

THE LIMITING LAW OF THE DISCRETE GAUSSIAN LEVEL-LINES

JOSEPH CHEN AND EYAL LUBETZKY

ABSTRACT. Consider the $(2+1)$ D Discrete Gaussian ($\mathbb{Z}\text{GFF}$, integer-valued Gaussian free field) model in an $L \times L$ box above a hard floor. Bricmont, El-Mellouki and Fröhlich (1986) established that, at low enough temperature, this random surface exhibits entropic repulsion: the floor propels the average height to be poly-logarithmic in L . The second author, Martinelli and Sly (2016) showed that, for all but exceptional values of L , the surface has a plateau whose height concentrates on an explicit integer $H(L)$, and fills nearly the full square. It was conjectured there that the boundary of this plateau—the top level-line of the surface—should have random fluctuations of $L^{1/3+o(1)}$.

We confirm this conjecture of [LMS16] and further recover the limiting law of the top level-line: there exists an explicit sequence $N = L^{1-o(1)}$ such that the distance of the top level-line from I , the interval of length $N^{2/3}$ centered along the side boundary, converges, after rescaling it by $N^{1/3}$ and the width of the interval by $N^{2/3}$, to a Ferrari–Spohn diffusion. In particular, the level-line fluctuations at, say, the center of I , have a limit law involving the Airy function rescaled by $N^{1/3}$. This gives the first example of one of the $(2+1)$ D $|\nabla\phi|^p$ models (approximating 3D Ising and crystal formation) where a Ferrari–Spohn limit law of its level-lines is confirmed ($\mathbb{Z}\text{GFF}$ is the case $p = 2$).

More generally, we find the joint limit law of any finite number of top level-lines: rescaling their distances from the side boundary, each by its $(N_n^{2/3}, N_n^{1/3})$, yields a product of Ferrari–Spohn laws. These new results extend to the full universality class of $|\nabla\phi|^p$ models for any fixed $p > 1$.

1. INTRODUCTION

The $(2+1)$ D Discrete Gaussian model, also known as the integer-valued Gaussian free field and denoted here $\mathbb{Z}\text{GFF}$, is a random surface model extensively studied in the context of the roughening transition in crystals (see the work of Chui and Weeks [15] in 1976, and the related models in [6] dating back to the 1950’s). It is dual to the Villain XY model [41], and as such undergoes a Kosterlitz–Thouless phase transition (see [30])—one of two models (along with Solid-On-Solid) where this transition was established by Fröhlich and Spencer in their celebrated works [19, 20].

For $\beta > 0$ (the inverse-temperature, which in our context will be taken fixed and large enough), the $(2+1)$ D $\mathbb{Z}\text{GFF}$ model with a floor (or a hard wall) at height 0 is a probability distribution over functions that assign nonnegative integer heights to the sites of the square grid $\Lambda = \llbracket 1, L \rrbracket^2$. Writing $x \sim y$ for the nearest-neighbor relation in \mathbb{Z}^2 , the probability of $\phi : \Lambda \rightarrow \mathbb{Z}_+$ is given by

$$\pi_\Lambda^0(\phi) \propto \exp \left(-\beta \sum_{x \sim y} |\phi_x - \phi_y|^2 \right), \quad (1.1)$$

with zero boundary conditions ($\phi_x = 0$ for all $x \notin \Lambda$). Define $\widehat{\pi}_\Lambda^0$ as the analogue of π_Λ^0 in the no-floor setting (i.e., when ϕ accepts values in \mathbb{Z} as opposed to \mathbb{Z}_+), and define the infinite-volume measure $\widehat{\pi}_\infty$ as the weak limit of $\widehat{\pi}_\Lambda^0$ as $L \rightarrow \infty$. The aforementioned $\mathbb{Z}\text{GFF}$ phase transition, occurring in $\phi \sim \widehat{\pi}_\infty$ at a critical β_R (empirically, $\beta_R \approx 0.665$), can be demonstrated at ϕ_o as follows: for $\beta \leq \beta_R$ (delocalized regime) the surface is rough, in that $\lim_{L \rightarrow \infty} \text{Var}(\phi_o) = \infty$, while for $\beta > \beta_R$ (localized regime) the surface is rigid, in that $\text{Var}(\phi_o) = O(1)$ (rigidity was shown in [4] at large β , moreover yielding exponential tails of $|\phi_o|$ in that regime; roughness at small β was established in [19, 20]; the continuity of the phase transition was recently established by Lammers [32]).

We focus on $\phi \sim \pi_\Lambda^0$ in the regime of β large enough, where most $x \in \Lambda$ would have $\phi_x = 0$ under the no-floor measure $\widehat{\pi}_\Lambda^0$, yet the floor in π_Λ^0 induces a nontrivial surface (see Fig. 1). This entropic repulsion effect was identified in a pioneering work of Bricmont, El-Mellouki and Fröhlich [5]: despite the penalizing boundary conditions, the surface is propelled to height at least $c\sqrt{\log L}$, where it gains entropy from extra (permitted once at this height) downward fluctuations.

