute at most Cx_j^2 to any sum starting at j. Thus $\sum_{i\ge j} x_i^2 \le Cx_j/\lambda^2$, and $\sum_{i\ge j} x_i^3 \le Cx_j^2/\lambda^2$. The transient obeys $\sum_{i\ge j} U d_0^i/x_j \le CU/x_0$, because $x_{i+1}/x_i \ge a_0' > d_0$ as in the proof of Theorem 6.2. The bulk head gives a geometric sum of size at most $C\lambda x_j (R/b_j)^2$. The merged terms divided by x_0 are bounded by $C\lambda + C\lambda^3 x_0 \log(2 + M_R)$. Consequently the cumulative bound divided by x_j is at most

$$C(\lambda + x_0/\lambda^2 + U/x_0 + \lambda^3 x_0 \log(2 + M_R)). \tag{6.19}$$

Here $x_0/\lambda^2 = o(\lambda)$ follows from $a_* < 1/10$, and the initial ratios in Section 5.1 give $U/x_0 = o(\lambda)$. The merged logarithmic term is $o(\lambda)$. Thus the bound is $O(\lambda) = o(\Lambda_R)$ in dimension four. □

The backward solution specifies two microscopic target values. To obtain a physical trajectory, both must agree with their bare values at the same temperature and counterterm. Signs on the four faces of a rectangle suffice for this two-parameter crossing.

Corollary 6.7 (Simultaneous thermal and kinetic matching)**.** *In Proposition 6.5 let the maps depend continuously on a real temperature β and a real external counterterm parameter a in a closed rectangle. Suppose the two bare coefficients are continuous there. Let the thermal bare coefficient minus ν_0^* have strict opposite signs on the two temperature faces, uniformly in a. Let the kinetic bare coefficient minus z_0^* have strict opposite signs on the two counterterm faces, uniformly in β. Then there are parameters in the rectangle at which the trajectory is physical and both terminal conditions hold. The kinetic sign condition holds, for example, if the kinetic bare coefficient is $z_{\mathrm{bare}}(\beta) + k(\beta, a)a$, with $|z_{\mathrm{bare}}(\beta)| \le C\lambda x_0$ and k of fixed sign with $k_* = \inf |k| > 0$, provided the counterterm interval $[a_-, a_+]$ has $a_- < 0 < a_+$ and*

$$k_* \min\{-a_-, a_+\} > (C\lambda + \vartheta_R) \sup_{(\beta,a)} x_0. \tag{6.20}$$

The supremum is over the chosen rectangle.

Proof. Multiply the two continuous matching functions by signs so each is negative on its lower coordinate face and positive on its upper coordinate face. Subtract this vector of functions from the identity on the rectangle and project coordinatewise to the rectangle. Brouwer's theorem gives a fixed point. The strict face signs force each coordinate of the fixed point to be interior, where it must be a zero of its matching function. At this point the full physical initialization agrees with the shot trajectory. For the sufficient criterion, $|z_{\mathrm{bare}}(\beta) - z_0^*| \le (C\lambda + \vartheta_R)x_0$ by Proposition 6.5. At each endpoint the term $k(\beta, a_\pm)a_\pm$ strictly dominates this error. Its signs on the two faces are opposite and uniform in β. □

For the finite-propagation family, the sign on the counterterm faces comes from the bulk kinetic coefficient of the microscopic density. The curvature matched by this construction uses the terminal slice of the chosen reference. Section 7.5 will impose the condition for the original split.

Corollary 6.8 (Physical matching for the tunable reference)**.** *Use the reference family in Proposition 4.4 with $c' = c^0$ fixed and the transfer decomposition of Section 4.3 at the original cut. Suppose a uniform one-step chart for this family satisfies Proposition 6.5 and the terminal curvature hypotheses of Proposition 7.3, and use the source bound of Lemma 6.6. Then the physical parameters can be matched with $|\delta| \le C\lambda^2 x_0$ for every sufficiently large n at fixed R. The prescribed terminal curvature is the one using the Gaussian slices of this family.*

Proof. Set $a = \delta/\lambda$. The initial moment estimates of Proposition 2.16 give a bare bulk kinetic coefficient of size $C\lambda x_0$ at $a = 0$. For the transfer choice of c, the bulk two-jet calculation in Section 4.3 uses $(c' - 1)/\zeta = O((R/n)^2)$ in place of ξ_1. Its kinetic coefficient therefore has slope in a negative and bounded away from zero, uniformly on the rectangle. The initial cumulative covariance is zero, so the compensated and local initial kinetic coefficients agree. The exact initialization uses the same heads and tails after addition of this density difference, by its weighted short-kernel bound. The change in the normalized bare mass is bounded by $C|a|(b_0/n)^2$. The thermal crossing of Section 6 therefore holds uniformly in the counterterm window by enlarging the minimum size $n_0(R)$ at fixed R. Lemma 6.6 gives $\vartheta_R = O(\lambda)$, so the kinetic slope and the bound on z_0^* give opposite signs on the faces $a = \pm A\lambda G_{\mathrm{ref}}$ for fixed sufficiently large A, since $x_0 \asymp G_{\mathrm{ref}} = A_{\mathrm{sc}}^4 g_{\mathrm{mic}}$ uniformly on the thermal window. Corollary 6.7 matches the two coordinates. Multiplying the counterterm window by λ gives the asserted bound on δ. □

7 Terminal integration and return to spins

At the last scale, the quartic potential confines the principal coordinate to a range of order $(V/g)^{1/4}$. The principal profile has overlap of order $\sqrt{V}$ with the uniform source. Together they give a magnetic second moment of order $V^{3/2}g^{-1/2}$. Proposition 7.4 makes this argument precise and restores the source in the integrated slices. Its dependence on the dimension enters through $n^s = \sqrt{V}$. In this section a_* denotes the terminal amplitude $(V/g)^{1/4}$.

Assumption 7.1 (Terminal estimates)**.** At the matched point let $g = g_N \to 0$, $h = g^{-1/4}V^{-1/4}$, and $a_* = (V/g)^{1/4}$. There is a test norm E_n controlling the supremum norm and a quotient seminorm $q_n \le CE_n$. For the conditional decomposition of Section 1, assume

$$\sigma_T^2 \asymp_{d,R} \sqrt{V}/g, \qquad \|\ell\|_2 \le C, \qquad \|u_T\|_{E_n} \le CV^{-1/2}, \qquad |\mathbf{1}^{\mathsf{T}} u_T| \ge c\sqrt{V}, \qquad \mathbb{E}e^{c\sqrt{V}\|Y\|_{E_n}^2} \le C. \tag{7.1}$$

Constants in this assumption may depend on d, R. Assume also that $v = C_<\mathbf{1}$ satisfies $\|v\|_{E_n} \le C\sqrt{V}$, where $C_<$ is the integrated covariance. The exact real terminal interaction is nonnegative and can be written

$$F_T = e^{-V_4} + \mathcal{E}, \qquad V_4 = \frac{g}{4}\sum_{x\in\Lambda}(T\phi)_x^4. \tag{7.2}$$

The filter T is uniformly bounded in ℓ^∞, with uniformly bounded inverse in ℓ^4. For $0 \le k \le 2$, assume the pointwise bound

$$|D_{v_1}\cdots D_{v_k}\mathcal{E}(\phi)| \le \epsilon_n H(\phi)G(\phi)\prod_{i=1}^{k}\frac{\|v_i\|_{E_n}}{h}, \qquad \epsilon_n = o(1). \tag{7.3}$$

Here $H = e^{-\kappa\|\phi\|_2^2/(Vh^2)}$, $G = e^{c_h q_n(\phi)^2/h^2}$, and c_h/κ is sufficiently small in terms of the principal-profile constants. The physical initial interaction and F_T are related by exact Gaussian integration, up to a field-independent positive factor.

Proposition 7.2 below gives $\epsilon_n \le C_R(\sqrt{g} + g^{-9/4}\|K_N\|)$ from the terminal cell factors. For the source, Section 2.3 gives $\|C_j\mathbf{1}\|_{F_{n,p}} \le C_R n^p b_j^{s-p}$. Since $p < s$, the geometric sum up to the terminal scale is $O_R(n^s) = O_R(\sqrt{V})$.

7.1 Terminal comparison from a finite-cell RG output

At fixed terminal aspect the cube contains a bounded number of cells. Expanding their circle product gives the pointwise error in Assumption 7.1 from the retained coefficients and the activity norm. We use the adapted chart, where the terminal test norm is $E_n = F_{n,p}$ and the terminal quotient can be taken to be $q_n = F_{n,p}$. Here $p < 1$ and $2p > s - 1$.

The regulator ratio c_h/κ must be fixed before the terminal aspect is enlarged to control the sum over cells. This order requires a principal-profile bound uniform in the choice of aspect. Whenever the annular estimates hold in a weight $p' > p$ and to a few more orders than needed by $F_{n,p}$, the bound $\|u_T\|_{F_{n,p}} \le CV^{-1/2}$ can use a constant independent of the terminal aspect, for all sufficiently large n at each fixed aspect. To see this, first use one fixed comparison aspect and let $c \uparrow c_{*,\Lambda}$ with n, R fixed. The conditional covariance and low-tail variance bounds apply uniformly in this limit. Since the integrated variance stays bounded while the full soft variance diverges, the conditional profile tends to $u_\star = B_\Lambda^{1/2}e$. The profile bound in the higher norm $F_{n,p'}$ therefore holds for $u_\star$ with the constant from this fixed comparison aspect. For the actual aspect the higher norm of $V^{1/2}(u_T - u_\star)$ is bounded, possibly with a larger constant. Its normalized ℓ^2 norm tends to zero

by the exact profile identity in Section 2.3. At depth at most δn, the difference of the weights gives a factor $C\delta^{p'-p}$ in the $F_{n,p}$ norm. In the interior, the higher norm gives bounded rescaled first differences. A value bounded away from zero would persist on a cube of side $c_\delta n$ and contradict the normalized ℓ^2 bound. To include higher derivatives of order k, apply the discrete Taylor interpolation on an interior cube with mesh $\lfloor \eta n \rfloor$ and a fixed order $P' > k$. In rescaled units the resulting bound is $C_\delta(\eta^{-k}\|\cdot\|_\infty + \eta^{P'-k})$, by the same fixed-order interpolation matrix as in the annular estimate. Send $n \to \infty$, then $\eta \downarrow 0$, then $\delta \downarrow 0$. Thus $\|V^{1/2}(u_T - u_\star)\|_{F_{n,p}} \to 0$. This permits the choice of the output coefficient c_h/κ before enlarging the terminal aspect.

Proposition 7.2 (Pointwise terminal error from the terminal coordinates)**.** *Let $b \asymp n/A$ with fixed large A, and use the adapted cell quotient and actual disjoint cell labels of Section 3. Suppose the terminal circle product, with vacuum normalized to one, has the following output.*

1. *The primary quartic is $V_4 = \frac{g}{4}\sum_x (T\phi)_x^4$ with cellwise anchor factors. The filter T has a uniformly bounded absolute row norm, bounded ℓ^4 inverse, and a bounded finite-range forward filter B_{rec} of range $O(R)$ with $\|I - B_{\mathrm{rec}}T\|_{2\to 2} = o_R(1)$.*

2. *The retained quadratic coefficients have size $O_R(g)$ in the cell field norm at radius $\ell_b = b^{-s/2}$. The nonprimary polynomial factors are $1 + P$, of degree at most six, with coefficients of size $O_R(g^2)$ in that norm. There are a bounded number per cell.*

3. *With $h_b = g^{-1/4}b^{-s/2}$ the large-field bound on each connected activity, through two derivatives, is*

$$|D_{v_1}\cdots D_{v_k}K(X,\phi)| \le C\delta_K \prod_{U\in X} \exp\left\{-\kappa_0 \frac{\|\phi_U\|_2^2}{b^d h_b^2} + c_0 \frac{q_{U^*}(\phi)^2}{h_b^2}\right\} \times \prod_{i=1}^{k} \frac{\sup_{U\in X} F_{b,U^*}(v_i)}{h_b}, \qquad 0 \le k \le 2, \tag{7.4}$$

   where $\delta_K = o(1)$ and $\kappa_0, c_0 > 0$ are fixed. Here U^* is a fixed field-support enlargement of U.

Then the pointwise terminal bound of Assumption 7.1 holds with $\epsilon_n \le C_{A,R}(\sqrt{g} + \delta_K)$. The output constants $\kappa, c_h > 0$ can be fixed before taking A sufficiently large. The conclusion applies to the terminal value of the controlled path whenever $\delta_K = O(g^{-9/4}\|K_N\|)$ and its retained coefficients have the stated sizes.

Proof. For the truncated Neumann filter used in the microscopic and bulk charts the filter assumptions follow directly from $T = \sum_{j=0}^{M} Q_\Lambda^j$. Up to the fixed scalar normalization, $B_{\mathrm{rec}} = I - Q_\Lambda$ and $T^{-1} = (I - Q_\Lambda)(I - Q_\Lambda^{M+1})^{-1}$. The row and column bounds give $\|Q_\Lambda\|_{4\to 4} \le q < 1$, and $\|I - B_{\mathrm{rec}}T\|_{2\to 2} \le q^{M+1}$.

Set $h = g^{-1/4}n^{-s/2}$, $Z = \|\phi\|_2/(\sqrt{V}h)$, and $Q_n = F_{n,p}(\phi)/h$. Restriction to cell neighbourhoods does not increase the raw norm by more than a fixed factor. The radius conversion from ℓ_b to h_b is $g^{-1/4}$. Thus a terminal quadratic and its first two direction derivatives cost $C_A\sqrt{g}(1+Q_n)^2$, with each tested direction contributing $F_{n,p}(v)/h$. A degree-k polynomial correction costs $C_A g^{2-k/4}(1+Q_n)^k$ in these same jets, hence a power at least $g^{1/2}$. There are at most C_A cells. In the term with no occupied cells, ℓ^4 invertibility gives $V_4 \ge cZ^4$. Writing $e^{-U} - 1 = -U\int_0^1 e^{-tU}\,dt$ for the quadratic corrections and expanding the finite product of $1+P$'s about one yields the error $C_{A,R}\sqrt{g}\, e^{-\kappa Z^2 + c_h Q_n^2}$ in the two tested jets. For this absorption use the quartic decay, the bound $e^{C_A\sqrt{g}(1+Q_n)^2}$, and a fixed fraction of the output quotient reserve.

Let X be the set of occupied cells in a nonempty term, and write $Z_X^2 = \|\phi_{\cup X}\|_2^2/(Vh^2)$ and $S^2 = \|(T\phi)_{(\cup X)^c}\|_2^2/(Vh^2)$. Outside $\cup X$ and its range-$O(R)$ strips, recovery of the original field

uses only the unoccupied filtered rows. Those strips have volume $O((R/b)V)$. On their interior sites use $\|\phi\|_\infty \le hQ_n$; on ordinary unoccupied sites use the operator estimate for $I - B_{\mathrm{rec}}T$. These estimates give

$$(1 - 2\|I - B_{\mathrm{rec}}T\|_{2\to 2}^2)Z^2 \le Z_X^2 + CS^2 + C(R/b)Q_n^2. \tag{7.5}$$

The coefficient of Z_X^2 comes from retaining the occupied part of the ℓ^2 norm exactly. The activity damping costs $-\kappa_0(n/b)^s Z_X^2$. The unoccupied quartic supplies any fixed negative multiple of S^2 at a fixed constant cost, since $Vgh^4 = 1$. For the quotient factors, the summed restriction estimate of Section 3 gives

$$\sum_U \frac{q_{U^*}(\phi)^2}{h_b^2} \le CA^{s-1-2p}Q_n^2. \tag{7.6}$$

Choose A so the decay on Z_X^2 exceeds the required multiple and $c_0 CA^{s-1-2p} < c_h/4$. Then take R large and R/b small. The recovery inequality leaves a strict surplus in both exponents, which absorbs $C_A\sqrt{g}(1+Q_n)^2$ and the finite-degree polynomials from the two field derivatives of the reference. The term is bounded in these jets by $C_{A,R}\delta_K^m e^{-\kappa Z^2 + c_h Q_n^2}$, where m is the number of connected occupied components. Summing over the bounded number of subsets proves the claim. □

7.2 The terminal mass graph from small-field coefficients

The shooting theorem requires the mass as a function of the other terminal coordinates. Varying the mass at fixed activity and fixed Gaussian covariance gives a leading curvature derivative equal to the squared norm of the rescaled principal profile in mass units. The next proposition bounds the error in this derivative, which allows us to solve the scalar curvature equation uniformly.

Proposition 7.3 (Terminal curvature estimate in the cube chart)**.** *Use a fixed terminal aspect and the canonically normalized tail. Let $x = x_N$, $f = f_N$, and Λ_R be as in Section 6, and assume $(x^2 + f)/(\Lambda_R x) = o(1)$. Suppose the exact terminal circle-product chart is holomorphic and even in the field on an enlarged complex coefficient box and satisfies the following small-field estimates, with the vacuum removed.*

1. *The primary quartic uses a fixed real filter T with uniformly bounded absolute rows, and its coefficient is of order x in the stable sector. The total quadratic has the form $\frac{1}{2}\nu b^{-s}\|\phi\|_2^2 + U_{\mathrm{st}}$, where U_{st} has size $O(\Lambda_R x)$ at $\ell_b = b^{-s/2}$. The remaining polynomial factors are $1 + P$, of bounded degree and coefficient size $O(x^2)$.*

2. *The activity and its first two field jets satisfy the small-field bound Cf times the product of growing cell regulators. In a unit normalized activity direction the same estimate, with Cf, bounds the variation. The growing regulator on this fixed number of cells is bounded by $\exp(C\|\phi\|_{E_n}^2/n^{-s})$.*

3. *In the terminal chart, the ν derivative of $V_4 + U_{\mathrm{st}}$ has size $o(1)$ at the small-field radius. This holds if these raw coordinates are independent of ν, or if the inverse normal-form chart has a small first variation there. The test directions satisfy $\|\rho_N\|_{E_n} = O(n^{-s/2})$ for $\rho_N = n^{s/2}u_T$.*

For each fixed stable box, take the mass radius to be $C_M\Lambda_R x$, with C_M larger than a fixed multiple of the bound for $|q_N(0)|/(\Lambda_R x)$ divided by the lower bound on $c_N = b^{-s}\|\rho_N\|_2^2$. The coefficient bounds must hold with a further fixed-fraction collar for Cauchy's estimates. Suppose the hard field has the exponential moment $\mathbb{E}e^{cn^s\|Y\|_{E_n}^2/\lambda} \le C$, and $n^s/\sigma_T^2 = O(\lambda x)$. Then the exact terminal curvature has the graph in Assumption 6.1, with normalized activity derivative $O(f/(\Lambda_R x))$.

Proof. Write F_T for the terminal interaction, and $Z(a) = \mathbb{E}F_T(a\rho_N + Y)$, using the rescaled coordinate a. At $a = 0$ we use two field jets in the direction ρ_N. The terminal number of cells is bounded, and $\|\rho_N\|_{E_n}/\ell_b = O(1)$. If $y = \|Y\|_{E_n}/\ell_b$, the integrated jets are bounded uniformly whenever their pointwise majorant is a fixed polynomial in y times e^{Cy^2}, by the hard Gaussian moment and small λ.