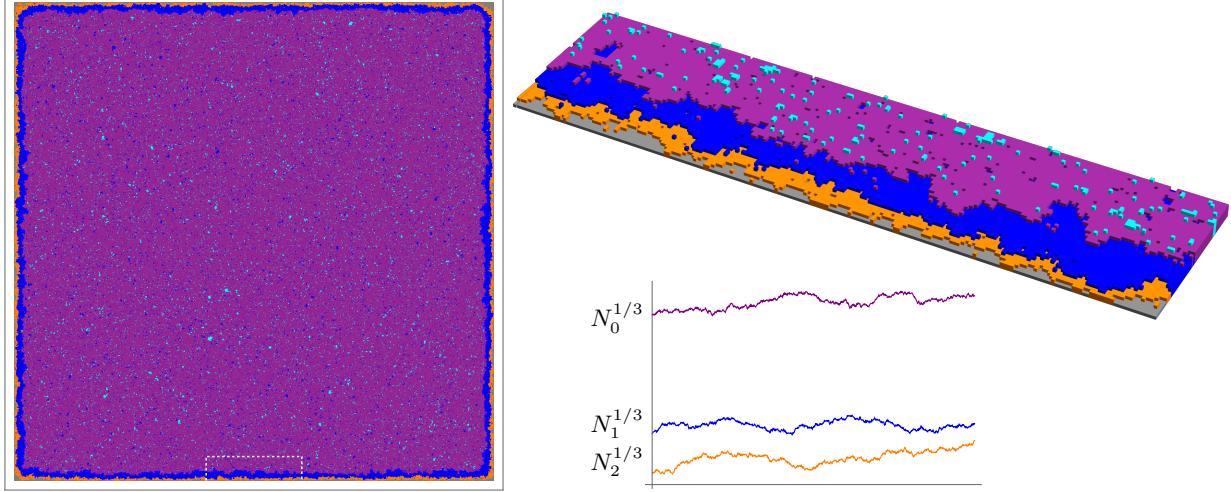


FIGURE 1. Illustration of the low temperature ZGFF on $\llbracket 1, L \rrbracket^2$, where we look at the law of the top level-lines in a rectangle (magnified on right) along the center of the bottom side. On bottom right: independent Ferrari–Spohn diffusions, the limiting law of these level-lines.

The study of entropic repulsion in the paper above considered two closely-related models: the ZGFF and Solid-On-Solid (sos), where the term $|\phi_x - \phi_y|^2$ in Eq. (1.1) is replaced by $|\phi_x - \phi_y|$. Caputo et al. [12–14] made significant progress in characterizing the shape of the sos surface above a floor at low temperature. They showed that the macroscopic level-lines form a sequence of nested loops, with the height of the top one concentrating on one of the two values $\lfloor \frac{1}{4\beta} \log L \rfloor - 1, \lfloor \frac{1}{4\beta} \log L \rfloor$. These works further identified the deterministic scaling limits of the top level-lines in $[0, 1]^2$: these are the boundaries of the shapes formed by dilating the associated Wulff shape to explicit radii, then taking the union of all their translations in $[0, 1]^2$. So, the limits are curved near the corners of the box (where the dilated Wulff shape is visible), and overlap the boundary in intervals surrounding the center-sides (the “flat portions” of the limit), where the random fluctuations are of interest.

The top level-line in sos was shown in [14] to have random fluctuations of at most $L^{1/3+o(1)}$ from the flat portions of its limit. Key to that was showing that the law of the top level-line in the relevant region resembles that of a random walk, conditioned to be nonnegative, and penalized exponentially in the area below it. The latter is known [28] to have $L^{1/3}$ fluctuations, and moreover its scaling limit is known to have the law of a Ferrari–Spohn diffusion [18], defined below (in fact, the continuous analogue—a Brownian excursion tilted by its area—is equivalent via a Girsanov transformation to Brownian motion staying above a parabolic barrier, which was studied in [18]). It was recently established [7] that the order of the sos top level-line fluctuations is at least $L^{1/3}$, and one expects that to be the correct order, with a scaling limit that is *not* Ferrari–Spohn but a variant of it, due to the interaction between the different level-lines (if the lower level-lines were instead deterministic—achievable via suitable boundary conditions—the resulting top level-line would yield a Ferrari–Spohn limit, as was shown in [7] via adapting the recent proof [25] of a Ferrari–Spohn limit for the 2D Ising model with critical prewetting). Rather, the conjectured sos limit should take after the line-ensemble of non-crossing random walks with geometric area tilts (see, e.g., [1, 9–11, 16, 24, 39] for studies of this line-ensemble in the discrete (random walks) and continuous (Brownian polymers) setting, which for a single curve reverts to a Ferrari–Spohn limit).

For ZGFF, the second author, Martinelli and Sly [35] showed that, as in the case of sos, the surface heights concentrate on two values $H, H + 1$, and for “most” values of L the surface forms a plateau of area $(1 - \varepsilon_\beta)L^2$ at height H (for ε arbitrarily small as β increases). In more detail, let

$$H(L) := \max \left\{ h : \widehat{\pi}_\infty(\phi_o = h) \geq \frac{5\beta}{L} \right\}. \quad (1.2)$$

Further define the h level-line(s) of $\phi \sim \pi_\Lambda^0$, for each $h \geq 0$, as the collection of loops obtained from edges dual to nearest-neighbors $x \sim y$ such that $\phi_x < x$ and $\phi_y \geq h$ (for concreteness, in case 4 dual-edges share an endpoint, one “splits” them into two loops along the NORTHEAST diagonal). Call a loop γ *macroscopic* if its length is at least $\log^2 L$ (or one could use a threshold of $(C/\beta) \log L$).

Theorem ([35]). *Fix β large, and consider the \mathbb{Z} GFF $\phi \sim \pi_\Lambda^0$ per Eq. (1.1) and $H(L)$ per Eq. (1.2). Then with high probability, there is a unique macroscopic h level-line loop for each $h = 0, \dots, H$, and there are no macroscopic h level-line loops for any $h \geq H+2$. This sequence of loops is nested; the loops for $h \leq H-1$ have area $(1-o(1))L^2$, and the loop for $h = H$ has area at least $(1-\varepsilon_\beta)L^2$.*

As noted in [35, Remark 1.3], for all but an exceptional set¹ of values of L , the area of the H level-line line loop is in fact $(1-o(1))L^2$, and there are no macroscopic $(H+1)$ level-line loops (making the one at height H the top level-line). An open problem left in that work (see [35, §1.5]) was to prove that this top level-line should have $L^{1/3+o(1)}$ fluctuations away from the corners of Λ .

Our main result identifies the correct scale of the fluctuations of the top level-line from an interval set at the center of the side boundary—which as conjectured, is indeed $L^{1/3+o(1)}$, yet with a necessary $L^{o(1)}$ correction—and moreover recovers its Ferrari–Spohn limit law, defined as follows. Let $\text{Ai}(x)$ be the Airy function (of the first kind), i.e., the solution to $y''(x) = xy$ with the initial condition $y = 0$ at $x = \infty$. For $\sigma > 0$, let $\varphi_\sigma(x) := \text{Ai}((2/\sigma^2)^{1/3}x - \omega_1)$, where ω_1 is the “first” zero of Ai , that is, $\omega_1 = \min\{x > 0 : \text{Ai}(-x) = 0\}$. The stationary Ferrari–Spohn diffusion we consider (see [18, 28]) is the diffusion FS_σ on $(0, \infty)$ with Dirichlet boundary condition at 0 and generator

$$\mathsf{L}_\sigma = \frac{\sigma^2}{2} \frac{d^2}{dx^2} + \sigma^2 \frac{\varphi'_\sigma(x)}{\varphi_\sigma(x)} \frac{d}{dx}. \quad (1.3)$$

This diffusion is ergodic and reversible with respect to the probability density $\frac{(2/\sigma^2)^{-\frac{1}{3}}}{\text{Ai}'(-\omega_1)^2} \varphi_\sigma(x)^2 \mathbb{1}_{\{x > 0\}}$, where the rescaled Airy function φ_σ given above is also the first eigenfunction of the operator L_σ . Our main result addresses side-lengths L excluding an explicit set \mathcal{B} of integers of logarithmic density zero (see Remark 1.3), about which the plateau transitions from one height to the next.

Theorem 1.1. *Fix $\beta > 0$ large enough and consider $\phi \sim \pi_\Lambda^0$, the (2+1)D \mathbb{Z} GFF model on $\Lambda = \llbracket 1, L \rrbracket^2$ for $L \in \mathbb{Z} \setminus \mathcal{B}$ (with \mathcal{B} as in Remark 1.3) with a floor and zero boundary conditions. Set $H(L)$ per Eq. (1.2), and consider the bottom boundary interval $I = \llbracket \frac{L}{2} - N^{2/3}, \frac{L}{2} + N^{2/3} \rrbracket$ for*

$$N := \frac{1}{\widehat{\pi}_\infty(\phi_o = H)} \quad (= L^{1-o(1)}). \quad (1.4)$$