The empty-activity term of F_T is $e^{-V_4-U}\prod(1+P)$, with $U = \frac{1}{2}\nu b^{-s}\|\phi\|_2^2 + U_{\mathrm{st}}$. In the complex box, the primary quartic exponential has modulus at most one on real fields. Taylor's integral formula for the exponential bounds the integrated two-jet difference from $1 - V_4 - U$ by Cx^2. Indeed each term in that remainder carries two factors of $V_4 + U$ or their field derivatives, and the adverse exponential from U is at most $e^{C\Lambda_R x(1+y)^2}$. The factors $1 + P$ produce errors of order x^2 by the same integrated jet bound. Every nonempty activity term is $O(f)$ in these jets, by the product rule and the small-field regulator. The entire error is holomorphic on the enlarged box, and changing only the normalized activity by a vector of unit norm changes its first variation by at most Cf. This follows by differentiating the finite circle product, using the stated activity variation and integrating the same majorant.

In particular $Z(0) = 1 + O(x + f)$, so its scalar inverse is uniform. Let $c_N = b^{-s}\|\rho_N\|_2^2 \asymp 1$. The second a derivative of the uniform mass is $c_N\nu$. The leading quartic contribution is bounded by $C\lambda x$ since

$$\left|\mathbb{E}D^2_{\rho_N}V_4(Y)\right| \le Cx\sum_{z\in\Lambda}|(T\rho_N)_z|^2\mathbb{E}|(TY)_z|^2 \le C\lambda x\, V n^{-2s} = C\lambda x. \tag{7.7}$$

By evenness $Z'(0) = 0$. Thus the curvature $q_N = n^s/\sigma_T^2 - Z''(0)/Z(0)$ has the form $c_N\nu + r_N$, where $r_N = O(\Lambda_R x)$. At $\nu = 0$ the bound depends on the fixed stable box and has activity contribution $O(f) = o(\Lambda_R x)$.

For the ν derivative of $Z''(0)$, Cauchy's formula on the larger mass box differentiates the $O(x^2+f)$ error with cost $C/(\Lambda_R x)$. The assumed small first variation of $V_4 + U_{\mathrm{st}}$ shows $\partial_\nu Z''(0) = -c_N + o(1)$ and $\partial_\nu Z(0) = O(1)$. Consequently $\partial_\nu q_N = c_N + o(1)$. The normalized γ, w derivatives of $r_N/(\Lambda_R x)$ are bounded by Cauchy's formula on their fixed enlarged boxes. The normalized activity derivative is $O(f/(\Lambda_R x))$ by the integrated jet variation. These bounds give a contraction of $\nu \mapsto -c_N^{-1}r_N$ on a disc of radius a fixed multiple of $\Lambda_R x$, after choosing that multiple from the zero-mass bound and then taking R large. Its fixed point is holomorphic and real for real data. Implicit differentiation gives the required normalized derivatives. □

The ratio in the proposition is implied by $x_0/\Lambda_R \to 0$ and $U/(\Lambda_R x_0) \to 0$, both of which follow from the initial scales in Section 6. Indeed $f_j/(\Lambda_R x_j) \le x_0^2/\Lambda_R + [U/(\Lambda_R x_0)](d_0/a_0')^j$, where $x_{j+1} \ge a_0' x_j$. The frozen resonance and the terminal low-tail variance estimate supply $n^s/\sigma_T^2 = O(\lambda x_N)$ before the mass is matched. The inputs still to be checked in each application are the retained coefficient estimates, the first variation of the inverse mass coordinate, and the small-field activity bound.

We can now use the quartic scale to estimate the magnetic second moment. The first bound below concerns the terminal field. The second estimates the change when a uniform source is coupled to the full auxiliary field, including the slices already integrated.

Proposition 7.4 (The terminal moment and Gaussian source restoration)**.** *Under Assumption 7.1,*

$$\mathbb{E}_T(\mathbf{1}^{\mathsf{T}}\phi)^2 \asymp V^{3/2}g^{-1/2}, \qquad |\mathbb{E}_0(\mathbf{1}^{\mathsf{T}}\phi)^2 - \mathbb{E}_T(\mathbf{1}^{\mathsf{T}}\phi)^2| \le CV^{3/2}. \tag{7.8}$$

Here $\mathbb{E}_T$ and $\mathbb{E}_0$ use the normalized terminal and initial auxiliary laws.

Proof. Put $a = a_* t$. Since $\ell^{\mathsf{T}} Y = 0$ and $\ell^{\mathsf{T}} u_T = 1$, one has $\|a u_T + Y\|_2 \geq |a|/\|\ell\|_2$. Also $a_* = \sqrt{V} h$; the quotient bound on u_T and the triangle inequality give

$$H(a_* t u_T + Y) G(a_* t u_T + Y) \leq e^{-ct^2} \exp(C c_h h^{-2} \|Y\|_{E_n}^2). \tag{7.9}$$

Since $h^{-2} = g^{1/2}\sqrt{V}$, the latter exponential has a uniform expectation with room for Cauchy–Schwarz. The Gaussian variance of $\mathbf{1}^{\mathsf{T}} Y$ is $O(V^{3/2})$, from the supremum norm bound and the exponential moment. Consequently the same expectation with insertion $|\mathbf{1}^{\mathsf{T}} Y|^k$, $k \leq 2$, is bounded by $C V^{3k/4} e^{-ct^2}$.

The same field-derivative estimate assumed for $\mathcal{E}$ holds with constant C for e^{-V_4}. Indeed $\|T\phi\|_4^4 \geq c\|\phi\|_4^4 \geq c V^{-1} \|\phi\|_2^4$. Differentiation produces fixed-degree polynomials in $g^{1/4}\|T\phi\|_4$ and in the factors $g^{1/4}\|T v_i\|_4 \leq C\|v_i\|_{E_n}/h$. Absorbing the former by a fraction of e^{-V_4} leaves H. Thus the preceding Gaussian envelope applies to the integrand and its two field derivatives.

For compact t-sets, $\mathbb{E} V_4(a_* t u_T + Y) \leq C(t^4 + g)$ by the supremum norm bound on T and the fourth moment of $\|Y\|_\infty$. Jensen's inequality yields a uniform positive lower bound for the quartic weight integrated over Y, and the error is $o(1)$ there. The Gaussian factor $e^{-a_*^2 t^2/(2\sigma_T^2)}$ is bounded below on compact t-sets since $a_*^2/\sigma_T^2 = O(\sqrt{g})$. Integrating the upper envelope and using, for the lower second moment, $1 \leq |t| \leq 2$, gives $Z_T = \mathbb{E}_{C_T} F_T \asymp a_*/\sigma_T$ and $\mathbb{E}_T a^2 \asymp a_*^2$. The weighted envelope also gives $\mathbb{E}_T(\mathbf{1}^{\mathsf{T}} Y)^2 = O(V^{3/2})$. The L^2 triangle inequalities and the overlap of u_T give the first conclusion.

For the second conclusion put $S = \mathbf{1}^{\mathsf{T}} \phi$. Gaussian translation in the earlier covariance turns the source $J\mathbf{1}$ into a factor $e^{J^2 \mathbf{1}^{\mathsf{T}} C_< \mathbf{1}/2}$ and the shift $F_T(\phi + Jv)$, with the terminal source e^{JS} retained. The same envelope permits two derivatives at $J = 0$. For fixed n, the shift incurs only a constant and an arbitrarily small loss in the two quadratic exponents as J varies near zero, and the linear source is absorbed by the remaining Gaussian reserve. Thus the difference of moments is

$$\mathbf{1}^{\mathsf{T}} C_< \mathbf{1} + 2 Z_T^{-1} \mathbb{E}_{C_T}[S D_v F_T] + Z_T^{-1} \mathbb{E}_{C_T} D_v^2 F_T. \tag{7.10}$$

The first term is at most $C V^{3/2}$. With $B_v = \|v\|_{E_n}/h \leq C V^{3/4} g^{1/4}$, the two other terms are bounded by $C V^{3/4} g^{-1/4} B_v$ and $C B_v^2$, respectively, by the derivative envelope and the normalization bound. Both are $O(V^{3/2})$. □

To return to spins we must account for the variation of $A_\Lambda \mathbf{1}$ near the boundary. The ferromagnetic two-point function is nonnegative, so uniform upper and lower bounds on the entries of this vector suffice.

Lemma 7.5 (Positive-row transfer to the Ising susceptibility)**.** *Let A_Λ be a fixed positive Gaussian split of a ferromagnetic Ising interaction such that A_Λ has nonnegative entries,* $\inf_x (A_\Lambda)_{xx} \geq a'_R > 0$ *and* $\sup_x \sum_y (A_\Lambda)_{xy} \leq p_R < \infty$*, uniformly in n at fixed R. Let the unscaled auxiliary field be $h_{\mathrm{phys}} = \sqrt{\beta} A_\Lambda \sigma + \eta$ with independent $\eta \sim N(0, A_\Lambda)$. At zero field,*

$$\beta (a'_R)^2 \mathbb{E}_\beta(\mathbf{1}^{\mathsf{T}} \sigma)^2 \leq \mathbb{E}_\beta(\mathbf{1}^{\mathsf{T}} h_{\mathrm{phys}})^2 \leq \beta p_R^2 \mathbb{E}_\beta(\mathbf{1}^{\mathsf{T}} \sigma)^2 + p_R V. \tag{7.11}$$

Proof. The vector $A_\Lambda \mathbf{1}$ has entries in $[a'_R, p_R]$. The two-point functions $\mathbb{E}_\beta \sigma_x \sigma_y$ are nonnegative, as follows by expanding the zero-field Ising partition functions in $\tanh(\beta J_{xy})$ and retaining the parity constraint. Multiplying these expectations by the two row weights and summing gives the bounds on the weighted spin term. The independent noise adds the nonnegative variance $\mathbf{1}^{\mathsf{T}} A_\Lambda \mathbf{1} \leq p_R V$. □

7.3 The four-dimensional matched point for the tunable reference

We first prove the susceptibility bound for the terminal split obtained by tuning the Gaussian reference. The temperature and counterterm solve the two matching equations on a common rectangle. Fix the original cut J, the pole parameter $c' = c^0$ and the transfer parameter $\eta > 0$ of Section 4.3. For each δ use $d_\delta = (\zeta + \delta)/\zeta$, $c = 1 + d_\delta(c' - 1)$ and $t_\delta = c/(d_\delta c')$. The terminal covariance is T_δ from (5.19). At the matched parameters this decomposition is held fixed when taking the curvature.

Theorem 7.6 (Susceptibility at the finite-propagation matched point)**.** *In dimension four, for all sufficiently large fixed R and $n \geq n_0(R)$, there are an external parameter $|\delta_*| \leq C\lambda^2 x_0$ and a physical temperature $\tau_\Lambda^{\mathrm{pc},\delta_*}$ for which the terminal effective potential of the transferred split* (5.19) *has zero curvature in its principal coordinate $a = \ell^{\mathsf{T}}\phi$, with terminal Gaussian covariance T_{δ_*}. The physical Ising susceptibility at this temperature satisfies*

$$c_R n^2 \sqrt{\log n} \leq \chi_\Lambda^F(\tau_\Lambda^{\mathrm{pc},\delta_*}) \leq C_R n^2 \sqrt{\log n}. \tag{7.12}$$

The constants are independent of n.

Proof. Use the covariance transfer decomposition in Section 4.3. Its ordinary annular covariance bounds, filtered bubble bounds and merged amplitude hold uniformly across the real matching window. The ordinary and transferred terminal data can use the same terminal aspect, chosen in the order of Section 7.1. The off-shell activity for the two scalar coordinates is given by Lemma 5.3. Its normalized initial mass and kinetic derivatives after division by f_0 are at most

$$C(U/x_0)\Lambda_R x_0/f_0 \leq C\Lambda_R. \tag{7.13}$$

The remaining prescribed stable coefficients have the entry bounds of the coupled chart. In particular after removing the bulk kinetic two-jet the fixed short head has effective size $O(\lambda x_0/M_R^2)$ at entry. The cut and the boundary tensors are fixed at each choice of the external parameters and continuously depend on those parameters in the common label charts.

Proposition 4.7 and Proposition 3.12 give the ordinary, collar and activity bounds on the same enlarged extraction disc. At radius f_j the activity contraction is the one in Corollary 3.16: choose L and the size margin so that the leading operator and the large-set error sum to less than $d_0/4$, then take R large for the mixed-norm decoration and Cauchy errors. The ratio $f_j/g_j^{9/4}$ tends to zero uniformly, and $f_{j+1} \geq d_0 f_j$. The normalized kinetic coordinate is $z_j/(\Lambda_R x_j)$. Its arrows into the stable, mass and activity rows are bounded by the coupled chart, and the derivative of the quartic error in that coordinate is $O(f_j)$. Proposition 4.5 and Lemma 6.6 give the backward kinetic norm $O(\lambda) = o(\Lambda_R)$, including first variations. Thus the estimates for the dynamics in Proposition 6.5 have the stated strict margins.

At the terminal scale the stable head and collar quadratic sizes are $O(\Lambda_R x_N)$. The raw kinetic coefficient satisfies the same bound by the inverse compensated chart and the terminal condition $z_N = 0$. Both collar polynomial degrees and the ordinary margin factors have size $O(x_N^2)$. At fixed stable and activity coordinates the inverse mass chart has the derivative of Lemma 3.14, so the scalar mass derivative is $1 + O(\lambda x_N)$ and the induced variation of the quartic and local kinetic terms is $O(\lambda x_N)$ in the small-radius units. The fixed heads and full collar tensors are independent raw coordinates in this inverse. The terminal activity and its first variations are in the stated mixed norm, on a fixed number of labels. The Gaussian estimates are given by Lemma 5.4. Hence Proposition 7.3 applies, also for kinetic inputs on the shrinking box by the same finite-jet argument. It gives the required continuous terminal mass graph and its bounded normalized kinetic derivative.

We may now apply the backward shooting proposition and Corollary 6.8. The latter uses the exact finite-head kinetic coefficient, which differs from the bulk two-jet used for its sign comparison by the exponentially small initial tail. This difference fits in the strict crossing margin. At the resulting temperature and counterterm both scalar bare coordinates match the original density, as in Lemma 5.3. The spin model at that temperature is the physical one, since the covariance change was compensated in its auxiliary density.

Finally $g_N \asymp_R (\log n)^{-1}$. The terminal bounds in the chart and Proposition 7.2 give the pointwise terminal error with $g_N^{-9/4}\|K_N\| = o(1)$. Lemma 5.4 supplies all the Gaussian bounds for Proposition 7.4, including the integrated-source estimate. It follows that the initial auxiliary magnetic second moment has order $V^{3/2}g_N^{-1/2}$. The density compensation preserves this initial auxiliary law. The spin comparison of Lemma 7.5 therefore applies with the same $O_R(V)$ error. Divide by $V = n^4$ to get the asserted bounds. □

7.4 Comparison of the two terminal curvature conditions

The physical law of the full auxiliary field is the same in the two references. The terminal coordinate depends on how its covariance is split. To recover its original marginal, we keep its conditional distribution given the full field. Gaussian regression then identifies the covariance that must remain after the preliminary integrations.

Lemma 7.7 (Original terminal coordinate under a changed reference)**.** *Let $C^0 = C^0_< + C^0_T > 0$ be the original Gaussian covariance, and let $a_0 = \ell^{\mathsf{T}}\phi^0_T$ be its terminal coordinate, with $\sigma_0^2 = \ell^{\mathsf{T}}C^0_T\ell > 0$. Put*

$$w_0 = (C^0)^{-1}C^0_T\ell, \qquad \eta_0^2 = \sigma_0^2 - w_0^{\mathsf{T}}C^0w_0. \tag{7.14}$$

In any Gaussian reference $\Sigma > 0$ with physical interaction factor F^Σ representing the same law of the full auxiliary field, set

$$V_{0,\Sigma} = \eta_0^2 + w_0^{\mathsf{T}}\Sigma w_0, \qquad v_{0,\Sigma} = \Sigma w_0, \qquad u_{0,\Sigma} = v_{0,\Sigma}/V_{0,\Sigma}. \tag{7.15}$$

Suppose that $C^\Sigma_{<,k}$ has been integrated by scale k, with exact interaction F^Σ_k, and write $C^\Sigma_{T,k} = \Sigma - C^\Sigma_{<,k}$. Define

$$B_{0,\Sigma,k} = C^\Sigma_{T,k} - \frac{v_{0,\Sigma}v_{0,\Sigma}^{\mathsf{T}}}{V_{0,\Sigma}}. \tag{7.16}$$

If $B_{0,\Sigma,k} \geq 0$, the density of the original coordinate a_0, up to a constant independent of a, is

$$e^{-a^2/(2V_{0,\Sigma})}\mathbb{E}_{B_{0,\Sigma,k}}F^\Sigma_k(au_{0,\Sigma} + Y). \tag{7.17}$$

In particular its zero-curvature condition can be imposed on the reference-Σ RG chart using this profile and hard covariance. The condition (7.16) *is equivalent to*

$$v_{0,\Sigma} \in \operatorname{ran} C^\Sigma_{T,k}, \qquad v_{0,\Sigma}^{\mathsf{T}}(C^\Sigma_{T,k})^\dagger v_{0,\Sigma} \leq V_{0,\Sigma}, \tag{7.18}$$

where the dagger denotes the inverse on the range.

Proof. In the original Gaussian split, the joint covariance of the full field ϕ and a_0 has cross column $C^0_T\ell$. Gaussian regression therefore represents $a_0 = w_0^{\mathsf{T}}\phi + \eta_0 Z$, with Z standard normal independent of ϕ. Weighting by the physical interaction, a function only of ϕ, preserves this conditional law. As a result, under the reference Σ one can first form the same linear expression with ϕ centered Gaussian of covariance Σ and then weight by $F^\Sigma(\phi)$. The proxy variance of the linear expression

is $V_{0,\Sigma}$ and its cross column is $v_{0,\Sigma}$. Conditioning on its value a gives mean $au_{0,\Sigma}$ and covariance $\Sigma - v_{0,\Sigma}v_{0,\Sigma}^{\mathsf{T}}/V_{0,\Sigma}$. If (7.16) is positive semidefinite, this covariance is the sum of $C^{\Sigma}_{<,k}$ and $B_{0,\Sigma,k}$. Gaussian addition gives (7.17). The range test follows by conjugating the rank-one inequality by the inverse square root of $C^{\Sigma}_{T,k}$ on its range and then testing on its kernel. The argument includes degenerate hard covariances. □

This gives a sufficient terminal estimate for matching the original coordinate via a different reference. Suppose the outgoing chart is available at a cut k a bounded number of scales before the original terminal scale, and that $B_{0,\Sigma,k}$ in (7.16) is positive and has the uniform hard exponential norm moment of Proposition 7.3. Suppose also that

$$\|n^{s/2}u_{0,\Sigma}\|_{E_n} = O(n^{-s/2}), \qquad b_k^{-s}\|n^{s/2}u_{0,\Sigma}\|_2^2 \asymp 1, \qquad n^s/V_{0,\Sigma} = O(\lambda x_k). \tag{7.19}$$

If the profile, covariance and proxy variance are continuous in the external parameters and independent of the variable terminal mass, the proof of Proposition 7.3 applies to (7.17). Its strict mass derivative is the displayed normalized profile square plus the same inverse-chart error. One can therefore use this scalar graph in the backward shooting instead of the graph for the modified terminal coordinate. The number of steps released into the terminal comparison is bounded, so the comparison sequence, the terminal aspect constants and the activity weights are unchanged up to fixed factors. The usual terminal moment proof still uses the natural decomposition of the modified reference. Its estimate needs the terminal coefficient box of order $\Lambda_R x_k$, which the new curvature graph supplies.