Let \mathfrak{L}_0 be the top macroscopic level-line of ϕ , and define its vertical distance from I via

$$\rho(x) = \min\{y \geq 0 : (\frac{L}{2} + x, y) \in \mathfrak{L}_0\} \quad \text{for} \quad -N^{2/3} \leq x \leq N^{2/3}.$$

Then $Y_0(t) := N^{-1/3} \rho(tN^{2/3})$ converges weakly to the stationary law FS_σ on $[-1, 1]$ as $L \rightarrow \infty$ for a fixed $\sigma > 0$. The analogous $\bar{Y}_0(t)$ w.r.t. $\bar{\rho}(x) = \max\{y \leq \frac{L}{2} : (\frac{L}{2} + x, y) \in \mathfrak{L}_0\}$ has $\|\bar{Y}_0 - Y_0\|_\infty \xrightarrow{\text{P}} 0$.

More generally, for fixed m , let $\mathfrak{L}_0 \subset \dots \subset \mathfrak{L}_{m-1}$ be the top m macroscopic level-lines, and

$$N_n := \frac{1}{\widehat{\pi}_\infty(\phi_o = H - n)} \quad (= L^{1-o(1)}) \quad (n = 0, \dots, m-1). \quad (1.5)$$

Let $I_n = \llbracket \frac{L}{2} - N_n^{2/3}, \frac{L}{2} + N_n^{2/3} \rrbracket$, and let $\rho_n(x) = \min\{y \geq 0 : (\frac{L}{2} + x, y) \in \mathfrak{L}_n\}$ denote the vertical distance of \mathfrak{L}_n from I_n . The joint law of $Y_n(t) := N_n^{-1/3} \rho_n(tN_n^{2/3})$ ($n = 0, \dots, m-1$) converges weakly to that of m independent stationary Ferrari–Spohn diffusions FS_{σ_n} on $[-1, 1]$ for fixed $\sigma_n > 0$. The analogous $\bar{Y}_n(t)$ w.r.t. $\bar{\rho}_n(x) = \max\{y \leq \frac{L}{2} : (\frac{L}{2} + x, y) \in \mathfrak{L}_n\}$ have $\max_{n \leq m} \|\bar{Y}_n - Y_n\|_\infty \xrightarrow{\text{P}} 0$.

¹Precisely, for all side-lengths L outside a set $\mathcal{B} \subset \mathbb{Z}$ of zero logarithmic density ($\sum_{k \in \mathcal{B} \cap \llbracket 1, n \rrbracket} \frac{1}{k} = o(\log n)$).

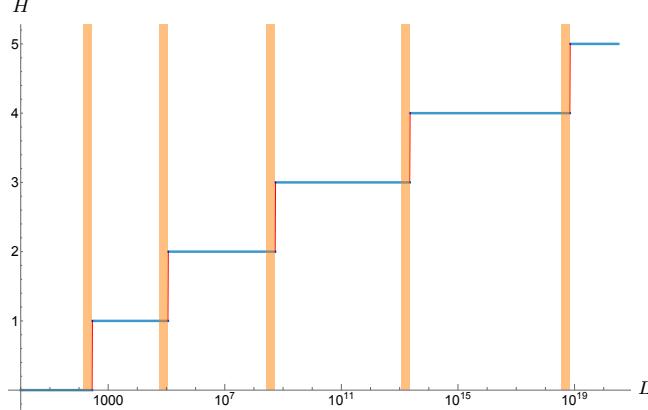


FIGURE 2. The exceptional set \mathcal{B} of values of L as per Remark 1.3, highlighted in orange, which delimits the intervals $[\lfloor L_h, \frac{3}{4}L_{h+1} \rfloor]$ where the top level-line concentrates on height h . As depicted in this log-plot, the set \mathcal{B} has zero logarithmic density.

Remark 1.2. The sequence $N = N_0$ from Eq. (1.4) satisfies $e^{-c\sqrt{\beta \log L / \log \log L}} \leq N/L \leq 1/(5\beta)$, and each of these two bounds gives the behavior of the fluctuations of \mathfrak{L}_0 for infinitely many values of L . Namely, the fluctuations of \mathfrak{L}_0 are of order $L^{1/3}$ for infinitely many values of L , yet they are $o(L^{-1/3})$ for infinitely many other values of L . As for $N_1 \geq N_2 \geq \dots \geq N_n$ from Eq. (1.5), while these scales are all $L^{1-o(1)}$ for any fixed n , they are of lower order: $N_n = N_{n-1} e^{-\Theta(\sqrt{\beta \log L / \log \log L})}$ (in particular, the fluctuations of $\mathfrak{L}_1, \dots, \mathfrak{L}_n$ have order $N_n^{1/3} = L^{1/3-o(1)} = o(L^{1/3})$).