A scalar example illustrates both issues in this comparison. A small change of reference can move the curvature zero and can make the candidate conditional covariance at the original cut negative.

Proposition 7.8 (The curvature and the hard covariance need not transfer)**.** *Even a scalar covariance change of the form $\Sigma_t = tC^0 + (1-t)A$ can alter the natural terminal curvature for a fixed physical law. It can also make* (7.16) *indefinite at the original terminal cut for arbitrarily small $|t-1|$.*

Proof. In dimension one for the auxiliary field, take

$$C^0 = 2, \qquad C_T^0 = 1, \qquad \ell = 1. \tag{7.20}$$

Then $w_0 = 1/2$ and $\eta_0^2 = 1/2$. For t near one take

$$\Sigma_t = 2t + \tfrac{1}{2}(1-t), \qquad C_{T,t} = t. \tag{7.21}$$

Both parts of this Gaussian split are positive. This is of the form of a scaled covariance plus a microscopic term. For the original conditional coordinate the proposed hard variance at the same cut is

$$H_t = t - \frac{\Sigma_t^2}{2+\Sigma_t}. \tag{7.22}$$

It vanishes at $t = 1$ and its derivative there is $-1/8$. Hence it is negative for all sufficiently small positive $t - 1$.

The same example shows that the natural curvatures can differ even for a fixed full physical law. Take the full law to be the even mixture of normals of variance one and means $\pm\sqrt{3}$. In a Gaussian reference of variance Σ_t, its positive density relative to that Gaussian defines the interaction factor.

The natural terminal coordinate is conditionally Gaussian given the full field, with mean $w_t\phi$ and variance η_t^2, where

$$w_t = t/\Sigma_t, \qquad \eta_t^2 = t - t^2/\Sigma_t. \tag{7.23}$$

Thus its marginal is the even mixture of normals of variance $v_t = w_t^2 + \eta_t^2$ and means $\pm w_t\sqrt{3}$. Its log density has second derivative $3w_t^2/v_t^2 - 1/v_t$ at zero. This vanishes exactly when $\eta_t^2/w_t^2 = 2$. The latter ratio is $\Sigma_t(\Sigma_t - t)/t$, equal to two at $t = 1$ and with derivative $1/2$ there. The curvature consequently changes under an arbitrarily small reference modification. □

For the transfer family on the cube we can estimate the variance and the profile throughout the counterterm window. The covariance moment then follows whenever the subtraction in (7.16) is positive.

Proposition 7.9 (Bounds for the original-coordinate comparison)**.** *Use the covariance transfer of Section 4.3 with the unperturbed parameter fixed there, in dimension four. Let k be a fixed number of scales before the original terminal cut. For every $1/2 < p < 1$, on the physical counterterm window and for all sufficiently large n at fixed R, the proxy variance and the profile in Lemma 7.7 satisfy the bounds required for its terminal mass graph. If $B_{0,\Sigma,k} \geq 0$, its hard Gaussian exponential moment in $F_{n,p}$ also has the required uniform bound. Thus positivity of this particular rank-one subtraction is sufficient for all the Gaussian inputs of that transfer at the chosen cut.*

Proof. Write $T = C_T^0$, $P = C_<^0$, $G = T + P$, $A_c = a_R A_\Lambda$, and $\sigma^2 = \ell^{\mathsf{T}} T \ell$. Then $\Sigma = tG + (1-t)A_c$, where $t = t_\delta$ is bounded above and below by positive constants. Put $u = T\ell/\sigma^2$ and $r = A_c w_0$, where $w_0 = G^{-1}T\ell$. The regression quantities satisfy

$$v_{0,\Sigma} = t\sigma^2 u + (1-t)r, \tag{7.24}$$

$$V_{0,\Sigma} = t\sigma^2 + (1-t)(\eta_0^2 + w_0^{\mathsf{T}} A_c w_0). \tag{7.25}$$

The old conditional variance satisfies $\eta_0^2 = \ell^{\mathsf{T}}(P - PG^{-1}P)\ell \leq C\lambda n^2$ by Section 2.3. Also $A_c \leq G$ and $w_0 = \ell - G^{-1}P\ell$, so

$$w_0^{\mathsf{T}} A_c w_0 \leq 2\ell^{\mathsf{T}} A_c \ell + 2\ell^{\mathsf{T}} P G^{-1} P \ell \leq C\lambda n^2. \tag{7.26}$$

Consequently $V_{0,\Sigma} = t\sigma^2 + O(|1-t|\lambda n^2)$ and $V_{0,\Sigma} \asymp n^2/(\lambda x_k)$.

For the profile estimate we use two elementary operator bounds. First an exponentially short kernel, with the clipping defect bounds of Proposition 4.4, maps $F_{n,p'}$ to $F_{n,p}$ with bounded norm for $p < p' < 1$ and $n \geq n_0(R)$. At a site of truncated boundary depth $\rho = \min(n, R + \operatorname{dist}(x, \Lambda^c))$, split its bulk row at a small multiple of ρ. Short translations in that row commute with each tested difference and preserve the local weights. On the long translations and the clipping defect, use the exponential weight before taking the fixed number of point differences. Their extra cost is bounded by

$$C(\rho/n)^{p'-p}(1+\rho)^P e^{-c\rho/R}. \tag{7.27}$$

This has supremum tending to zero as $n \to \infty$ at fixed R. The same estimate applies to products of the short kernels after taking an intermediate exponent. Second, for the integrated covariance,

$$\|Pf\|_{F_{n,p'}} \leq C\lambda n^2 \|f\|_{F_{n,p'}} \qquad (p' < 1). \tag{7.28}$$

To see this, apply the estimate to each individual slice at side b before merging. On a stencil at depth at most Cb, every read field value costs $C(b/n)^{p'}\|f\|_{F_{n,p'}}$. For a difference of order a on the covariance row, covariance Cauchy–Schwarz and the annular bounds give $C\lambda b^{-2}(\min(b,\rho)/b)^{p''}\min(b,\rho)^{-a}$,

with $p' < p'' < 1$. Summing the $O(b^4)$ read sites and multiplying by the raw test weight gives $C\lambda b^2 \|f\|_{F_{n,p'}}$. At larger depth the slice equals the bulk there by finite propagation. Move the differences to f and use the absolute row bound $C\lambda b^2$. The geometric sum over the integrated slices proves the displayed estimate.

The limiting pole profile of Section 7.1 is $B_\Lambda \ell$ in the common normalization. The short inverse estimate and the pole profile bound, using exponents strictly between p and one, therefore give $\|\ell\|_{F_{n,p'}} \le Cn^{-2}$ for a fixed $p' > p$. The exact inverse relation $A_c G^{-1} = I - c' A_\Lambda B_\Lambda^{-1}$ gives

$$r = A_c \ell - (I - c' A_\Lambda B_\Lambda^{-1}) P \ell. \tag{7.29}$$

The two operator bounds just proved show $\|r\|_{F_{n,p}} \le C\lambda$. The formula for the cross column now gives $\|u_{0,\Sigma} - u\|_{F_{n,p}} = o(n^{-2})$, uniformly on the counterterm window. In particular the normalized profile square has the same upper and lower bounds as the unperturbed one.

It remains to bound the hard moment, conditional on positivity. Let D be the sum of the unperturbed slices released by stopping at k. In the transferred decomposition $C_{T,k}^{\Sigma} \le t(T + D)$. Subtract the regression rank-one matrix, and write $T_h = T - \sigma^2 uu^{\mathsf{T}}$. The displayed variance and cross-column identities give the form bound

$$B_{0,\Sigma,k} \le t(T_h + D) + C\lambda n^2 uu^{\mathsf{T}} + \frac{C}{\lambda n^2} rr^{\mathsf{T}}. \tag{7.30}$$

Indeed the difference of the two soft rank-one coefficients is $O(\lambda n^2)$, and the mixed form of u, r is bounded by the last two terms by Cauchy–Schwarz. The remaining rr^{T} term from the subtraction is nonpositive. The old hard covariance has the required Gaussian moment. The bounded number of released slices have the annular moment at scales comparable to n. The square roots of the last two rank-one terms have profile norms at most $C\sqrt{\lambda} n^{-1}$. Combining these finitely many moments and applying Gaussian domination by conditional Jensen proves the required bound for $B_{0,\Sigma,k}$ when it is positive. All vectors and matrices in the comparison are continuous in the external parameter. □

7.5 The prescribed terminal curvature in dimension four

To impose the curvature of the prescribed split, we make two changes in the calculation. First we integrate in a microscopic field whose covariance with the original coordinate is exactly $T\ell$. This removes the term $(1 - t)A_c w_0$, which in Lemma 7.7 requires an inverse estimate at high frequencies. Then we stop one scale earlier. The slice retained at this cut supplies a variance margin of order λn^2, enough to make the prescribed regression covariance positive throughout the matching rectangle.

Throughout this subsection $T = C_T^0$, $G = C^0$, $A_c = a_R A_\Lambda$, $\sigma^2 = \ell^{\mathsf{T}} T \ell$, and w_0, η_0^2 are the original regression data of Lemma 7.7. The original pole parameter $c' = c^0$, cut J, and coordinate a_0 are held fixed throughout the matching rectangle. Write $t = 1 - \varepsilon$ and $\Sigma = tG + \varepsilon A_c$. This is the transfer family of Section 4.3 with $\delta = \zeta\varepsilon c'/(1 - \varepsilon c')$. The preterminal integrations will use its finite-range slices.

The field transformation must preserve the short-range estimates used at initialization. Its dependence on the external parameter must also allow the bare kinetic coefficient to cross the value obtained by backward shooting.

Lemma 7.10 (Microscopic change of field)**.** *Let $S = G\Sigma^{-1}$. Use ψ as the field in the RG map and represent the original full field by $\phi = S\psi$. For $|\varepsilon| \le H\lambda^2 x_0(0)$ with fixed H, the transformed physical interaction has the initial activity bounds and the off-shell mass and kinetic first-variation bounds of Lemma 5.3. Its bare kinetic coefficient is strictly increasing in ε/λ, with slope bounded above and away from zero. Its scalar mass satisfies the thermal crossing estimate of Section 6.*

Proof. In the following matrix identities A, B are the actual cube matrices. Set $D = B - \varepsilon c' A$. From the old precision we get

$$\Sigma G^{-1} = I - \varepsilon c' A B^{-1} = D B^{-1}, \qquad S = B D^{-1}, \qquad B^{-1} S = D^{-1}. \tag{7.31}$$

If $F^0 = e^{-U^0}$ is the original full interaction relative to the Gaussian with covariance G, the interaction in the ψ field relative to the Gaussian with covariance Σ is, up to a constant,

$$F_{\mathrm{in}}(\psi) = \exp\{-U^0(S\psi) - \tfrac{1}{2}\psi^{\mathsf{T}} H_\varepsilon \psi\}, \qquad H_\varepsilon = S^{\mathsf{T}} G^{-1} S - \Sigma^{-1} = \frac{\varepsilon c'}{a_R} D^{-1}(B - c'A) D^{-1}. \tag{7.32}$$

Indeed $S^{\mathsf{T}} G^{-1} = \Sigma^{-1}$, $S - I = \varepsilon c' A D^{-1}$, and $a_R \Sigma^{-1} A = D^{-1}(B - c'A)$. Thus the old Hessian is conjugated by S, and the old Hessian-subtracted spin pressure is evaluated at D^{-1} applied to the new primary field in physical units.

For both signs of ε write $D = d_\varepsilon (I - K)$, where

$$d_\varepsilon = 1 - \varepsilon c' A(0), \qquad K = d_\varepsilon^{-1}\big[Q_\Lambda + \varepsilon c'(A - A(0) I)\big]. \tag{7.33}$$

The bulk off-diagonal kernel of Q dominates a fixed multiple of D_R, and the off-diagonal entries of A are bounded by a fixed multiple of D_R. Hence K is nonnegative for large R. Its absolute row and column sums, including a small fixed exponential weight at length R, are bounded by a constant strictly below one. This uses the original row gap and $\varepsilon = o_R(1)$. The inverse D^{-1}, its first ε-derivative and the products in the Hessian identity consequently have uniform weighted bounds. Each term in a clipping discrepancy has a short walk reaching the exterior, which supplies the additional factor $e^{-c\,\mathrm{dist}(x,\Lambda^{\mathrm{c}})/R}$ in row x. A differentiated inverse marks one step and costs only a polynomial in the walk length. In particular the added Hessian has weighted size $C|\varepsilon|/a_R$, and its derivative has size C/a_R. The derivatives of the two conjugating factors of the old residual Hessian preserve its weighted bound $O_{A_0}(g_{\mathrm{mic}}/a_R)$ on a fixed thermal window of width $A_0 g_{\mathrm{mic}}$.

The finite filter used on a label is $T_{\mathrm{f}} = d_\varepsilon^{-1} \sum_{i=0}^{m} K^i$, where m is the microscopic Neumann cutoff of Section 5.1. It has nonnegative entries, and $D T_{\mathrm{f}} = I - K^{m+1}$. This supplies the recovery identity on occupied and unoccupied label sets and the bounded ℓ^4 inverse at the terminal scale. The bulk tail and its clipping defect have the bounds of (5.4). The pressure atoms on each induced spin graph depend on their filtered arguments as in Proposition 5.1. The new filter bounds supply the two ingredients in its real-field estimate, namely bounded test directions on the enlarged stencils and recovery on the available rows. The same exponentially small tails on the two polydiscs give its complex-field estimate and its connected path expansion. The bulk and boundary quadratic heads have the weighted bounds of Proposition 2.16, since $|\varepsilon| = O(g_{\mathrm{mic}})$. Cutting these heads and then removing their bulk second moments gives the interface terms on the same actual collar labels as in Lemma 5.3. Multiplication on both occupied and unoccupied labels by the off-shell mass and kinetic factors therefore gives activity bounded by CU and partial derivatives bounded by CU/x_0 on the smaller discs of that lemma.

Normalize the bulk truncation by its row sum, absorbing the fourth power of this sum into the initial quartic. The sum is bounded above and away from zero and all fixed absolute moments of order j of the normalized filter are $O_j(R^j)$, by geometric summation of the walks. It has the cubic symmetries used in the ordinary transport and has normalized Fourier multiplier bounded below on the low-frequency bubble band. In particular the initial quartic $x_0(\varepsilon) \asymp x_0(0)$ and its comparison sequence obey the bounds used in the shrinking chart. The coefficients are continuous in the external parameters. The pole displacement still uses the original reference sequence.

Finally $B - c'A = \zeta J_\Lambda + (1-c')A$. With $p = \widehat{A}(0) = \widehat{B}(0)$ in bulk and $j_2 = \frac{1}{2}\sum_z z_1^2 D_R(z) \asymp R^2$, the ε-derivative of the coefficient of $|k|^2$ in the added bulk Hessian is

$$\frac{c'\zeta j_2}{a_R p^2}\frac{1+\varepsilon c'}{(1-\varepsilon c')^3} + O\left(\frac{R^2}{a_R}(R/n)^2\right). \tag{7.34}$$

This follows by expanding the scalar bulk symbol $\varepsilon c'(\widehat{B} - c'\widehat{A})/[a_R(\widehat{B} - \varepsilon c'\widehat{A})^2]$. The old conjugated Hessian contributes $O_{A_0}(g_{\mathrm{mic}}R^2/a_R)$ to this derivative. Cutting the weighted tails before extracting the kinetic coefficient costs $O((R^2/a_R)e^{-cb_0/R})$ in the derivative. Since $a_R = \lambda R^2$, the slope in ε/λ has the asserted strict lower bound. The old scalar row is multiplied by $(1-\varepsilon c')^{-2}$. The new normalized bare mass from H_ε is temperature independent and is $O(|\varepsilon|\lambda^{-1}(b_0/n)^2)$, since $|1-c'| = O((R/n)^2)$. These estimates preserve the strict thermal crossing for $n \geq n_0(R)$. □

At the last integrated scale the auxiliary precision has a spectral interval on which both the old terminal covariance and the released slice are large. The generalized principal eigenvector gives a test with uniformly positive spectral mass on this interval.

Lemma 7.11 (A variance margin after releasing one scale)**.** *Let J be the original cut, let $k \leq J-1$ be a bounded number of scales earlier, and write $T = C_T^0$, $\sigma^2 = \ell^{\mathsf{T}} T \ell$. Set $W_k = t^{-1}C_{T,k}^{\Sigma}$ for the transfer decomposition. Then $T\ell \in \operatorname{ran} W_k$ and, uniformly in the physical window,*

$$(T\ell)^{\mathsf{T}} W_k^{\dagger}(T\ell) \leq \sigma^2 - c\lambda n^2. \tag{7.35}$$

The constant can depend on the fixed terminal aspect.

Proof. Write $\mathcal{Q} = \mathcal{Q}_{c',0}$, $y = \mathcal{O}^*\ell = (\mathcal{L}_\Lambda \ell, 0)$, and let μ be the spectral measure of $\mathcal{Q}$ at y. Denote the old terminal multiplier by f, the sum of the released multipliers by d, and set $h(s) = (1-\eta s)(f(s)+d(s))$. Then $T = a_R \mathcal{O} f(\mathcal{Q})\mathcal{O}^*$ and $W_k = a_R \mathcal{O} h(\mathcal{Q})\mathcal{O}^*$. For every physical test z, spectral Cauchy–Schwarz gives

$$|\langle z, T\ell\rangle|^2 \leq \langle z, W_k z\rangle\, a_R \int \frac{f(s)^2}{h(s)}\, d\mu(s), \tag{7.36}$$

with zero quotients interpreted as zero. This proves the range assertion and bounds the inverse quadratic form by the integral. Since $\eta\|\mathcal{Q}\| \leq 1/2$ and $sf(s) \leq 1$,

$$\sigma^2 - (T\ell)^{\mathsf{T}} W_k^{\dagger}(T\ell) \geq a_R \int \frac{f(s)d(s)}{f(s)+d(s)}\, d\mu(s) - 2\eta a_R \|y\|^2. \tag{7.37}$$

Put $r_J = L^{-2J} \asymp M_*(R/n)^2$. On $0 < s \leq \epsilon r_J$, the bound $\sum_{i<J} q_i(s) \leq C/r_J$ gives $f(s) \geq 1/(2s)$ if ϵ is small. The low-frequency lower bound for q_{J-1} gives $d(s) \geq c/r_J$ there. To bound the spectral mass, take $U = (\mathcal{L}_\Lambda \ell, \ell)$. The identity $A_\Lambda \ell = c_{*,\Lambda}^{-1} B_\Lambda \ell$ and the normalization $\ell^{\mathsf{T}} B_\Lambda \ell = 1$ give

$$\mathcal{Q}U = \left(0, \frac{c_{*,\Lambda} - c'}{c' c_{*,\Lambda}} B_\Lambda \ell\right), \qquad \langle y, U\rangle = c_{*,\Lambda}^{-1}. \tag{7.38}$$

Both vectors have bounded norm, and the last scalar is bounded below. The component of U above ϵr_J has norm at most $\|\mathcal{Q}U\|/(\epsilon r_J) = O(\lambda^2 \bar{g}_N)$. Hence $\mu([0, \epsilon r_J]) \geq c > 0$ for all sufficiently large n at fixed R. The integral is at least c/r_J, while $\|y\|^2 = O(1)$. Using $a_R/r_J \asymp \lambda n^2$ proves the claim. □

The preceding lemma estimates the inverse quadratic form at the earlier cut. We now compare its margin with the change in the variance of the original scalar. This gives a positive Gaussian covariance on which the curvature calculation can be performed.