Remark 1.3. As one increases the side-length L of the box, once the top level-line concentrates on height h (at some L_h), it will remain so until L further increases by a factor of $\exp(c\beta \frac{h}{\log h})$. Precisely, the transition marking the onset of level $h+1$ occurs at $L \approx (4 + \varepsilon_\beta)\beta/\widehat{\pi}_\infty(\phi_o = h+1)$. As such, one can define the set \mathcal{B} of exceptional values of the side-lengths L (see Fig. 2) via

$$\mathcal{B} = \bigcup_{h \geq 1} [\lfloor \frac{3}{4}L_h, L_h \rfloor] \quad \text{where} \quad L_h = \lceil 5\beta/\widehat{\pi}_\infty(\phi_o = h) \rceil \quad (h = 1, 2, \dots). \quad (1.6)$$

With this choice, $\sum_{k \in \mathcal{B} \cap [1, n]} \frac{1}{k} = O(\sqrt{\log n \log \log n}) = o(\log n)$ (i.e., zero logarithmic density), and the $\mathbb{Z}\text{GFF}$ surface will concentrate on height $H(L) = h$ for $L \in [\lfloor L_h, \dots, \frac{3}{4}L_{h+1} \rfloor]$ ($h = 1, 2, \dots$).

Remark 1.4. Taking the intervals $I_n = [\frac{L}{2} - KN_n^{2/3}, \frac{L}{2} + KN_n^{2/3}]$ for any fixed $K > 0$ (in lieu of $K = 1$), one has convergence to independent stationary Ferrari–Spohn diffusions FS_{σ_n} on $[-K, K]$.

The new results extend to the family of $|\nabla \phi|^p$ models, defined as follows. In lieu of Eq. (1.1), the probability of $\phi : \Lambda \rightarrow \mathbb{Z}_+$ (with 0 boundary conditions and inverse-temperature $\beta > 0$) is given by

$$\pi_\Lambda^{(p),0}(\phi) \propto \exp \left(-\beta \sum_{x \sim y} |\phi_x - \phi_y|^p \right), \quad (1.7)$$

so that $p = 2$ is the $\mathbb{Z}\text{GFF}$ and $p = 1$ is the SOS model. It was shown in [35] that, for all $p > 1$, and all “typical” values of L , the surface is typically a plateau (with microscopic fluctuations) at a single deterministic height H (that scales differently with L for different values of p), as is the case for SOS and $\mathbb{Z}\text{GFF}$. As for the limit law of the top level-lines, the pictures in SOS and $\mathbb{Z}\text{GFF}$ are different (a Ferrari–Spohn limit in the latter, vs. a conjectured variant thereof in the former).

In what follows, put $\pi_\Lambda^{(p)}$ instead of $\pi_\Lambda^{(p),0}$ for brevity (0 boundary conditions by default), and let

$$H^{(p)}(L) := \max \left\{ h : \widehat{\pi}_\infty^{(p)}(\phi_o = h) \geq \frac{5\beta}{L} \right\}, \quad N_n^{(p)} := \frac{1}{\widehat{\pi}_\infty^{(p)}(\phi_o = H^{(p)} - n)} \quad (= L^{1-o(1)}), \quad (1.8)$$

the analogues of H and N_n defined in Eqs. (1.2) and (1.5) for the $\mathbb{Z}\text{GFF}$.

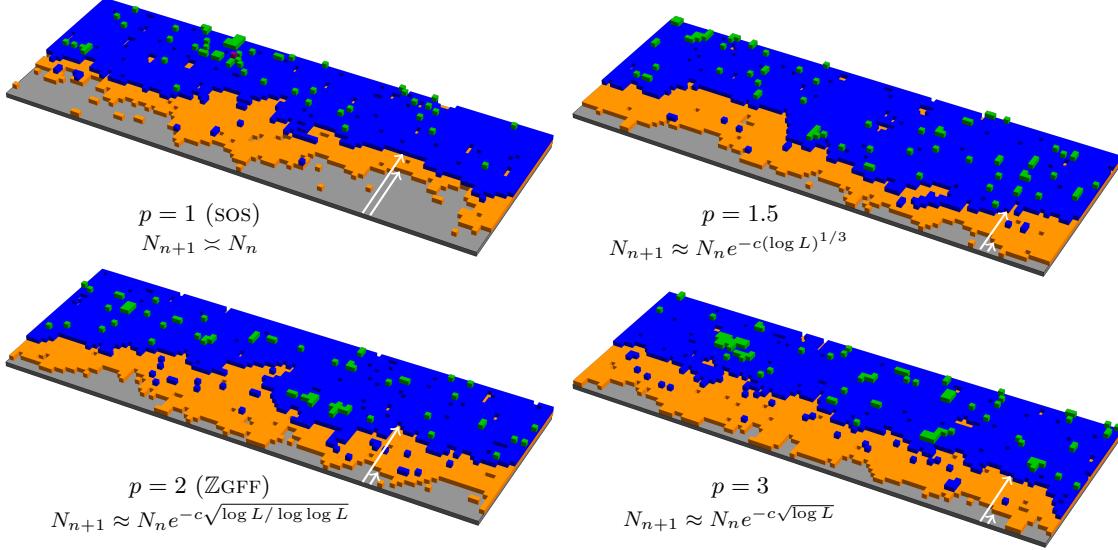


FIGURE 3. Comparison of the scales $N_{n+1}^{(p)} \asymp 1/\hat{\pi}_\infty^{(p)}(\phi_o = L)$ in the $|\nabla\phi|^p$ -model for different values of p . As the scaling for the level-line \mathfrak{L}_n is $((N_n^{(p)})^{2/3}, (N_n^{(p)})^{1/3})$, the $|\nabla\phi|^p$ models enjoy a scale separation between the level-lines for $p > 1$, unlike the sos model ($p = 1$).

The next theorem shows that for all $p > 1$, the level-lines have the same limit law as for $\mathbb{Z}\text{GFF}$.

Theorem 1.5. *Fix $p > 1$, and consider the setting of Theorem 1.1 with the $|\nabla\phi|^p$ model $\pi_\Lambda^{(p)}$ replacing π_Λ^0 and $H^{(p)}, N_n^{(p)}$ replacing H, N_n . Then the conclusion of Theorem 1.1 holds: the law of $(Y_n)_{n=0}^{m-1}$, the rescaled vertical distances of the top m level-lines from the bottom boundary intervals, converges weakly to the law of m independent stationary Ferrari–Spohn diffusions FS_σ on $[-1, 1]$.*

Remark 1.6. Analogously to Remark 1.2, for $2 < p < \infty$, the sequence $N_0^{(p)}$ from Eq. (1.8) satisfies $e^{-c\sqrt{\beta \log L}} \leq N_0^{(p)}/L \leq 1/(5\beta)$ and $N_n^{(p)} = N_{n-1}^{(p)} e^{-\Theta(\sqrt{\beta \log L})}$. For $1 < p < 2$, one has that $N_0^{(p)}$ from Eq. (1.8) satisfies that $e^{-c\beta^{1/p}(\log L)^{(p-1)/p}} \leq N_0^{(p)}/L \leq 1/(5\beta)$ and $N_n^{(p)} = N_{n-1}^{(p)} e^{-\Theta(\beta^{1/p}(\log L)^{(p-1)/p})}$. In both cases, as in Remark 1.2, the upper and lower bounds on $N_0^{(p)}/L$ give the correct behavior of the fluctuations along infinitely many values of L : the fluctuations of \mathfrak{L}_0 have order $L^{1/3}$ in the former case and $L^{1/3-o(1)} = o(L^{1/3})$ in the latter. The lower level-lines, $\mathfrak{L}_1, \dots, \mathfrak{L}_n$ for any fixed n , all have fluctuations $L^{1/3-o(1)} = o(L^{1/3})$.

Remark 1.7. Analogously to Remark 1.3, one defines the exceptional set \mathcal{B} for the $|\nabla\phi|^p$ model in Theorem 1.5 exactly as in Eq. (1.6) but with $\hat{\pi}_\infty^{(p)}$ replacing $\hat{\pi}_\infty$ there. The fact that \mathcal{B} has zero logarithmic density extends to all $1 < p < \infty$, owing to the fact that the large deviation rate of $\hat{\pi}_\infty^{(p)}$ is super-linear: $\lim_{h \rightarrow \infty} -\frac{1}{h} \log \hat{\pi}_\infty^{(p)}(\phi_o = h) = \infty$ (see Remark 7.2 for more on this).

1.1. Proof ideas. Our starting point is the observation that the level-lines in the $\mathbb{Z}\text{GFF}$, and more generally, in the $|\nabla\phi|^p$ model for any $p > 1$, ought to be “separated” from one another. Indeed, \mathfrak{L}_n , the n -th level-line from the top, is expected to behave like a random walk tilted by a term of $\exp[\lambda_{L,n}\mathcal{A}(\mathfrak{L}_n)]$, where $\lambda_{L,n} = \hat{\pi}_\infty^{(p)}(\phi_o = H^{(p)} - n)$ and $\mathcal{A}(\mathfrak{L}_n)$ is the area of its interior. Consequently, as such area-tilted random walks are known to have a Ferrari–Spohn limit law after rescaling their width and height by $(\lambda^{2/3}, \lambda^{1/3})$ respectively, one expects the distance of \mathfrak{L}_n from (say) the bottom boundary to be of order $\hat{\pi}_\infty^{(p)}(\phi_o = H^{(p)} - n)^{-1/3}$, i.e., $(N_n^{(p)})^{-1/3}$ for $N_n^{(p)}$ from Eq. (1.8). In the case of sos ($p = 1$), the fact that $\hat{\pi}_\infty^{(1)}(\phi_o = h) \asymp e^{-4\beta h}$ would have every level-line \mathfrak{L}_n , for $n \geq 0$ fixed, be at the same scale: all would be found at $\Theta(L^{1/3})$ from the bottom boundary, as seen in the top left of Fig. 3 (see Section 1.3 for more on this).