Proposition 7.12 (Positive hard covariance for the original coordinate)**.** *Perform the finite-range integrations in ψ and stop one scale before the prescribed cut. Write $C_k = C_{T,k}^{\Sigma}$. In this representation the original coordinate has Gaussian variance $\widetilde{V} = \eta_0^2 + w_0^{\mathsf{T}} S\Sigma S^{\mathsf{T}} w_0$ and cross covariance $v = T\ell$ with ψ. Throughout the matching rectangle, for sufficiently large fixed R and $n \geq n_0(R)$,*

$$v \in \operatorname{ran} C_k, \qquad \widetilde{V} - v^{\mathsf{T}} C_k^{\dagger} v \geq c\lambda n^2, \qquad \widetilde{B}_k = C_k - \frac{vv^{\mathsf{T}}}{\widetilde{V}} \succeq \frac{c\lambda n^2}{\widetilde{V}} C_k. \tag{7.39}$$

The profile $v/\widetilde{V}$, variance $\widetilde{V}$, and hard covariance $\widetilde{B}_k$ satisfy the Gaussian hypotheses of Proposition 7.3 for the curvature of the original terminal marginal.

Proof. Under the physical law one may sample the original coordinate given ϕ as $w_0^{\mathsf{T}}\phi + \eta_0 Z$, where Z is an independent standard normal. After the change of variable the conditional mean is $w_0^{\mathsf{T}} S\psi$, also when the physical interaction is included. Under the unweighted Gaussian in ψ the cross column is $\Sigma S^{\mathsf{T}} w_0 = Gw_0 = T\ell$. The identity $tG\Sigma^{-1}G = G - \varepsilon S A_c$ gives

$$t\widetilde{V} = \sigma^2 - \varepsilon(\eta_0^2 + w_0^{\mathsf{T}} S A_c w_0). \tag{7.40}$$

For the matrix in the error, put $M = G^{-1/2} A_c G^{-1/2}$, so $0 \preceq M \preceq I$. Then $SA_c = G^{1/2}(tI + \varepsilon M)^{-1} M G^{1/2}$ is positive and bounded above by CA_c. The variance bounds in Proposition 7.9 give $\eta_0^2 + w_0^{\mathsf{T}} A_c w_0 \leq C\lambda n^2$. Hence

$$|t\widetilde{V} - \sigma^2| \leq C|\varepsilon|\lambda n^2. \tag{7.41}$$

Lemma 7.11 applied to $W_k = C_k/t$ proves the first two estimates in (7.39), after the terminal aspect is fixed and R is large enough that $C|\varepsilon|$ fits in its margin. The range estimate and Cauchy–Schwarz on the support of C_k give the third. The lower bound is uniform in variance units. The relative factor $\lambda n^2/\widetilde{V}$ is of order $\lambda^2 x_k$.

Set $u = T\ell/\sigma^2$. The new mean profile $v/\widetilde{V} = (\sigma^2/\widetilde{V})u$ is a bounded positive multiple of the original one. Since $\ell^{\mathsf{T}} u = 1$, $\|\ell\|_2 \leq C$ and $\|u\|_{F_{n,p}} \leq Cn^{-2}$, the profile at the rescaled coordinate has $F_{n,p}$-norm $O(n^{-1})$ and $b_k^{-2}\|nv/\widetilde{V}\|_2^2 \asymp 1$. Moreover $n^2/\widetilde{V} = O(\lambda x_k)$, since shortening the horizon by one step preserves comparability with $\bar{g}_N$. For the hard field let D_{rel} be the released old slice and let $T_h = T - \sigma^2 uu^{\mathsf{T}}$. The transferred covariance is bounded above by $t(T + D_{\mathrm{rel}})$, by its spectral multiplier. Thus

$$0 \preceq \widetilde{B}_k \preceq t(T_h + D_{\mathrm{rel}}) + C|\varepsilon|\lambda n^2 uu^{\mathsf{T}}. \tag{7.42}$$

The old hard covariance and the released slice at a fixed fraction of scale n have the raw Gaussian norm moment of Section 2.3. A Gaussian with the last rank-one covariance has norm bounded by $C\sqrt{\lambda}\, n^{-1}|Z|$. Gaussian addition and conditional Jensen's inequality give $\mathbb{E}\exp(cn^2\|Y\|_{F_{n,p}}^2/\lambda) \leq C$ for $Y \sim N(0, \widetilde{B}_k)$.

Gaussian regression followed by addition of the integrated covariance now gives the original scalar density, up to a positive constant,

$$\exp[-a^2/(2\widetilde{V})]\, \mathbb{E}_{\widetilde{B}_k} F_k(av/\widetilde{V} + Y), \tag{7.43}$$

where F_k is the exact RG output for F_{in} in the ψ field. Indeed the conditional covariance of the full ψ equals $(\Sigma - C_k) + \widetilde{B}_k$. The regression data here are continuous in the external parameter

and independent of the running mass. In the physical ϕ coordinates the same hard covariance is $S\widetilde{B}_k S^{\mathsf{T}}$, for the representation with total covariance $S\Sigma S^{\mathsf{T}}$. This identifies precisely the subtraction estimated in (7.39). In the direct comparison using ϕ itself with covariance Σ, the column would still be Σw_0. □

The original scalar density is now expressed in the variables of the coupled RG chart. We use that density for mass matching and the natural terminal field of the same chart for the magnetic moment.

Corollary 7.13 (The prescribed matched point in dimension four)**.** *For all sufficiently large fixed R and $n \geq n_0(R)$, there is a physical temperature at which the original prescribed terminal marginal has zero curvature and $\chi_\Lambda^F \asymp_R n^2\sqrt{\log n}$.*

Proof. Use scales up to k for the backward shooting with terminal kinetic coordinate $z_k = 0$. For each fixed external pair (β, ε), Lemma 7.10 puts the physical microscopic family and its two off-shell coordinates in the joint initial chart. The RG map in ψ uses the transfer slices themselves. In particular it has the finite propagation, filtered bubble and annular estimates used in Proposition 4.7. The normalized microscopic filter of Lemma 7.10 has the same moment and recovery bounds needed for that proposition. Its coupled collar transport, the connected extraction of Proposition 3.12 and the summable kinetic source of Lemma 6.6 therefore apply as in Theorem 7.6. They give the same terminal coefficient box and inverse mass derivative, with $x_k \asymp_R (\log n)^{-1}$. The comparison coordinate used at the terminal step is given by (7.43). By Proposition 7.12 all its Gaussian inputs satisfy Proposition 7.3. In particular the leading derivative of its rescaled curvature with respect to the terminal mass is $b_k^{-2}\|nv/\widetilde{V}\|_2^2 \asymp 1$, with an $o(1)$ error. The terminal scalar graph has the required stable and activity derivatives. In the residual chart the inverse compensated kinetic and mass coordinates have the same first-variation bounds used in Theorem 7.6. Cauchy estimation on the enlarged kinetic disc gives the bounded normalized kinetic derivative of this scalar graph.

Proposition 6.5 gives continuous bare targets ν_0^*, z_0^*. Fix the thermal width as in Section 6, then take H large enough that the two kinetic faces enclose z_0^*, using the positive slope in Lemma 7.10 and the estimates $|z_0^*| + |z_0^{\mathrm{bare}}(\beta, 0)| \leq C_{A_0}\lambda x_0(0)$. The constants in the backward target bound can be taken before choosing H. In the bulk source of Proposition 4.5 the fixed microscopic head enters through M_0/x_0 and $M_j R/(b_j x_j)$. Its raw second moment is $O_{A_0,H}(\lambda x_0)$. The initial ratio is therefore small, and at later steps $(x_0/x_j)R/b_j$ stays bounded by the reciprocal comparison and geometric growth of b_j. The separately generated filtered head has the uniform bounds of Proposition 4.7. In the initial stable box a collar contribution $O_{A_0,H}(\lambda x_0)$ costs $O_{A_0,H}(M_R^{-2})$. Thus each fixed choice of widths fits the same box and source estimates by taking R large afterwards. The added scalar perturbation tends to zero in the thermal face margin as n grows at fixed R. The uniform thermal slope and the bounded normalized target consequently give the opposite thermal face signs. Corollary 6.7 sets both bare discrepancies to zero. At this point the off-shell reweightings restore F_{in} exactly. Its density in (7.43) has zero curvature for the original a_0.

For the magnetic moment return to the natural terminal field in ψ at cut k. Lemma 5.4 applies here because the aspect changes by a fixed factor and the pole displacement fixed at J still has the required order at k. More explicitly, conditioning $T + D_{\mathrm{rel}}$ on ℓ gives a hard covariance bounded above by

$$T_h + (I - u\ell^{\mathsf{T}})D_{\mathrm{rel}}(I - \ell u^{\mathsf{T}}). \tag{7.44}$$

This follows by using u as a candidate linear predictor. The added projected slice has the terminal moment, since $\ell^{\mathsf{T}} D_{\mathrm{rel}}\ell = O(\lambda n^2)$. Adding the released slice changes the old mean profile by $o(n^{-2})$ in the raw norm by covariance Cauchy–Schwarz and $\sigma^2 \asymp n^2/(\lambda x_k)$. The transferred natural

hard moment, profile and integrated unit source bounds then follow by the multiplier argument of that lemma. Propositions 7.2 and 7.4 use the coefficient bounds at the matched mass and give $\mathbb{E}(\mathbf{1}^{\mathsf{T}}\psi)^2 \asymp_R V^{3/2} g_k^{-1/2}$ for the full, interacting field. These moment bounds only use the size of the matched mass. They apply with the mass chosen by the original curvature graph.

In the physical field this linear functional is paired with $m = S^{-\mathsf{T}}\mathbf{1} = (I - \varepsilon c' B^{-1} A)\mathbf{1}$. The absolute short row bounds give $1 - C|\varepsilon| \le m_x \le 1 + C|\varepsilon|$ uniformly in the cube. Thus Am has entries bounded above and away from zero at fixed R. In physical units the conditional Hubbard–Stratonovich field has mean $\sqrt{\beta} A\sigma$ and Gaussian covariance A. Positivity of the spin two-point functions as in Lemma 7.5 compares the weighted spin variance to $\mathbb{E}(\mathbf{1}^{\mathsf{T}}\sigma)^2$, and the added noise variance is $m^{\mathsf{T}} A m = O_R(V)$. Converting the fixed field normalization and dividing by $V = n^4$ proves the assertion. □

7.6 Proof of the susceptibility theorem

Proposition 7.14 (Physical terminal data in the fractional dimensions)**.** *Fix $d \in \{1, 2, 3\}$. For all sufficiently large fixed R and $n \ge n_0(R)$, there is a physical temperature at which the prescribed terminal curvature vanishes. At this temperature the terminal interaction satisfies Assumption 7.1 with $g = g_N \asymp_{d,R} (\log V)^{-1}$.*

Proof. In dimension one take $p = 1/8$, so $p < s/2 = 1/4$ and $2p > s - 1 = -1/2$. The raw field norm with this weight admits the slice, principal-profile and integrated-source bounds in Section 2.3. The initial quadratic estimates in Proposition 2.16 and the interval specialization of Proposition 3.15 therefore use the same weight as the terminal integration.

Proposition 3.15 and Corollary 3.16 give the one-step and terminal coefficient estimates for the exact initial family of Proposition 5.1. The terminal graph follows from Proposition 7.3 and the conditioned Gaussian bounds of Section 2.3. Apply Theorem 6.2 and Corollary 6.3 to the joint physical initialization. The resulting real temperature has zero terminal curvature. The reciprocal increments give $g \asymp_{d,R} (\log V)^{-1}$. The activity estimate gives $g^{-9/4}\|K_N\| \le C g^{3/4} + C U d_0^N g^{-9/4} = o(1)$. Proposition 7.2 gives the pointwise terminal error and its first two field derivatives. The finite Neumann filter has the row and inverse bounds proved there. The prescribed principal variance is of order $\sqrt{V}/g$ at fixed R. By Lemma 2.8 the integrated slices contribute only $O_R(\sqrt{V})$ to it, so $\sigma_T^2 \asymp_R \sqrt{V}/g$. The short-range row gap gives $\|\ell\|_2 \le \|B_\Lambda^{-1/2}\| \le C$. Section 2.3 gives the exponential moment of the conditioned Gaussian field, and Corollaries 2.12 and 2.13 give the profile and the integrated uniform-source bounds at the same weights. At the bare mass the initial density is physical, and exact Gaussian integration preserves positivity. This verifies Assumption 7.1. □

Proof of Theorem 1.1. In dimensions one through three, take the physical temperature of Proposition 7.14. Proposition 7.4 gives the initial auxiliary moment of order $V^{3/2} g^{-1/2}$. The fixed canonical field rescaling changes only the constants for fixed R. The microscopic split has a positive diagonal and a bounded row, so Lemma 7.5 applies with β bounded above and below by positive constants. Its $O_R(V)$ error is of lower order. Dividing the spin moment by V proves the result.

In dimension four, Corollary 7.13 supplies a physical temperature for the prescribed coordinate and gives $\chi_\Lambda^F \asymp_R n^2 \sqrt{\log n}$. Since $V = n^4$, this has the asserted order. □

A Polynomial covariance slices and terminal conditioning

A.1 Scalar slices

Fix a common spectral upper bound S for the auxiliary precisions and write $4t/S = 2 - 2\cos\theta$, $0 \le \theta \le \pi$. We use the scalar periodization of [32, Lemma 2.1]. Let χ be a nonzero, nonnegative, even smooth function supported in $[-1/2, 1/2]$, $F = |\widehat{\chi}|^2$, and $I_F = \int_0^\infty TF(T)\,dT$. Define

$$P_T(\theta) = \sum_{h\in\mathbb{Z}} F(T(\theta + 2\pi h)), \qquad q_i(t) = \frac{4}{SI_F}\int_{I_i} TP_T(\theta)\,dT, \quad I_0 = (0, L], \quad I_i = [L^i, L^{i+1}] \ (i \ge 1). \tag{A.1}$$

Put $r_i = L^{-2i}$. For $J \ge 1$, write $r = L^{-2J}$, $g_<(t) = \sum_{i<J} q_i(t)$, $f(t) = t^{-1} - g_<(t)$, and $s_J(t) = tf(t) = 1 - tg_<(t)$.

Lemma A.1 (Polynomial covariance slices and their tail)**.** *The polynomials q_i are strictly positive on $[0, S]$, satisfy $q_i(t) \ge q_i(0)$ for $t < 0$, and sum to $1/t$ on $0 < t \le S$. For each fixed integer $m \ge 1$,*

$$\deg q_i \le C_L L^i, \qquad q_i(t) \le C_m r_i^{2m-1}(t + r_i)^{-2m}, \qquad q_i(t) \ge c/r_i \quad (0 \le t \le c_* r_i). \tag{A.2}$$

The upper bound holds on $[0, S]$. The polynomial s_J has degree at most $C_L r^{-1/2}$, and on $0 < t \le S$,

$$g_<(t) \asymp (t + r)^{-1}, \qquad 0 \le s_J(t) \le C_m\big(r/(t + r)\big)^m, \tag{A.3}$$

$$0 \le f(t)(1 - s_J(t)) \le C_m r^{2m-1}(t + r)^{-2m}. \tag{A.4}$$

Proof. Poisson summation expresses P_T as an even trigonometric polynomial of degree at most T. Its Fourier coefficients are proportional to the autocorrelation of χ at integer multiples of $1/T$, so they are nonnegative and vanish outside this degree. Integrating each period and using $\sum_h (\theta + 2\pi h)^{-2} = (2 - 2\cos\theta)^{-1}$ gives $\sum_{i\ge 0} q_i(t) = 1/t$ for $0 < t \le S$. The definition gives $q_i \ge 0$ on $[0, S]$. For $t < 0$ each cosine in its expansion becomes a hyperbolic cosine, so $q_i(t) \ge q_i(0)$. For $T \ge 1$, rapid decay gives $P_T(\theta) \le C_a(1 + T^2 t)^{-a}$ for every $a > 0$. For $T < 1$ it is a constant of order T^{-1}. Integration over the scale interval gives the degree and upper bounds in (A.2). The lower bound follows from $F(0) > 0$. In fact the integral of the zeroth period is strictly positive for every $0 \le \theta \le \pi$, since the entire Fourier transform of χ cannot vanish on a nonempty interval. Continuity then gives $\min_{[0,S]} q_i > 0$ at each fixed i.

The degree bound for s_J follows from that for the slices. The partial sum satisfies $g_< \le \min(t^{-1}, C/r)$. For $t \ge Dr$, geometric summation of (A.2) gives $f(t) \le C_m r^{2m-1} t^{-2m}$, hence $g_< \ge (2t)^{-1}$ for fixed large D. For $t \le Dr$ the zeroth period in the last integrated slice gives $q_{J-1}(t) \ge c_{D,L}/r$. After setting $T = L^J z$, this follows from the positive minimum of $\int_{1/L}^1 zF(yz)\,dz$ for $0 \le y = L^J\theta \le C_D$. Finally $0 \le tf \le 1$ and $1 - s_J = tg_<$ control the residual for $t \le r$, and the bound on f controls it for $t \ge r$. □

A.2 Comparison of the joint Gaussian fields

In this appendix the joint Gaussian fields have unnormalized covariances, so observed covariances are multiplied by a_R at the end. Write $Q_c = \mathcal{Q}_c$, $D = \mathcal{D}$, $U = \mathcal{L}_\Lambda$, $A = U^* D^{-1} U$, and $B = B_\Lambda$. Let $Z = X_< + X_T$ be the sum of independent Gaussian fields with covariances $g_<(Q_c)$ and $f(Q_c)$. Write $a = \ell^{\mathsf{T}}\mathcal{O}Z$, $Y = \mathcal{O}X_T$, and $C_T = \mathcal{O}f(Q_c)\mathcal{O}^*$.

Lemma A.2 (Conditioning the joint Gaussian field)**.** *For $1 \le c < c_{*,\Lambda}$, on tests of finite inverse energy, $\mathrm{Cov}(Z \mid a = 0) \le CQ_1^{-1}$ uniformly in c. The observed terminal field satisfies*

$$\mathcal{S}_\ell(C_T) \le \mathrm{Cov}(Y \mid a = 0) = \mathcal{O}\big[s_J(Q_c)\,\mathrm{Cov}(Z \mid a = 0)s_J(Q_c) + f(Q_c)(I - s_J(Q_c))\big]\mathcal{O}^*. \tag{A.5}$$

Proof. The quadratic-form identities follow first on regularized or spectrally truncated tests and then by Gaussian L^2 limits. Regress X_T on Z to obtain $X_T = s_J(Q_c)Z + \eta$, where η is independent of Z and has covariance $f(Q_c)(I - s_J(Q_c))$. This gives the displayed identity after conditioning on $a = 0$. The lower bound follows by rank-one regression, since $\mathrm{Cov}(Y, a) = C_T\ell$ and $\mathrm{Var}\, a \geq \ell^{\mathsf{T}} C_T \ell$.

To compare the conditioned joint covariance to Q_1^{-1}, condition first on the full observed field $h = U^* Z_a$, where Z_a denotes the auxiliary component. Given h, the two blocks are independent. Their conditional means are $D^{-1}UA^{-1}h$ and $cB^{-1}h$; their covariance blocks are $R_A = D^{-1} - D^{-1}UA^{-1}U^*D^{-1}$ and cB^{-1}. Indeed, completing the physical square gives a tilt of the auxiliary Gaussian depending only on h, and ordinary Gaussian regression gives the first mean and covariance. Denote these combined mean and covariance operators by T_c and J_c, respectively. Then

$$\mathrm{Cov}(Z \mid a = 0) = J_c + T_c \mathcal{S}_\ell(V_c) T_c^*, \qquad Q_1^{-1} = J_1 + T_1 V_1 T_1^*. \tag{A.6}$$

Section 2.3 gives $\mathcal{S}_\ell(V_c) \leq CV_1$, $\|V_1\| \leq C\varrho^{-s}$ and $|c - 1| \leq C\varrho^s$. We have $J_c \leq CJ_1$, while $(T_c - T_1)V_1(T_c - T_1)^*$ is confined to the physical block and bounded there by $C\varrho^s B^{-2} \leq CB^{-1}$. Expanding $T_c = T_1 + (T_c - T_1)$ and using Cauchy–Schwarz thus gives $\mathrm{Cov}(Z \mid a = 0) \leq CQ_1^{-1}$. □

A.3 Inverse energy on polynomial tests

In the fractional channel let $N = N_{\mathrm{sm}}$, and let W have continuous auxiliary component $\sqrt{w_\alpha(u)} D_R^{N/2} P(u, \cdot)$, where P has degree less than K in u. Suppose the coefficient vectors of P, the white component, and the physical component are supported in a common cube of side $M \geq R$. With $0 < \alpha = d/4 < 1$, we claim

$$\langle W, Q_1^{-1} W \rangle \leq C\{K^2 + K^{2-2\alpha}(M/R)^{2\alpha}\}\|W\|^2. \tag{A.7}$$

For the polynomial part, write $t = 1 - u$ and $\nu(t) = t^\alpha (1 - t)^{N-\alpha-1}$. For any scalar polynomial p of degree below K,

$$\int_0^1 \frac{|p|^2}{t} \nu \, dt \leq CK^2 \|p\|^2_{L^2(\nu)}, \qquad \int_0^1 \frac{|p|}{t} \nu \, dt \leq CK^{1-\alpha} \|p\|_{L^2(\nu)}. \tag{A.8}$$

Indeed, for $a = N - \alpha - 1$, the shifted Jacobi polynomial $P_j^{(\alpha,a)}(1 - 2t)$ has value $\binom{j+\alpha}{j} = O((j+1)^\alpha)$ at zero and squared norm comparable to $(j+1)^{-1}$. Its ratio to the endpoint value is $\sum_{h=0}^{j} (-j)_h (j + \alpha + a + 1)_h t^h / ((\alpha + 1)_h h!)$. For $j < K$ and $t \leq K^{-2}$, the absolute sum is bounded by $\sum_{h \geq 0} C^h / (h!)^2$. Cauchy–Schwarz in the orthonormal expansion gives $\sup_{t \leq K^{-2}} |p(t)| \leq CK^{1+\alpha} \|p\|_{L^2(\nu)}$. Integrate this estimate up to K^{-2}. Above K^{-2} use $t^{-1} \leq K^2$ for the squared estimate, and Cauchy–Schwarz with $\int_{K^{-2}}^1 t^{\alpha-2}(1 - t)^a \, dt \leq CK^{2-2\alpha}$ for the linear estimate.

For the spatial part, a vector v supported in a cube of side $M \geq R$ satisfies, with normalized Fourier measure,

$$\int \omega(k)^{-\alpha} |\widehat{v}(k)|^2 \, dk \leq C(M/R)^{2\alpha} \|v\|_2^2. \tag{A.9}$$

On $|k| \leq M^{-1}$ use $\|\widehat{v}\|_\infty^2 \leq CM^d \|v\|_2^2$ and integrate $(R|k|)^{-2\alpha}$, which is integrable since $d > 2\alpha$. On the complement use the supremum of the multiplier. By the variational formula for inverse energy the finite-volume form is bounded by its bulk counterpart on the zero extension of the physical component. Completing the bulk auxiliary square gives

$$\langle W_a, D^{-1} W_a \rangle + \big\| (B_\infty - A_\infty)^{-1/2} (W_p + \mathcal{L}^* D^{-1} W_a) \big\|_2^2. \tag{A.10}$$

The first polynomial inequality bounds the first term by $CK^2\|W_a\|^2$, since $1-u\widehat{D}_R \geq 1-u$. Put $d_k = \widehat{D}_R(k)$. The second polynomial inequality bounds the squared modulus of the continuous contribution to $\widehat{\mathcal{L}^* D^{-1} W_a}(k)$ by

$$CK^{2-2\alpha}|d_k|^{2N} \int_0^1 w_\alpha(u)|\widehat{P}(u,k)|^2\,du. \tag{A.11}$$

Use $|d_k|^{2N} \leq |d_k|^N$ and apply the spatial bound to $D_R^{N/2}P(u,\cdot)$, supported in a cube of side $M + C_N R$. Since $B_\infty - A_\infty = \zeta(I-D_R)^\alpha$, integration gives the second term of (A.7). The same spatial bound applies directly to the physical and white sources.

Corollary A.3 (Resolvent bound for the conditioned terminal field)**.** *Let $1 \leq d \leq 4$, $1 \leq c < c_{*,\Lambda}$, and take the terminal cut $r = L^{-2J} \asymp M_*(R/n)^2$ of Section 2.3.2. Let Y and C_T have the unnormalized convention of Lemma A.2. For each fixed integer $m \geq 1$, uniformly in c,*

$$a_R \mathcal{S}_\ell(C_T) \leq a_R \operatorname{Cov}(Y \mid a = 0) \leq C_m a_R r^{2m-1} \mathcal{O}(Q_c + rI)^{-2m}\mathcal{O}^*. \tag{A.12}$$

Proof. For the terminal cut choose $K \asymp_L r^{-1/2}$ larger than the degree of s_J plus one. For every physical test v the vector $W = s_J(Q_c)\mathcal{O}^* v$ has the stated form. Each application of the continuous diagonal multiplies the polynomial by $I - uD_R$, and the off-diagonal input adds a term independent of u. Each block increases the spatial support by at most $C_N R$, so $M \leq n + C_N RK \leq C_{M_*,L,N} RK$ at the terminal aspect. Thus (A.7) gives $\langle W, Q_1^{-1} W\rangle \leq Cr^{-1}\|W\|^2$. For the local split $D = I$, the Schur complement $B - A = \zeta(I-D_R)_\Lambda$ has lowest eigenvalue at least $c(R/n)^2$ by Section 2.3. The couplings are bounded, so completing the same square yields $\|Q_1^{-1}\| \leq C(n/R)^2$, which gives the same inequality at this aspect. Insert this bound in the regression formula of Lemma A.2. Use $s_J(t)^2 \leq C_m(r/(t+r))^{2m}$ in its first term and (A.4) in its second term. Multiplication by a_R gives the stated bound. □

B Resolvent derivatives at the local endpoint

Lemma 2.11 bounds discrete derivatives on a stencil at depth $\rho_b(x)$ from a face. The calculation below uses the notation of Section 2.3.5.

Proof of Lemma 2.11. Put $M_r = \mathcal{Q}_c + rI$, $D = B/c + rI$, and $H = D - A/(1+r)$ in the bulk. Write

$$H = d_p I - K_r, \quad d_p = c^{-1} + r, \quad K_r = Q/c + A/(1+r). \tag{B.1}$$

Here Q has zero diagonal. The bulk Fourier multiplier of H is comparable to $r + \min\{|Rk|^2, 1\}$ by the local Schur comparison, and the same lower bound holds for the smallest eigenvalue of the bulk matrix $M_r(k)$. In fact the block triangular factorization giving the Schur complement has uniformly bounded factors and inverses.

Let $U_l = M_r^{-l} f$ with $\|f\|_2 \leq 1$, and write u_l for its physical component and $w_l = \mathcal{L}_\Lambda^*(U_l)_{\mathrm{aux}}$ for its observation. Extend physical vectors by zero. Set $\widetilde{u}_l = u_l - d_p^{-l} f_{\mathrm{phys}}$. For $l \geq 2$,

$$\|\widetilde{u}_l\|_\infty + \|w_l\|_\infty \leq C_l R^{-2} r^{1-l}. \tag{B.2}$$

At $l = 1$ the upper bound $CR^{-2}(1 + \log(1/r))^{1/2}$ suffices. To prove these bounds, put $D_0 = \operatorname{diag}((1+r)I, d_p I)$ on all channels. The Duhamel expansion of $e^{-t(D_0 - K)}$ in the entrywise nonnegative coupling $K = D_0 - M_r$ shows that $M_r^{-l} - D_0^{-l}$ is entrywise nonnegative and bounded by its bulk matrix elements. Indeed, the removed term is the path with no coupling; the other terms are

positive and finite-volume paths are a subset of bulk paths. Integrate the expansion with weight $t^{l-1}/(l-1)!$. At $|Rk| \le c_1$, the squared bulk row norm is controlled by integrating $(r+|Rk|^2)^{-2l}$, which gives $C_l R^{-4} r^{2-2l}$ for $l > 1$ and $CR^{-4}(1+\log(1/r))$ at $l = 1$. At the remaining frequencies the uniform matrix gap and the resolvent identity bound $\|M_r(k)^{-l} - D_0^{-l}\|$ by $C_l(|\widehat{Q}(k)| + \|\widehat{\mathcal{L}}(k)\|)$. Its squared integral is $O_l(R^{-4})$ by the size and smoothing of the split. This proves the row estimates for $U_l - D_0^{-l} f$. The absolute row norm of $\mathcal{L}^*$ is bounded and the ℓ^2 norm of each row is $O(R^{-2})$; applying $\mathcal{L}_\Lambda^*$ gives the asserted bound on w_l.

Let $d_x = \min\{b, \operatorname{dist}(x, \Lambda^c)\}$. The absolute-leg bound proved above already suffices if $d_x \le C_{m,\gamma} R$. Its proof also applies to every site when $b = O(R)$, since only the local support and weights of the observation leg are used. At interior sites the block equations give

$$\begin{aligned} H\widetilde{u}_l &= \widetilde{u}_{l-1} + (1+r)^{-1} w_{l-1} + d_p^{-l} K_r f_{\mathrm{phys}}, \\ H w_l &= (1+r)^{-1}\big(A\widetilde{u}_{l-1} + D w_{l-1} + d_p^{-(l-1)} A f_{\mathrm{phys}}\big). \end{aligned} \tag{B.3}$$

For the second equality first use $(1+r)w_l = Au_l + w_{l-1}$ and $Du_l = u_{l-1} + w_l$, and commute the bulk scalar convolutions where the stencils lie in the cube. Let $k_0 > |\gamma|/2$ be a fixed integer. Iterating the displayed equations for $j \le k_0$, in a ball of radius $d_x/2$ with a reserved $O(k_0 R)$ collar, gives

$$\|H^j \widetilde{u}_l\|_{\infty,\, B(x,d_x/2)} \le C_{j,l} R^{-2} r^{1-l+j}, \qquad l - j \ge 2. \tag{B.4}$$

Indeed, the finite-range convolutions multiplying the two previous components commute with H in this interior calculation and have bounded absolute row norms. The source terms, including any subsequent H-factors on them, are bounded by $C_{j,l} R^{-2}$ using the row ℓ^2 norms of K_r, A and the bounded absolute row norms of H. Thus the exponent follows from (B.2).

We convert the interior bounds on powers of H to derivative bounds using a cutoff Green function. Set $\delta = (R/d_x)^2 \ge r$ and $\widetilde{H} = H + \delta I$, and use a smooth lattice cutoff χ equal to one on $B(x, d_x/4)$ with support in $B(x, d_x/2)$ and differences of order j bounded by $C_j d_x^{-j}$. Let $h_x = \widetilde{H}^{-k_0} A \nabla^{\gamma,*} \delta_x$ in the bulk. Then

$$\|h_x\|_1 \le C\delta^{-k_0} d_x^{-|\gamma|}, \qquad \|\nabla^j[(1-\chi)h_x]\|_1 \le C_j \delta^{-k_0} d_x^{-|\gamma|-j}, \quad 0 \le j \le 2k_0. \tag{B.5}$$

For these kernel bounds, decompose near zero in the Fourier variable $\theta = Rk$ into smooth dyadic bands of size $T \ge \sqrt{\delta}$ and a final ball of size $\sqrt{\delta}$. The symbol of $\widetilde{H}$ is comparable to $\delta + |\theta|^2$ there; its derivatives at scale T are bounded by $C_a T^{2-a}$, with the zero-th order bound also applying on the final ball. The resulting row $\nabla^j h_x$ on one band is bounded by

$$C_J R^{-4-|\gamma|-j} T^{4+|\gamma|+j-2k_0} (1 + T|y-x|/R)^{-J}. \tag{B.6}$$

This follows by integration by parts after rescaling the band, including the discrete difference multipliers. At high frequencies $\widetilde{H}$ has a uniform gap and bounded derivatives in θ; the factor $\widehat{A}$ gives arbitrarily many fixed powers of $(1+|\theta|)^{-1}$ by choosing the split smoothing order large enough. The exponentially small white term is harmless on the scaled torus. Thus the high-frequency row is bounded by $C_J R^{-4-|\gamma|-j}(1+|y-x|/R)^{-J}$. Summing the low-frequency ℓ^1 bounds at $j = 0$ uses $2k_0 > |\gamma|$. For $j > 0$, in the tail $|y-x| \gtrsim d_x$ reserve a large fixed power of $(1 + T d_x/R)^{-1}$ before summing. The discrete product rule with the cutoff gives (B.5).

Since K_r is even and has range $O(R)$, $\widetilde{H}$ is a constant of size $O(\delta)$ plus a bounded weighted sum of centered second differences over jumps of size $O(R)$. A product of j such differences costs at most $C_j R^{2j}$ times the largest ℓ^1 norm of a lattice difference of order $2j$, by telescoping each jump into coordinate steps. Thus the second bound in (B.5) gives

$$\|\widetilde{H}^{k_0}[(1-\chi)h_x]\|_1 \le C \sum_{j=0}^{k_0} \delta^{k_0-j} R^{2j} \delta^{-k_0} d_x^{-|\gamma|-2j} \le C d_x^{-|\gamma|}. \tag{B.7}$$

The cutoff identity

$$\nabla^\gamma A\widetilde{u}_l(x) = \langle \chi h_x, \widetilde{H}^{k_0}\widetilde{u}_l\rangle + \langle \widetilde{H}^{k_0}[(1-\chi)h_x], \widetilde{u}_l\rangle \tag{B.8}$$

now bounds its left side by $C_l R^{-2} r^{1-l} d_x^{-|\gamma|}$ when $l \geq k_0 + 2$. Indeed the first term uses $\|\widetilde{H}^{k_0}\widetilde{u}_l\|_{\infty,\,\mathrm{supp}\,\chi} \leq C_l \delta^{k_0} R^{-2} r^{1-l}$ by the interior powers and $r \leq \delta$, and the second term uses (B.2). All pairings are with zero extensions.

Finally, exactly on the cube,

$$w_m = (1+r)^{-m} w_0 + \sum_{l=1}^{m} (1+r)^{-(m-l+1)} A_\Lambda u_l. \tag{B.9}$$

For $l \geq k_0 + 2$, use the cutoff bound on the $\widetilde{u}_l$ part. For smaller l the bound $\|\nabla^\gamma A\|_1 \leq C_\gamma R^{-|\gamma|}$ and the first-resolvent or higher-resolvent row bounds suffice: the loss $(d_x/R)^{|\gamma|} \leq r^{-|\gamma|/2}$ and the first-resolvent logarithm are absorbed by taking $m > k_0 + 2 + |\gamma|/2$. The direct-source rows $\nabla^\gamma w_0$ and $d_p^{-l}\nabla^\gamma A f_{\mathrm{phys}}$ are bounded by $CR^{-2-|\gamma|}$, by the smoothing of the split. Altogether $|\nabla^\gamma w_m(x)| \leq C_m R^{-2} r^{1-m} d_x^{-|\gamma|}$. Squaring this dual norm and multiplying by $a_R r^{2m-1}$ gives the claim since $a_R = \lambda R^2$ and $\lambda R^{-2} r = \lambda b^{-2}$. □

C Gaussian integration at a fixed kinetic coefficient

We collect the identities and bounds for integrating a quadratic interaction when its kinetic coefficient ranges over a fixed disc. The square-completion formula also gives the terminal comparison in Section C.4. The trajectory of Section 4 uses the finite-propagation reference on its shrinking kinetic domain.

C.1 Exact bulk compensation of the pure quadratic terms in dimension four

For a purely quadratic input, completing the Gaussian square gives an exact version of the second-order compensation in Section 3.8. The first two Taylor coefficients of the bulk Fourier symbol determine the mass and kinetic coordinates. This calculation allows a fixed kinetic coefficient as the quartic coupling tends to zero and identifies the corresponding change in the covariance.

Lemma C.1 (Exact compensation of the bulk quadratic two-jet)**.** *Let C be a finite-range bulk slice and $W = W_j^{\mathrm{cov}}$ the integrated bulk covariance at step j. Write*

$$\widehat{H}(k) = b_j^{-2} u + z^{\mathrm{loc}}|k|^2 + O(|k|^4) \tag{C.1}$$

for the symbol of the incoming quadratic Hessian. Put $a = b_j^{-2}\widehat{W}(0)$ and $m = b_j^{-4}\sum_x x_1^2 W(x)$. The compensated mass and kinetic coefficients are

$$\nu = \frac{u}{1-au}, \qquad z = \frac{z^{\mathrm{loc}} - \frac{1}{2}mu^2}{(1-au)^2}. \tag{C.2}$$

Under pure quadratic Gaussian integration and with $W^+ = W + C$, the compensated dimensionful mass $b_j^{-2}\nu$ and the compensated kinetic coefficient z are unchanged. In particular there is only the engineering rescaling of ν when b_j is replaced by b_{j+1}. The identities are uniform and analytic on the low-jet domain $|au| \leq 1/4$, $|z^{\mathrm{loc}}| \leq z_0$, wherever the Gaussian quadratic integration exists. The corresponding compensated coefficient of a single quartic insertion, to first order in that insertion, is $g = g^{\mathrm{loc}}(1-au)^{-4}$.