For all $p > 1$, however, one has $\widehat{\pi}_\infty^{(p)}(\phi_o = h+1)/\widehat{\pi}_\infty^{(p)}(\phi_o = h) = o(1)$, whence the above heuristic implies that the scales $N_n^{(p)}$ satisfy $N_{n+1}^{(p)} = o(N_n^{(p)})$ for all n ; i.e., each level-line is at a microscopic scale compared to the one above it (see Fig. 3). Let us focus on the $\mathbb{Z}\text{GFF}$ model for simplicity. There, the above ratio of $\widehat{\pi}_\infty(\cdot)$ terms is $e^{-\Theta(h/\log h)}$. For $h \asymp H(L)$, this is $\exp[-(\log L)^{1/2-o(1)}]$, e.g., we expect the second level-line \mathfrak{L}_1 to be at an $\exp[-(\log L)^{1/2-o(1)}]$ fraction of the “height” (distance from the bottom boundary) of the top level-line \mathfrak{L}_0 , and similarly for the height of \mathfrak{L}_{n+1} vs. \mathfrak{L}_n . Intuitively, this should mean there is no interaction between the level-lines; if these are effectively independent area-tilted RWs, their limit would be a product of Ferrari–Spohn diffusions.

Showing this separation of the level-lines rigorously is nontrivial since the typical height of \mathfrak{L}_{n+1} ought to be $N_n^{1/3}e^{-(\log L)^{1/2-o(1)}}$, so one needs to first control \mathfrak{L}_n and show that it is concentrated about its height of $N_n^{1/3}$ to a finer degree. However, even for the SOS level-lines ($p = 1$), where much more was known on $\{\mathfrak{L}_n\}$, the best-known estimates on the height of \mathfrak{L}_n had an $L^{o(1)}$ error... Moreover, we emphasize that, even after one manages to show that the level-lines are well separated, the interactions between could still have a huge impact on their limiting laws. A notoriously challenging obstacle in these and related models is the *pinning* problem, whereby a level-line might opt to stick to its predecessor/successor (in this case, random; in other settings, it can be the boundary of the box) due to said interactions. Handling these delicate interactions (and associated pinning issues) is where most of our efforts in the proof are concentrated.

Our crucial idea here is to separate the analysis into two cases: one showing that \mathfrak{L}_n stochastically dominates a Ferrari–Spohn diffusion, and another showing it is stochastically dominated by one with the same parameters. If one were to try to show directly that \mathfrak{L}_n behaves as an area-tilted random walk, current techniques break down due to the combined interactions (and pinning hazards) with random level-lines below and above. For instance, the depinning result of [27] is quite delicate, even failing when the exponential decay of interactions has rate $\leq \beta/2$ (for us it is β , a non-issue). Specifically, it is only valid when the boundary of the domain is flat, and it appears that extending it to “nearly flat” domains is highly nontrivial, if at all possible using their method. In our setting, even the top level-line \mathfrak{L}_0 faces pinning issues from below (to \mathfrak{L}_1) and from above (to microscopic holes along the boundary of the mesoscopic rectangle of interest, where we wish to derive a Ferrari–Spohn law). Luckily, the stochastic domination route allows us to employ FKG adjustments that only go in the right direction for the respective side of the bound (upper/lower) currently studied. This does away with the pinning issues on one of the two sides of \mathfrak{L}_n (the interactions on the remaining side are handled differently in the upper and lower bounds, as we explain below).

Upper bound on the limit law of \mathfrak{L}_n . The upper bound is proved in three steps.

- (i) We show that each \mathfrak{L}_n does not exceed height $N_n^{1/3}(\log L)^C$, a sufficiently fine degree of accuracy to obtain the aforementioned separation. This is achieved via a “growth gadget,” that is, a result showing that conditional on the fact that the level-line loop \mathfrak{L}_n exceeds a certain area, it is likely to exceed it farther.
- (ii) Then, disregarding lower level-lines via the separation established above, as well as upper level-lines via monotonicity, we show that \mathfrak{L}_n continues to drop down from $N_n^{1/3}(\log L)^C$ to its equilibrium height $N_n^{1/3}$ using finer random walk estimates.
- (iii) Lastly, we show that the law of the level-line intersected in a small box is approximately that of an area-tilted RW, when its initial height is $O(N_n^{1/3})$.