Proof. For a finite Gaussian vector with covariance C and symmetric Hessian H, write $X = C^{1/2}Z$ with Z standard Gaussian. Completing the square in Z gives

$$\begin{aligned}\mathbb{E}_C\Big[e^{-\langle\phi+X,H(\phi+X)\rangle/2}F(\phi+X)\Big] &= \det(I+C^{1/2}HC^{1/2})^{-1/2}\\ &\quad\times e^{-\langle\phi,H^+\phi\rangle/2}\mathbb{E}_{\widetilde{C}}F\big((I+CH)^{-1}\phi+X\big),\\ H^+ &= H(I+CH)^{-1},\\ \widetilde{C} &= C^{1/2}(I+C^{1/2}HC^{1/2})^{-1}C^{1/2}.\end{aligned} \tag{C.3}$$

This holds when $I + C^{1/2}HC^{1/2}$ is positive definite and requires no inverse of C. It gives the bulk symbol identities by working with periodic convolution matrices. For scalar symbols, the identity

$$\frac{\widehat{H}^+}{1-(\widehat{W}+\widehat{C})\widehat{H}^+} = \frac{\widehat{H}}{1-\widehat{W}\widehat{H}} \tag{C.4}$$

holds near $k = 0$ whenever the denominators at zero are nonzero. Taking the Taylor coefficients and using $\widehat{W}(k) = b_j^2 a - \frac{1}{2}b_j^4 m|k|^2 + O(|k|^4)$ gives (C.2). For the quartic, the degree of the inserted polynomial is unchanged only by the term of the expectation without Gaussian contractions. Each zero-momentum external leg is then multiplied by $(1+\widehat{C}(0)\widehat{H}(0))^{-1}$. The same denominator identity to the fourth power yields the stated quartic compensation. Inversion of (C.2) gives

$$u = \frac{\nu}{1+a\nu}, \qquad z^{\mathrm{loc}} = \frac{z}{(1+a\nu)^2} + \tfrac{1}{2}mu^2. \tag{C.5}$$

Since $|a| + |m| \le C\lambda$, differentiation gives $\partial_\nu u = 1 + O(\lambda|u|)$, $|\partial_\nu z^{\mathrm{loc}}| \le C\lambda(|z| + |u|)$, and $|\partial_\nu g^{\mathrm{loc}}| \le C\lambda|g|$ uniformly on this domain. □

For the canonical bulk slice, the rational chart has a fixed kinetic disc when the Hessian is represented by the local Laplacian. Indeed let $\omega(k) = 1 - \widehat{D}_R(k)$ and $r = (R/b_j)^2$. The shifted-precision bound in Section 2.4 bounds the bulk auxiliary inverse by $C(r+\omega(k))^{-1}$ in operator norm in dimension four. Applying the scalar cutoff bound before integrating over k, and using boundedness of the observation vector, gives

$$0 \le \widehat{C}(k) \le C_m\lambda R^2 r^{2m-1}(r+\omega(k))^{-2m} \le C_m\lambda R^2(r+\omega(k))^{-1}. \tag{C.6}$$

The last bound also holds for the merged slice at its largest scale. To see this, sum the preceding bound with the corresponding geometric $r_i \ge r$. For $r_i \ge \omega(k)$ sum r_i^{-1}, and for $r_i < \omega(k)$ sum $r_i^{2m-1}\omega(k)^{-2m}$. Consequently, for $\widehat{H}(k) = b_j^{-2}u + z^{\mathrm{loc}}R^{-2}\omega(k)$ with a bounded adjustment of the normalization of z^{loc},

$$\sup_k |\widehat{C}(k)\widehat{H}(k)| \le C\lambda(|u| + |z^{\mathrm{loc}}|). \tag{C.7}$$

For small λ, this bounds the denominator of the Gaussian identity uniformly away from zero in a kinetic disc of radius independent of g_j, including for the holomorphic finite-dimensional Gaussian integrals on that disc.

In particular the pure mass–kinetic contraction need not be summed as a source proportional to $\lambda|u||z|$. It is absorbed exactly by the chart. However, completing the square in a bulk slice with a quadratic input generally changes its propagation. For a nonzero even finite-range convolution C and a nonzero even finite-range Hessian H of zero row sum, $\widehat{C}/(1+\widehat{C}\widehat{H})$ cannot be a Laurent polynomial. If it were, $1+\widehat{C}\widehat{H}$ would divide $\widehat{C}$ in the Laurent polynomial ring, and hence divide $1 = (1+\widehat{C}\widehat{H}) - \widehat{H}\widehat{C}$. A Laurent polynomial unit is a monomial. Symmetry under inversion of all

variables makes this monomial constant, and its value at $k = 0$ is one, which contradicts $\widehat{C}\widehat{H} \neq 0$. Thus this exact compensation of the projected quadratic contribution does not justify using the dressed covariance in the connected finite-range step.

Lemma C.2 (Scale-local expansion of the dressed ordinary slice)**.** *Let C be an ordinary covariance slice on a four-dimensional cube at scale $b \geq R$. Fix $1/2 < p < 1$, put $\rho(x) = \min\{b, R + \operatorname{dist}(x, \Lambda^c)\}$ and $w(x) = (\rho(x)/b)^p$, and use the raw test norm F_b. Suppose the slice has range Kb and the annular stencil variance bounds of Section 2.3 with this weight and up to at least $\max\{2, P\}$ orders. With $J_R = R^{-2} P_\Lambda L_R P_\Lambda^*$, set*

$$H = zJ_R + b^{-2}uI, \qquad t = |z| + |u|. \tag{C.8}$$

For $C_0\lambda t < 1$, the dressed covariance in (C.3) *has the absolutely convergent expansion*

$$\widetilde{C} = \sum_{m \geq 0} C^{(m)}, \qquad C^{(m)} = (-CH)^m C. \tag{C.9}$$

Each term has range at most $K_0(m+1)b$ and satisfies

$$|\nabla_x^\gamma \nabla_y^\eta C^{(m)}(x,y)| \leq C_0\lambda b^{-2}(C_0\lambda t)^m \frac{w(x)w(y)}{\rho(x)^{|\gamma|}\rho(y)^{|\eta|}}, \qquad |\gamma|, |\eta| \leq P. \tag{C.10}$$

The same assertion for one unit coefficient derivative in (u, z) holds with $(C_0\lambda t)^m$ replaced by $m(C_0\lambda)^m t^{m-1}$, for $m \geq 1$. In particular both series converge with an arbitrarily prescribed fixed exponential weight in the range index m when λt is sufficiently small. For real coefficients in this disc the Gaussian with covariance $\widetilde{C}$ has the following joint fluctuation bound on any finite family $\mathcal{P}$ of different scale-b tiles. For $Y_P = F_{b,P^*}(X)/\ell_b$ on fixed tile enlargements and bounded nonnegative integers k_P, one has

$$\mathbb{E}_{\widetilde{C}}\left[\prod_{P \in \mathcal{P}} Y_P^{k_P} \exp\left(\frac{c}{\lambda}\sum_{P \in \mathcal{P}} Y_P^2\right)\right] \leq C^{|\mathcal{P}|}\lambda^{\sum_P k_P/2}. \tag{C.11}$$

Proof. Let $\|f\|_w = \sup_x |f(x)|/w(x)$. Covariance Cauchy–Schwarz gives the product of the two annular stencil standard deviations in every mixed entry. In particular

$$|\nabla_x^\gamma C(x,y)| \leq C\lambda b^{-2}\frac{w(x)}{\rho(x)^{|\gamma|}}w(y). \tag{C.12}$$

For $\nabla_x^\gamma(CJ_R)(x,y)$, extend the second variable by zero. At depth larger than $2\sqrt{d}R$, write the symmetric walk as an average of centered second differences of steps of length at most $\sqrt{d}R$. Each such difference is bounded, by twice telescoping in coordinate directions, by CR^2 times the largest second stencil bound within distance $\sqrt{d}R + 2$ of y, where w and ρ are comparable to their values at y. Closer to the boundary, use the zeroth bound on the extended row and $\rho(y) \asymp R$. The normalization R^{-2} gives in both cases

$$|\nabla_x^\gamma(CJ_R)(x,y)| \leq C\lambda b^{-2}\frac{w(x)}{\rho(x)^{|\gamma|}}\frac{w(y)}{\rho(y)^2}. \tag{C.13}$$

The scalar mass term has the same bound since $b^{-2} \leq \rho(y)^{-2}$. All these rows have support within a fixed multiple of b from the stencil anchor. A ball of this radius meets $O(b^3)$ sites at each given boundary depth below b, and has at most $O(b^4)$ sites of larger depth. Therefore

$$\sum_{|y-x| \leq K_1 b}\frac{w(y)^2}{\rho(y)^2} \leq Cb^3 b^{-2p}\sum_{r=1}^{\lceil b \rceil}(r+R)^{2p-2} + Cb^2 \leq Cb^2. \tag{C.14}$$

Multiplying the row estimate by $|f(y)| \le w(y)\|f\|_w$ gives $\|CHf\|_{F_b} \le C\lambda t\|f\|_w$, and hence the same bound with F_b on the left replaced by $\|\cdot\|_w$.

At the right end of $(CH)^m C$, the $\|\cdot\|_w$ norm of the η-stencil of column y of C is at most $C\lambda b^{-2} w(y)\rho(y)^{-|\eta|}$. Apply the weighted bound to the intervening factors, and apply the F_b bound for the leftmost factor. This proves (C.10), including the original mixed bound for $m = 0$. Each factor has range bounded by a fixed multiple of b. Marking one of the m factors and using the same estimate on it with $t = 1$ proves the derivative assertion.

For real coefficients, the nonzero eigenvalues of CH are those of the symmetric $C^{1/2}HC^{1/2}$ by the determinant identity for products. They have absolute value at most the weighted operator norm of CH. It follows that the hard resolvent is positive and $\widetilde{C} \le (1 - C_0\lambda t)^{-1} C$. Absorb each power of $Y_P/\sqrt{\lambda}$ into $C\exp(cY_P^2/\lambda)$. The remaining exponential of the sum of squared norms is convex. Add an independent centered Gaussian to attain the dominating covariance and apply conditional Jensen to that exponential. For the dominating multiple of C, color the tiles in a bounded number of residue classes so that the fixed enlargements within one color are separated by more than the finite range. Hölder over the colors and independence within colors give a product of single-tile bounds for small enough c, proving the joint bound. □

The expansion controls a single dressed covariance link and its coefficient variation at ordinary scales. For real coefficients it also gives a joint Gaussian majorant. To estimate products with activities, we express the quadratic tilt as a product of local densities under the original Gaussian and keep track of the shifted background separately.

C.2 Local kinetic insertions in the Gaussian integral

A local insertion of the quadratic density contains two fluctuation fields. Its cost is small at either field radius, uniformly as the quartic coupling tends to zero. This permits us to use the original finite-range Gaussian integral and the regulator estimates of Lemma 3.10 for products of insertions.

Lemma C.3 (A kinetic insertion costs $O(\lambda)$ at both radii)**.** *Let $d = 4$, $1/2 < p < 1$, and use the covariance, raw norm and tile geometry of Lemma C.2. Write $H = zJ_R + b^{-2}uI$, with $J_R = R^{-2}P_\Lambda L_R P_\Lambda^*$, and let $|z|+|u| \le t_0$ on a fixed complex domain. There are quadratic polynomials E_P, supported on the fixed fine stencils P^*, such that*

$$\tfrac{1}{2}\langle\zeta, H\zeta\rangle = \sum_P E_P(\zeta). \tag{C.15}$$

Let $v_P = e^{-E_P} - 1$. In any product satisfying the hypotheses of Lemma 3.10, insert v_P on each tile of a set A of distinct fine labels, without changing the field factor on that label. At both radii the product bound of that lemma holds, with its strict reserve and an additional factor $(C_L\lambda(|z|+|u|))^{|A|}$. The derivative of the inserted product in a unit (z,u) direction has the bound obtained by replacing this factor, for $|A| \ge 1$, by

$$|A|(C_L\lambda)^{|A|}(|z|+|u|)^{|A|-1}. \tag{C.16}$$

If the product of the original field factors depends on the coefficients, its derivative is estimated separately by the same product bound.

Proof. Extend ζ by zero off Λ. Assign to P the contribution

$$\frac{z}{4R^2}\sum_{x\in P, y\in\Lambda} D_R(x-y)(\zeta_x-\zeta_y)^2 + \frac{z}{2R^2}\sum_{x\in P, y\notin\Lambda} D_R(x-y)\zeta_x^2 + \frac{u}{2b^2}\sum_{x\in P}\zeta_x^2. \tag{C.17}$$

The ordered interior pairs and the two orientations of the exterior pairs give the claimed exact sum. Write $M = F_{b,P^*}(\zeta)$, $Y_P = M/\ell_b = bM$. For pairs within $O(R)$ of the physical boundary, both fields are bounded by $CM(R/b)^p$, including the zero extension. A tile meets at most Cb^3R sites in this region. Away from this collar, with depth parameter $\rho = \min\{b, R + \operatorname{dist}(x, \Lambda^c)\}$, summing first differences bounds the pair difference by $CMRb^{-p}\rho^{p-1}$. The row sum of D_R is one. As in the energy estimate of Lemma C.9, layer summation bounds the kinetic contribution by

$$C|z|M^2\big[b^2(R/b)^{2p-1} + b^{3-2p}\textstyle\sum_{r=1}^{b}(R+r)^{2p-2} + b^2\big] \le C|z|Y_P^2. \tag{C.18}$$

The mass term is at most $C|u|Y_P^2$, by the tile volume bound. This proves both $|E_P| \le C(|z|+|u|)Y_P^2$ and $|D_{z,u}E_P| \le CY_P^2$. The integral formula for an exponential difference gives the extra polynomial cost $C(|z| + |u|)Y_P^2$ times $e^{Ct_0Y_P^2}$ at each mark. Its first variation replaces the coefficient by C. These insertions have no background field derivatives. In the proof of Lemma 3.10, the additional exponential in the fluctuation seminorm is absorbed by its $O(\lambda)$ Gaussian moment. The same coloring and Hölder argument there retains a factor $C_L\lambda$ for each marked Y_P^2. No background regulator is used by the new insertions. Summing the mark that carries the coefficient derivative proves the last estimate. □

Here the complex variation is in the density at a fixed field argument. For real coefficients set $S = (I + CH)^{-1}$. The square-completion identity can be written as

$$\mathbb{E}_C\big[e^{-\langle\phi+\zeta,H(\phi+\zeta)\rangle/2}F(\phi+\zeta)\big] = e^{-\langle\phi,H^+\phi\rangle/2}\mathbb{E}_C\big[e^{-\sum_P E_P(\zeta)}F(S\phi+\zeta)\big]. \tag{C.19}$$

The expectation of the density at zero background gives the determinant in the square-completion formula. Holding $\psi = S\phi$ fixed as an independent real field permits complex variation of the coefficients in the density. The expansion then uses finite-range links. Returning to the original field requires an additional estimate on $S\phi$, beginning with its restriction to a single tile.

Lemma C.4 (Local tests of the transformed background)**.** *Under the assumptions of Lemma C.2, fix a target fine tile P. Let $P^{(m)}$ contain the $O(mb)$-neighbourhood of its tested stencils, clipped at the physical boundary. With the same fine depth and scale weights on these sets,*

$$F_{b,P^*}\big((S-I)\phi\big) \le \sum_{m\ge1}(C\lambda(|z|+|u|))^m F_{b,P^{(m)}}(\phi). \tag{C.20}$$

One coefficient derivative has the corresponding bound with the summand prefactor replaced by $m(C\lambda)^m(|z|+|u|)^{m-1}$.

Proof. The proof of Lemma C.2 bounds CH from the weighted supremum norm to F_b, with operator size $C\lambda(|z|+|u|)$ and range $O(b)$. Apply the weighted supremum bound to the intermediate factors of $(CH)^m$, and the derivative bound to the outermost operator on the output stencil. All intervening field sites are in $P^{(m)}$. Summation of the Neumann series gives the estimate, and differentiating the series marks one of its m factors. □

This bound controls the field on a tile by the field in successive neighbourhoods. A neighbourhood of radius mb can contain order m^4 labels, while the decay in the estimate is exponential in m. To pull the background substitution back into the polymer norm one would have to assign its external field arguments to connected sets of comparable total length, with the same fixed number of tested derivatives.

C.3 Connected Gaussian integration on a fixed kinetic disc

For independent background ψ, the local densities of Lemma C.3 give a connected expansion under the original finite-range Gaussian. Tilting the Gaussian also changes its vacuum normalization. Taking a joint cumulant of the active factors cancels the vacuum contribution on every nonempty coefficient and preserves the fixed order of field derivatives.

Lemma C.5 (Normalized connected expansion for a quadratic tilt)**.** *Write C for an ordinary finite-range covariance and $H = \sum_i H_i$ for localized symmetric quadratic matrices. The label i of each matrix has a fixed-size stencil, and the stencil graph of the fine labels has bounded degree and includes the covariance range. Put*

$$Q_i(\zeta) = \tfrac{1}{2}\langle \zeta, H_i \zeta\rangle, \qquad v_i(\zeta) = e^{-Q_i(\zeta)} - 1. \tag{C.21}$$

Assume that the quadratic bounds are uniform on a fixed complex parameter disc and that on fixed stencils

$$|Q_i(\zeta)| \le C t_0 \sum_{P\sim i} Y_P^2, \qquad \sup_P \mathbb{E}_C e^{aY_P^2/\sigma} \le C_0, \qquad Y_P = F_{b,P^*}(\zeta)/\ell_b. \tag{C.22}$$

Here t_0 is fixed, $a > 0$, and σ is sufficiently small. Let the active factors be a finite family of translated activities or local insertions, each with a connected cover containing its field stencil. Their actual occupied labels are pairwise disjoint. Any label included only for the cover keeps its assigned reference. Use active factors and fillers of the class of Lemma 3.10, with a strict regulator reserve. A product of Gaussian expectations uses independent copies of the original Gaussian, one in each expectation. The multiplicity of polynomial insertions per label in each copy remains bounded.

Then normalized expectation with covariance $\widetilde{C} = C^{1/2}(1 + C^{1/2}HC^{1/2})^{-1}C^{1/2}$ has an exact expansion into connected covers in which a quadratic-tilt insertion costs $C\sigma$ at either radius. A component may use such insertions at unoccupied labels. The sum of all terms involving a tilt insertion gains $C\sigma$, uniformly in volume, after any fixed additional exponential cover weight, provided the active factors also have a sufficiently small weighted anchored sum. Only terms with at least one active factor occur in the normalized connected expansion. The estimate is holomorphic in the coefficients on the indicated disc and has the same first variation bound on smaller fixed discs. All bounds hold through the original fixed order of background field jets.

Proof. Each v_i satisfies

$$|v_i(\zeta)| \le |Q_i(\zeta)| e^{|Q_i(\zeta)|}. \tag{C.23}$$

For a set of distinct tilt insertions in one expectation the local quadratic powers have bounded multiplicity at a given stencil. Finite-range coloring and Hölder's inequality therefore retain one factor $C\sigma$ per insertion, as in Lemma 3.10. More explicitly, color the stencil graph so vectors in any one color are independent. Absorb each bounded power of $Y_P/\sqrt{\sigma}$ in a fixed fraction of its single-stencil exponential moment, and absorb the added quadratic exponential in another fixed fraction for small σ. The constants are exponential in the total number of read labels. Each tilt factor is independent of the background field, so differentiation in a tested background direction does not increase its Gaussian polynomial loss.

For clarity, the reserve estimate works with the multiple copies just described. On the actual occupied fine labels use the occupied decay and quotient estimate of Lemma 3.10; the disjointness of those labels permits each background square to be counted only once. On each unoccupied row use the stable quartic on that row, either at the background or with its assigned fluctuation. A convex combination of real quartic exponents gives the same lower bound with a convex combination of the row displacements. In the recovery estimate, bound the sum of squares of these displacements by the

sum of $C_L(Y_P^{(\alpha)})^2$ over pairs consisting of a read tile P and a Gaussian copy α. This follows from the absolute filter row bound and bounded local stencil size, without requiring one displacement on all rows. Recovery therefore gives the same background inequality and strict reserve. The occupied-field triangle inequalities are unchanged, and the additional fluctuation exponents and polynomial powers are paid by coloring within each copy. Independence across copies multiplies the moment constants. Their product is bounded by a constant exponential in the total number of read labels, counted with multiplicity over the copies.

Fix a nonempty family M of distinct active factors f_A, and introduce scalar variables t_A. Their joint cumulant in the normalized tilted expectation is

$$\kappa_{\widetilde{C}}(f_A : A \in M) = \left[\prod_{A\in M} t_A\right] \log \mathbb{E}_C \left[\prod_i (1 + v_i(\zeta)) \prod_{A\in M} (1 + t_A f_A)\right]. \tag{C.24}$$

For real parameters the finite-dimensional Gaussian identity follows from completing the square, also if C is singular. The identity extends holomorphically within the tilted disc, using domination by the indicated Gaussian moments. Taking a nonempty coefficient in (C.24) subtracts the logarithm of the vacuum denominator automatically. No logarithm of an activity or of a reference factor is used. At fixed real background the logarithm at $t = 0$ can be evaluated as a power series in t; it suffices to bound its squarefree coefficients in the field norm.

Here is the connected support expansion for that coefficient. Call a tilt insertion or an active factor an elementary object, with its chosen connected cover. Expand the two products inside the expectation. Components of the resulting objects are determined by touching covers. Distinct components use independent Gaussian stencils. Consequently the expectation is the hard-core gas of polymers G, where G is a nonempty connected set of elementary objects and its weight is their joint expectation. The hard-core condition is that two covers do not touch. One instance of any elementary object is allowed within a polymer. In the Mayer expansion of its logarithm, a term is a connected family of incompatible polymers. An elementary object can occur in distinct members of this list. Taking the coefficient in (C.24) makes every active factor occur exactly once. Tilt insertions can repeat in different members. These members are products of expectations, so the repeated insertions use different Gaussian copies. Within each expectation the tilt multiplicity is bounded as before.

For completeness one can obtain absolute summability from an anchored estimate alone. Give each elementary object O a nonnegative majorant $r(O)$ containing its two-radius cost and the fixed exponential in read labels. Write $Y(O)$ for its cover. For any required exponential cover weight, leave a fixed extra margin and suppose that

$$\sup_i \sum_{O: i\in Y(O)} r(O) e^{a_1 |Y(O)|} \le \delta_1 \tag{C.25}$$

with a_1 sufficiently large and δ_1 small. The sum of connected elementary configurations through a root label, weighted by the product of the object majorants, is at most $2\delta_1$ after consuming one fixed part of the size margin. Indeed choose a rooted spanning tree of touching objects and charge each child to a site in the closed neighbourhood of its parent cover. A parent of cover size k has at most Dk possible attachment sites. Allow ordered lists of children at each site. If the rooted sum at the previous tree height is bounded by $2\delta_1$, their cost is at most $(1 - 2\delta_1)^{-Dk}$, which is absorbed by the reserved factor e^k for small δ_1.

For the Mayer coefficient of a list of q polymers, the absolute value of the alternating sum over connected subgraphs of its incompatibility graph is at most the number of its spanning trees. To see this, fix an order on the edges and group connected graphs by the greedy spanning tree. For a

given tree the graphs are obtained by including an arbitrary subset of extra edges of larger order than every edge on their replacement tree path. The alternating sum on such an interval is zero or has absolute value one. Root the tree at a polymer containing the prescribed label. The factor $1/q!$ in the Mayer formula then becomes $1/(q-1)!$ after paying for the choice of root. Ordering the children of every rooted tree allows this last factor to pay for the labels of nonroot vertices. The same rooted-tree bound, with the remaining size margin, now applies to the polymers as objects. This proves absolute convergence uniformly in the volume. The formal Mayer coefficients agree with the Taylor coefficients at the vacuum by the usual expansion of the logarithm into connected incompatibility graphs. Equivalently that identity follows by partitioning every graph into its connected components and exponentiating the connected series. Absolute convergence after slightly inflating the object majorants justifies evaluation of all tilt variables at one in (C.24).

To turn these connected coefficients into reblocked weights, use the moment-cumulant identity

$$\mathbb{E}_{\widetilde{C}} \prod_{A\in M} f_A = \sum_{\mathcal{P} \text{ a partition of } M} \prod_{E\in\mathcal{P}} \kappa_{\widetilde{C}}(f_A : A \in E). \tag{C.26}$$

This follows by exponentiating the polynomial logarithm in the same squarefree coefficient algebra. Give each cumulant term its connected union cover and group touching covers. Their products are bounded by one more application of the rooted-tree bound with a further fixed size margin. For each such group include the reference factors at the unused labels in its cover, exactly once. The reference on labels outside the groups stays unchanged. With this assignment, the original multiplicative expansion on the fine labels gives the outgoing circle product. Each cover includes the labels that determine compatibility of its activity or extraction choice. Thus separated groups can be assembled independently in the reverse expansion. The construction uses actual occupied labels for regulator damping and uses covers for connected support. In particular a cover joining two activities across vacant labels never removes their fillers.

All terms at zero tilt are the original finite-range expectations. In the difference there is at least one tilt insertion. Inflate its activity from $C\sigma$ to a fixed small constant when making the anchored estimate and keep one factor $C\sigma$ outside the sum. This proves the difference bound. The same operation can keep any prescribed fixed number of costs belonging to other kinds of objects when those costs are factored out before making the anchored estimate. Uniform majorants on the complex disc justify coefficientwise Cauchy differentiation. Since its radius is independent of σ, the first coefficient variation of the tilt difference still has factor $C\sigma$. Gaussian expectations and products of expectations require only the tested background jets of their factors. This proves the lemma. □

Effect on the mixed activity estimates. One can use this expansion in the chart in which the background argument of a fluctuation factor is the independent variable ψ of the Gaussian square completion. The connected estimate for terms containing an activity or an extraction cover can then be obtained by the size comparison of Lemma 3.9. In the large-radius bound the activity cost is $\tau g^{-9/4}$, and the insertion cost is $O(g^{1/4})$ for each first-order centered ordinary difference. In the small-radius bound these costs are τ and $O(g)$. The output size weight $(A\varrho)^{P(U)}$ is paid by multiplying the object weights by $(A\varrho)^{|\overline{Y(O)}|+c}$, with a fixed c per connecting edge. Indeed a contact across the covariance range needs only a bounded coarse path. A root on such a path has a participating cover within bounded coarse distance. For a long input polymer X, take L large so that $|\overline{Y(X)}| \le C + |X|/4$. The factor $\varrho^{-P(X)}$ then pays a fixed exponential reserve in its fine size and the anchored count of its shapes. For small marks, ordinary decorations and tilt objects the larger weight costs a fixed constant, since their covers have bounded coarse size. The constants may depend on L, ϱ, which are fixed before taking the bare parameters small.

Thus, after retaining one activity or extraction factor, the mixed norm is $O(\tau)$, and the difference from the corresponding un-tilted connected terms with the same active factors and fillers is $O(\sigma\tau)$. The same statements hold in an activity direction by substituting one varied mark or by Cauchy's estimate on the activity ball. For pure terms of perturbative order at least three retain three powers of g in the small-radius estimate and three powers of $g^{1/4}$ at the large radius, where the output factor $g^{9/4}$ gives the same $O(g^3)$ estimate. A second-order local factor is counted with two powers at either radius. In the tilt difference retain in addition one factor of σ, giving $O(\sigma g^3)$. All statements use fixed kinetic discs and involve the same coefficient weights as the un-tilted estimates. In particular, long covariance links by themselves do not force an estimate on unboundedly high derivatives of an activity.

For pure terms of order one and two the one- and two-mark contributions are the dressed mean and covariance on the covers in (C.24). They must be subtracted at those covers and simultaneously included in the retained polynomial reference. The estimates above bound their supported expansions with the corresponding first- and second-order costs. A bound of order g^3 for the full pure remainder additionally requires this subtraction and a uniformly bounded outgoing reference realizing it. The substitution of the actual background $\psi = (I + CH)^{-1}\phi$ also requires a compatible support reassignment, since an incoming factor on one cover is then a function of background values outside that cover.

The long second-order terms in the dressed reference. The connected estimate at fixed kinetic coefficient requires a second-order subtraction on connected supports of unbounded length. This requirement already appears in the Gaussian heat of the stable quartic in an independent background.

Proposition C.6 (Dressed quartic pairs and their supports)**.** *Use the hypotheses of Lemma C.2. In the independent real background ψ, let $V(\psi) = \sum_x \psi_x^4/4$ and use the dressed covariance $\widetilde{C}$. The coefficient of $\psi_x^3\psi_y^3$, for $x \neq y$, in the second-order term of $\log \mathbb{E}_{\widetilde{C}} \exp[-gV(\psi + X)]$ is $g^2\widetilde{C}(x, y)$. For any prescribed exponential size weight $e^{A|U|}$, the part using $\widetilde{C} - C$ admits a connected-cover allocation with anchored sum of normalized sextic coefficients at most $C_A g^2\lambda^2(|z| + |u|)$ at radius ℓ_b, and first derivative at most $C_A g^2\lambda^2$, for small $\lambda(|z| + |u|)$. A fixed finite filter on the two external vertices has the same upper bounds.*

If the bulk C is nonzero and the kinetic perturbation is nonzero, of finite range and of zero row sum, a reference whose logarithm contains only fixed-diameter stencils cannot have remainder $O(g^3)$ for this example uniformly as $g \downarrow 0$. A separately factored quadratic is also permitted in this statement.

Proof. Taylor expansion for real $g \geq 0$ gives the first and second terms $-g\mathbb{E}V(\psi + X)$ and $g^2 \operatorname{Var}(V(\psi + X))/2$. The expansion is valid in zero-field jets through degree six on each fixed finite cube: the third right derivative and the required field jets are bounded on bounded real backgrounds by Gaussian polynomial moments, using $V \geq 0$. In the variance the product of the two linear fluctuation terms at x, y is $\psi_x^3\psi_y^3\widetilde{C}(x, y)$; no term with more fluctuation fields at either vertex has this monomial. The two orientations cancel the factor $1/2$.

For each $m \geq 1$ use the term $C^{(m)}$ of Lemma C.2. For each endpoint pair record the endpoint tiles and a shortest grid path between them, contained in the cube. Include the filter stencils by a fixed enlargement at the two ends, if there is a filter. The path has at most $C(m + 1)$ labels. A tested sextic coefficient has a factor at most $C\ell_b^6 = Cb^{-6}$. For each pair of tiles there are at most Cb^8 site pairs, and the kernel factor is bounded by $C\lambda b^{-2}(C\lambda(|z| + |u|))^m$. Thus the normalized contribution for one tile pair is bounded by $Cg^2\lambda(C\lambda(|z| + |u|))^m$. For a prescribed root tile the

anchored count of endpoint tile pairs is at most $C(m+1)^8$, since both endpoints are within $C(m+1)$ tiles of that root. The size weight contributes at most $e^{CA(m+1)}$. The geometric sum gives the claimed bound. Marking one Neumann factor gives the derivative bound. The finite filter consumes only its bounded absolute row norms and bounded enlargement of the covers.

For the last claim take the unfiltered example in a large periodic cube. The Fourier argument following Lemma C.1 shows that the dressed covariance has entries at arbitrarily large separations in infinite volume. The finite torus coefficients approach these by Fourier Riemann sums (or by periodizing the absolutely summable Neumann series). Fix a torus large enough to have a nonzero entry at sites farther apart than the reference stencils. A sum of logarithms of factors on those stencils has zero $\psi_x^3\psi_y^3$ coefficient, as has any quadratic. On this fixed torus a circle-product error of order g^3 in the mixed norm gives order g^3 in its small-radius jets near zero, also after taking the local logarithm. Here the reference has normalized vacuum bounded away from zero and bounded small-field jets, so the finite-jet inverse is bounded for this fixed torus. This would force $g^2\widetilde{C}(x,y) = O(g^3)$, contradicting the chosen nonzero entry. □

This computation bounds the long sextic coefficients in the size norm in which they would have to be retained. A local reference that includes only a finite collection of momentum jets has no such coefficients. Treating the long second-order terms as part of the reference still requires their connected transport together with the nonlocal background shift.

The bubble under a fixed change of kinetic coefficient. For a fixed kinetic coefficient, the filtered bubble changes by a uniformly small relative amount. The comparison uses the same coefficient at both ends of a scale increment. This qualification matters when the kinetic coefficient itself varies along a trajectory.

Lemma C.7 (Bulk bubble comparison under kinetic dressing)**.** *Let $d = 4$, $q(k) = R^{-2}(1 - \widehat{D}_R(k))$, and W_j be the unfiltered cumulative bulk covariance of the finite-propagation slices, excluding the terminal covariance. For small real z, define*

$$W_j^{[z]}(k) = \frac{\widehat{W}_j(k)}{1 + zq(k)\widehat{W}_j(k)}. \tag{C.27}$$

Let $\mathfrak{b}_j^{[z]}$ be the filtered bubble increment with $\widehat{W}_j, \widehat{W}_{j+1}$ replaced by $W_j^{[z]}, W_{j+1}^{[z]}$ in its Fourier formula. Then, with the undressed bubble $\mathfrak{b}_j$,

$$\mathfrak{b}_j^{[z]} = (1 + O(\lambda|z|))\mathfrak{b}_j, \qquad |\partial_z \mathfrak{b}_j^{[z]}| \le C\lambda\mathfrak{b}_j. \tag{C.28}$$

The constants are uniform in the ordinary and merged scales. This comparison uses the same z in both terms of an increment.

Proof. The local Schur and polynomial bounds of Section 2.4 give for each individual bulk slice at side b

$$0 \le \widehat{C}_b(k) \le C_m\lambda b^2(1 + b^2q(k))^{-m} \qquad (m > 1). \tag{C.29}$$

Indeed in the local split the auxiliary bulk precision has a positive lower bound of order $r + R^2q(k)$ after shifting by κr, where $r = (R/b)^2$. This follows from the Schur complement with its upper block bounded and uniformly positive. The observation has uniformly bounded norm. Applying the polynomial majorant, then multiplying by $a_R = \lambda R^2$, gives the estimate. The constants can include the bounded number of late scales with a larger displacement relative to r. For $q(k) > 0$, split the sum over geometric scales at $b^2q(k) = 1$. The two parts of the sum for $q(k)\widehat{W}_j(k)$ are bounded by

geometric sums with terms $C\lambda b^2 q(k)$ and $C\lambda(b^2 q(k))^{1-m}$. Thus $0 \le q(k)\widehat{W}_j(k) \le C\lambda$ for all j, k, including the merged scales. At $q = 0$ the same bound holds identically.

Write $v = \widehat{W}_j(k)$, $u = \widehat{W}_{j+1}(k)$. Both are nonnegative and $v \le u$. Pointwise,

$$\left(\frac{u}{1+zqu}\right)^2 - \left(\frac{v}{1+zqv}\right)^2 = \int_v^u \frac{2t}{(1+zqt)^3}\,dt. \tag{C.30}$$

The integrand differs from $2t$ by at most $C\lambda|z|\,2t$, and its z-derivative has magnitude at most $C\lambda\,2t$. Multiplying by the nonnegative filter weight $|\widehat{\widehat{T}}(k)|^4$ and integrating proves both estimates. □

Exact first-order quartic packet in the bulk. The change in covariance also alters the filtered square produced by the first heat of the quartic. We identify the entire first-order polynomial before projecting onto its local mass and kinetic parts.

Lemma C.8 (Transport of the centered quartic at fixed kinetic coefficient)**.** *Work in a periodic bulk calculation (with period larger than the original covariance ranges) in dimension four. Let W_j be the cumulative covariance, let J_R have symbol q as in Lemma C.7, and fix a sufficiently small real z. Put*

$$H_* = zJ_R, \qquad T_j = T(I + W_j H_*)^{-1}, \qquad v_j = (T\Sigma_j T^*)(0), \qquad \Sigma_j = W_j(I + H_* W_j)^{-1}, \tag{C.31}$$

where T is the normalized even bulk filter. Fix a horizon M ending before the terminal slice and put $v_* = v_M$. For $\mathcal{H}_4(t;a) = t^4 - 6at^2 + 3a^2$ one has

$$\frac{P_{W_j}\left[e^{-\langle\phi,H_*\phi\rangle/2}\mathcal{H}_4((T\phi)_x; v_*)\right]}{P_{W_j}\left[e^{-\langle\phi,H_*\phi\rangle/2}\right]} = \mathcal{H}_4((T_j\phi)_x; v_* - v_j). \tag{C.32}$$

After the merged step, $0 \le v_* - v_j \le C\lambda b_j^{-2}$; the merged input allows the bound $C\lambda R^{-2}$. Thus the first-order tadpole can be transported with its actual filtered square under a fixed kinetic coefficient, with no additional first-order source if this entire polynomial is retained.

Proof. Apply (C.3) with W_j in place of C. Its use is justified by the bound $\sup_k |zq(k)\widehat{W}_j(k)| \le C\lambda|z|$ from Lemma C.7. Square completion sends the observed mean to $(T_j\phi)_x$ and gives added Gaussian variance v_j. The coefficient of t^4 in $\exp(tu - at^2/2)$, multiplied by $4!$, is $\mathcal{H}_4(u;a)$. Gaussian integration of this generating function changes a to $a - v_j$, proving the identity. For k fixed, the derivative of $w \mapsto w/(1 + zq(k)w)$ is positive and bounded by a uniform constant on the interval of cumulative covariances. The filtered diagonal of the difference is therefore bounded by the sum of the filtered slice diagonals between j and M. The ordinary scale bound sums to $C\lambda b_j^{-2}$. The merged bound follows by summing the individual ultraviolet scales from R. The semigroup identity $P_C P_{W_j} = P_{W_j + C}$ gives the transport assertion. □

The formula identifies the whole first-order reference before localizing any quadratic jet. Its field filter depends on the cumulative covariance and loses finite range. Using this formula in the nonlinear step would require a decomposition of that reference into local factors with controlled activities and coefficient variations.

C.4 Terminal absorption of a marginal quadratic form in dimension four

A terminal form comparable to a small multiple of the Dirichlet kinetic energy need not be small at the quartic radius. It can instead be included in the Gaussian law, provided the scalar mass is centered in the modified law. In this subsection use the normalized Dirichlet form

$$J_R = R^{-2} P_\Lambda L_R P_\Lambda^*, \qquad |f|_{J_R}^2 = \langle f, J_R f \rangle. \tag{C.33}$$

Multiplying J_R by a bounded positive constant gives bulk quadratic Taylor coefficient $|k|^2$. In particular this representative can be used for the local kinetic jet.

To integrate a small quadratic perturbation into the terminal Gaussian law, we need its action on the conditioned covariance, on the principal profile and on the integrated uniform source. The following energy estimates give all three controls in the same Dirichlet form.

Lemma C.9 (Terminal energy and hard covariance)**.** *Take $d = 4$, $1/2 < p < 1$, and $n \geq CR$. For the terminal covariance of Section 2.3, let $C_h = \mathcal{S}_\ell(C_T)$ be the covariance conditioned on $\ell^{\mathsf{T}}\phi = 0$. Then*

$$C_h \leq C\lambda J_R^{-1}, \qquad |f|_{J_R} \leq Cn \|f\|_{F_{n,p}}. \tag{C.34}$$

In particular $|u_T|_{J_R} \leq Cn^{-1}$ and $|C_< \mathbf{1}|_{J_R} \leq C_R n^3$.

Proof. Put $L = P_\Lambda L_R P_\Lambda^*$. The one-dimensional zero-extension walk with uniform positive and negative steps of size at most R has symbol bounded by $C \min\{R^2|k|^2, 1\}$ when considered in one coordinate. The Fourier lower bound for L_R therefore dominates the normalized form of this walk on the cube. For each $\lfloor R/4 \rfloor \leq j \leq \lfloor R/2 \rfloor$ follow each step-j chain in that coordinate to its first exterior point. Telescoping and Cauchy–Schwarz on a chain with $O(n/j)$ sites bound its squared ℓ^2 norm by $C(n/j)^2$ times its squared differences. Averaging over the chosen steps proves $J_R \geq cn^{-2} I$.

The kernel of the split A in Lemma 2.1 is nonnegative, is supported in the spread-out cube, and satisfies $A(x) \leq C D_R(x)$ for $x \neq 0$. If $p_A = \sum_x A(x)$, then $G = p_A I - A_\Lambda$ is the corresponding zero-extension jump form, so $0 \leq G \leq CL$. Since both forms have bounded operator norm,

$$\|L^{1/2} A_\Lambda f\|_2 \leq p_A \|L^{1/2} f\|_2 + C\langle f, Gf \rangle^{1/2} \leq C \|L^{1/2} f\|_2. \tag{C.35}$$

By duality $A_\Lambda L^{-1} A_\Lambda \leq CL^{-1}$. The Schur identity and $B_\Lambda - A_\Lambda = \zeta L$ give

$$V_1 = A_\Lambda + A_\Lambda (B_\Lambda - A_\Lambda)^{-1} A_\Lambda \leq CL^{-1}. \tag{C.36}$$

For positive covariances, conditioning is monotone in the quadratic-form order, by the formula $\langle f, \mathcal{S}_\ell(C) f \rangle = \inf_{t \in \mathbb{R}} \langle f + t\ell, C(f + t\ell) \rangle$. Use $C_T \leq a_R V_c$, $\mathcal{S}_\ell(V_c) \leq CV_1$, and $a_R = \lambda R^2$ to prove the asserted hard bound.

For the energy estimate set $M = \|f\|_{F_{n,p}}$. If x is within distance CR of the exterior and $|y - x|_\infty \leq R$, the difference of the zero-extended values is at most $CM(R/n)^p$. There are at most $Cn^3 R$ points in this collar, and the jump row sums are one. Division by R^2 gives an energy at most $CM^2 n^2 (R/n)^{2p-1}$. Elsewhere, with depth $r \geq CR$, telescoping along nearest-neighbor paths gives a difference at most $CRMn^{-p} r^{p-1}$. At each depth there are at most Cn^3 points, so the energy outside the collar is at most

$$CM^2 n^{3-2p} \sum_{r=1}^{n} r^{2p-2} \leq CM^2 n^2. \tag{C.37}$$

The two stated applications now use the profile and integrated-source bounds of Section 2.3. □

Absorbing a kinetic form changes the variance in the principal direction. A scalar mass shift restores this variance. The same shift also changes the linear source that must be used to recover the initial magnetic moment.

Proposition C.10 (Centering and source restoration with a terminal kinetic form)**.** *For real symmetric Q satisfying $-qJ_R \le Q \le qJ_R$, let $Q_m = Q + mn^{-2}I$. For sufficiently small q, there is a unique $m = m_c(Q)$ in a fixed neighbourhood of zero for which multiplication by $e^{-\langle\phi, Q_m\phi\rangle/2}$ preserves the scalar Gaussian variance σ_T^2 in the decomposition along $a = \ell^{\mathsf{T}}\phi$. Moreover, $|m_c(Q)| \le Cq$, and the resulting hard Gaussian and soft profile satisfy the bounds of Assumption 7.1 uniformly for Q in this form ball. The scalar recentering is real analytic.*

Suppose the exact terminal interaction factors as $F_T = e^{-\langle\phi, Q_{m_c}\phi\rangle/2}\overline{F}_T$, where $\overline{F}_T$ has the interaction and error bounds of Assumption 7.1. Then both the terminal and initial auxiliary magnetic second moments have order $V^{3/2}g^{-1/2}$.

Proof. Write $u = u_T$, $J = J_R$, and $M = Q_m$. Poincaré gives $-q'J \le M \le q'J$ with $q' = C(q + |m|)$. In particular $\|C_h^{1/2} M C_h^{1/2}\| \le C\lambda q'$. Apply (C.3) with covariance C_h to the interaction at each value of a. The hard covariance and profile become

$$C_{h,M} = (I + C_h M)^{-1} C_h, \qquad u_M = (I + C_h M)^{-1} u = u - C_{h,M} M u. \tag{C.38}$$

The added soft precision is $r(M) = \langle u, M(I + C_h M)^{-1} u\rangle$. The tilt determinant is independent of a, and $\ell^{\mathsf{T}} u_M = 1$ because $C_h \ell = 0$. For q' small, $C_{h,M} \le 2C_h$; consequently its Gaussian exponential norm moment follows by adding an independent centered Gaussian to attain covariance $2C_h$ and applying conditional Jensen to $f \mapsto \exp(cn^2\|f\|_{F_{n,p}}^2)$.

For each weighted unit test stencil D entering $F_{n,p}$, the same hard moment gives $\|C_{h,M}^{1/2} D^*\|_2 \le Cn^{-1}$. Using the form norm and Lemma C.9,

$$|D(C_{h,M} M u)| \le \|C_{h,M}^{1/2} D^*\|_2 \|C_{h,M}^{1/2} M u\|_2 \le Cn^{-1}\sqrt{\lambda}\, q' |u|_J \le Cq' n^{-2}. \tag{C.39}$$

Thus $\|u_M - u\|_{F_{n,p}} \le Cq'n^{-2}$, and the magnetic overlap remains of order n^2. These inequalities also give $|u_M|_J \le Cn^{-1}$. In particular $|r(Q)| \le Cqn^{-2}$. Differentiating the resolvent identity, for any symmetric perturbation D_M, gives

$$Dr(M)[D_M] = \langle u_M, D_M u_M\rangle. \tag{C.40}$$

Hence $\partial_m r(Q_m) = n^{-2}\|u_M\|_2^2 \asymp n^{-2}$. The upper bound follows from $\|u_M\|_{F_{n,p}} \le Cn^{-2}$ and the lower bound from $\ell^{\mathsf{T}} u_M = 1$. The intermediate value theorem yields the unique zero and the stated bound on m_c. Analyticity follows from the invertible hard resolvent and the implicit function theorem. In particular, for a form direction D_Q of relative J-bound q_D, implicit differentiation gives

$$Dm_c(Q)[D_Q] = -n^2 \frac{\langle u_M, D_Q u_M\rangle}{\|u_M\|_2^2}, \qquad |Dm_c(Q)[D_Q]| \le Cq_D. \tag{C.41}$$

The profile variation is $-C_{h,M}(D_Q + n^{-2} Dm_c(Q)[D_Q] I)u_M$, whose $F_{n,p}$ norm is at most $Cq_D n^{-2}$ by the same stencil estimate.

At this zero the full tilted law is Gaussian with decomposition $au_M + Y_M$ and the original soft variance. The proof of Proposition 7.4 applied with $\overline{F}_T$ now gives the terminal second moment. To restore the earlier source put $v = C_< \mathbf{1}$. Gaussian translation in the earlier slices gives, including the shifted quadratic factor, the exact generating function up to a positive constant independent of t:

$$e^{t^2(\mathbf{1}^{\mathsf{T}} C_< \mathbf{1} - \langle v, Mv\rangle)/2} \mathbb{E}_M\left[e^{t\langle \mathbf{1} - Mv, \phi\rangle} \overline{F}_T(\phi + tv)\right]. \tag{C.42}$$

Here $\mathbb{E}_M$ is Gaussian expectation in the tilted law. Lemma C.9 gives

$$|\langle v, Mv\rangle| \le C_R q' n^6,\ |\langle Mv, u_M\rangle| \ \le C_R q' n^2,\ \mathrm{Var}(\langle Mv, Y_M\rangle) \le C_R (q')^2 n^6. \tag{C.43}$$

Choose the form ball at fixed R so that the source overlap $|\langle \mathbf{1} - Mv, u_M\rangle|$ remains bounded below by cn^2. The hard part of the source has standard deviation $O_R(n^3)$. In the terminal envelope argument all insertions of this hard linear Gaussian are bounded by Cauchy–Schwarz with the same exponential moment reserve. Thus the terminal second moment of the corrected linear form is $\asymp V^{3/2} g^{-1/2}$. Differentiating twice at zero, the terms with a derivative of $\overline{F}_T$ are controlled exactly by its directional bounds, with $\|v\|_{F_{n,p}}/h \le C_R V^{3/4} g^{1/4}$. As in Proposition 7.4, these terms and the two constant prefactor contributions are $O_R(V^{3/2})$. This proves the initial bound. □

The mass shift is of order q and can exceed the terminal thermal window. Applying this comparison to a physical trajectory would require an outgoing chart that isolates the form and controls $\overline{F}_T$ and its first variation near the recentered mass. The four-dimensional proof uses backward kinetic shooting on the shrinking domain of Section 6.1.

AI Statement

This document is generated with a collaboration with ChatGPT-Astra 6.0.

References

[1] J. Park, *Gaussian and non-Gaussian finite-size scalings of the spread-out Ising model at the free-boundary pseudocritical point*, GitHub preprint, DOI:10.5281/zenodo.24167641 (2026).

[2] M. Aizenman, *Geometric analysis of φ^4 fields and Ising models. Parts I and II*, Comm. Math. Phys. **86** (1982), 1–48. doi:10.1007/BF01205659.

[3] M. Aizenman, *On the number of incipient spanning clusters*, Nucl. Phys. B **485** (1997), 551–582. doi:10.1016/S0550-3213(96)00626-8.

[4] M. Aizenman and H. Duminil-Copin, *Marginal triviality of the scaling limits of critical 4D Ising and ϕ^4_4 models*, Ann. of Math. **194** (2021), 163–235. doi:10.4007/annals.2021.194.1.3.

[5] M. Aizenman and R. Fernández, *Critical exponents for long-range interactions*, Lett. Math. Phys. **16** (1988), 39–49. doi:10.1007/BF00398169.

[6] R. Bauerschmidt, D. C. Brydges, and G. Slade, *Scaling limits and critical behaviour of the 4-dimensional n-component $|\varphi|^4$ spin model*, J. Stat. Phys. **157** (2014), 692–742. doi:10.1007/s10955-014-1060-5.

[7] A. Blanc-Renaudie and A. Nachmias, *Critical percolation on the discrete torus in high dimensions*, arXiv:2512.19672 (2025).

[8] C. Borgs, J. T. Chayes, R. van der Hofstad, G. Slade, and J. Spencer, *Random subgraphs of finite graphs. I. The scaling window under the triangle condition*, Random Structures Algorithms **27** (2005), 137–184. doi:10.1002/rsa.20051.

[9] C. Borgs, J. T. Chayes, R. van der Hofstad, G. Slade, and J. Spencer, *Random subgraphs of finite graphs. II. The lace expansion and the triangle condition*, Ann. Probab. **33** (2005), 1886–1944. doi:10.1214/009117905000000260.

[10] E. Brézin and J. Zinn-Justin, *Finite size effects in phase transitions*, Nucl. Phys. B **257** (1985), 867–893. doi:10.1016/0550-3213(85)90379-7.

[11] D. C. Brydges and T. Spencer, *Self-avoiding walk in 5 or more dimensions*, Comm. Math. Phys. **97** (1985), 125–148. doi:10.1007/BF01206182.

[12] F. Camia, J. Jiang, and C. M. Newman, *The effect of free boundary conditions on the Ising model in high dimensions*, Probab. Theory Related Fields **181** (2021), 311–328. doi:10.1007/s00440-021-01041-9.

[13] S. Chatterjee and J. Hanson, *Restricted percolation critical exponents in high dimensions*, Comm. Pure Appl. Math. **73** (2020), 2370–2429. doi:10.1002/cpa.21938.

[14] F. J. Dyson, *Existence of a phase-transition in a one-dimensional Ising ferromagnet*, Comm. Math. Phys. **12** (1969), 91–107. doi:10.1007/BF01645907.

[15] M. E. Fisher and M. N. Barber, *Scaling theory for finite-size effects in the critical region*, Phys. Rev. Lett. **28** (1972), 1516–1519. doi:10.1103/PhysRevLett.28.1516.

[16] M. E. Fisher, S.-k. Ma, and B. G. Nickel, *Critical exponents for long-range interactions*, Phys. Rev. Lett. **29** (1972), 917–920. doi:10.1103/PhysRevLett.29.917.

[17] E. Flores-Sola, B. Berche, R. Kenna and M. Weigel, *Role of Fourier modes in finite-size scaling above the upper critical dimension*, Phys. Rev. Lett. **116** (2016), 115701.

[18] J. Fröhlich, *On the triviality of $\lambda\phi_d^4$ theories and the approach to the critical point in $d \geq 4$ dimensions*, Nucl. Phys. B **200** (1982), 281–296. doi:10.1016/0550-3213(82)90088-8.

[19] J. Grimm, E. M. Elçi, Z. Zhou, T. M. Garoni, and Y. Deng, *Geometric explanation of anomalous finite-size scaling in high dimensions*, Phys. Rev. Lett. **118** (2017), 115701. doi:10.1103/PhysRevLett.118.115701.

[20] T. Hara, R. van der Hofstad, and G. Slade, *Critical two-point functions and the lace expansion for spread-out high-dimensional percolation and related models*, Ann. Probab. **31** (2003), 349–408. doi:10.1214/aop/1046294314.

[21] T. Hara and G. Slade, *Mean-field critical behaviour for percolation in high dimensions*, Comm. Math. Phys. **128** (1990), 333–391. doi:10.1007/BF02108785.

[22] T. Hara and G. Slade, *Critical behaviour of self-avoiding walk in five or more dimensions*, Bull. Amer. Math. Soc. (N.S.) **25** (1991), 417–423. doi:10.1090/S0273-0979-1991-16085-4.

[23] T. Hutchcroft, E. Michta, and G. Slade, *High-dimensional near-critical percolation and the torus plateau*, Ann. Probab. **51** (2023), 580–625. doi:10.1214/22-AOP1608.

[24] Y. Liu, *A general approach to massive upper bound for two-point function with application to self-avoiding walk torus plateau*, Electron. J. Probab. **30** (2025), article 110. doi:10.1214/25-EJP1371.

[25] Y. Liu, J. Park, and G. Slade, *Universal finite-size scaling in high-dimensional critical phenomena*, Phys. Rev. E **113** (2026), 024143. doi:10.1103/xq7j-zll9.

[26] E. Luijten and H. W. J. Blöte, *Classical critical behavior of spin models with long-range interactions*, Phys. Rev. B **56** (1997), 8945–8958. doi:10.1103/PhysRevB.56.8945.

[27] E. Michta, *The scaling limit of the weakly self-avoiding walk on a high-dimensional torus*, Electron. Commun. Probab. **28** (2023), 1–13. doi:10.1214/23-ECP531.

[28] E. Michta, J. Park, and G. Slade, *Boundary conditions and universal finite-size scaling for the hierarchical $|\varphi|^4$ model in dimensions 4 and higher*, Comm. Pure Appl. Math. **78** (2025), 2001–2118. doi:10.1002/cpa.22256.

[29] R. Panis, *Triviality of the scaling limits of critical Ising and φ^4 models with effective dimension at least four*, Ann. Probab. **54** (2026), 892–972. doi:10.1214/25-AOP1782.

[30] J. Park, *A renormalisation group map for short- and long-ranged weakly coupled $|\varphi|^4$ models in $d \geq 4$ at and above the critical point*, arXiv:2511.03495 (2025).

[31] J. Park, *Torus scaling limits and the plateau of the critical weakly coupled $|\varphi|^4$ model in $d \geq 4$*, arXiv:2511.06321 (2025).

[32] J. Park, *Correlation function of the fractional spread-out Ising models at their upper critical dimension and critical temperature*, GitHub preprint, DOI:10.5281/zenodo.24167641 (2026).

[33] J. Rudnick, G. Gaspari, and V. Privman, *Effect of boundary conditions on the critical behavior of a finite high-dimensional Ising model*, Phys. Rev. B **32** (1985), 7594–7597. doi:10.1103/PhysRevB.32.7594.

[34] G. Slade, *Critical exponents for long-range $O(n)$ models below the upper critical dimension*, Comm. Math. Phys. **358** (2018), 343–436. doi:10.1007/s00220-017-3024-5.

[35] G. Slade, *The near-critical two-point function and the torus plateau for weakly self-avoiding walk in high dimensions*, Math. Phys. Anal. Geom. **26** (2023), article 6. doi:10.1007/s11040-023-09447-8.

[36] M. Wittmann and A. P. Young, *Finite-size scaling above the upper critical dimension*, Phys. Rev. E **90** (2014), 062137. doi:10.1103/PhysRevE.90.062137.