(As a byproduct of Item (i), we refine the the corresponding estimate for SOS; see Remark 4.6.)

In each case, one aims to estimate the law of a level-line \mathfrak{L}_n in a mesoscopic box V with boundary conditions $H-n, H-n-1$, and show it is approximately governed by $\exp[-\beta|\mathfrak{L}_n| + \lambda_{L,n}\mathcal{A}(\mathfrak{L}_n)]$ where $\lambda_{L,n} = \widehat{\pi}_\infty(\phi_o = H-n)$. E.g., the prequel [14] on SOS featured a growth gadget analogous to Item (i) where the region V had area $L^{4/3+o(1)}$ (then applied to boxes of size $L^{2/3+o(1)} \times L^{2/3+o(1)}$). The paper [7] on SOS had an analysis analogous to Item (iii) in a box of size $(CL^{2/3}) \times (CL^{2/3})$.

Unfortunately, for the $\mathbb{Z}\text{GFF}$, we can afford to obtain such estimates at regions of area $L(\log L)^c$. The culprit is the two point probability $\widehat{\pi}_\infty(\phi_x \geq h, \phi_y \geq h)$ for a pair of nearest-neighbors $x \sim y$. In sos, this probability is $\approx e^{-6\beta h}$, so when $\widehat{\pi}_\infty(\phi_x \geq h) \approx e^{-4\beta h}$ is $1/L$, the two point probability is $L^{-3/2}$. This translates to a domain restriction of $|V| \ll L^{3/2}$, a computation made rigorous in [14]. In the $\mathbb{Z}\text{GFF}$ model, the two point probability turns out to be $L^{-1-o(1)}$, since the large deviation event $\phi_x \geq h$ is dominated by ϕ which climbs to h in the shape of a harmonic pinnacle [35], whereby its neighbors are already at height $h - O(h/\log h)$. The consequence is that we can only work on domains of size $|V| \leq L^{1+o(1)}$, which forces us to take boxes of say $L^{2/3}(\log L)^c \times L^{1/3}(\log L)^c$.

This is a serious restriction, due to the pinning issues. When the height of V is only $L^{1/3+o(1)}$, even though the probability of a random walk reaching height $L^{1/3+o(1)}$ on an $L^{2/3}$ interval is only $e^{-L^{o(1)}}$, the interactions between \mathfrak{L}_n and ∂V could potentially tilt the measure by a factor of $e^{\varepsilon_\beta L^{2/3}}$. As we mentioned, the case of a flat boundary was handled in [27]; however, our V is a perturbation of a rectangle with wiggly boundaries, for which that proof method breaks down.

In the direction of showing an upper bound on \mathfrak{L}_n , we circumvent these pinning issues by using FKG to forget the floor constraint on a subset of V , picking up an area tilt only on a set $F \subset V$ instead. This effectively removes the size restriction from V , allowing us to use $L^{2/3+o(1)} \times L^{2/3+o(1)}$ rectangles, where the top boundary is too far to induce pinning issues. We handle interactions with the bottom boundary of V by conditioning \mathfrak{L}_n to stay sufficiently far away from it (again by FKG).

As evident from this strategy, we moreover have that with high probability, if we condition on $\mathfrak{L}_{n+1}, \mathfrak{L}_{n+2}, \dots$, then the conditional limit law of \mathfrak{L}_n is stochastically dominated by a Ferrari–Spohn diffusion. One can view this as an induction revealing level-lines from exterior to interior, ultimately showing that the joint limit law of $\mathfrak{L}_0, \dots, \mathfrak{L}_{m-1}$ is dominated by independent Ferrari–Spohn laws.

Lower bound on the limit law of \mathfrak{L}_n . For the lower bound on the top level-line \mathfrak{L}_0 , we first remove the lower level-lines by FKG (intuitively, they only push \mathfrak{L}_0 upwards). Then, we study the restriction of \mathfrak{L}_0 to a domain V which is a perturbation of an $N_0^{2/3} \times N_0^{1/3}(\log L)$ rectangle, with a flat boundary at the bottom and wiggly boundary on the other three sides. All the monotonicity workarounds from the upper bound now no longer apply, being in the wrong direction. In particular, the top boundary of V cannot be placed far away anymore. Instead, we condition \mathfrak{L}_0 to stay away from the top boundary, and handle the (now flat) bottom boundary by extending the depinning results of [27] to our setting. Once we establish the depinning, we can reduce the problem to the area-tilted 2D random walk setting of prior works [7, 25] which show convergence to Ferrari–Spohn. This proves a lower bound for \mathfrak{L}_0 , which in particular shows that its typical height is $O(N_0^{1/3})$.

Turning to the next level-line, \mathfrak{L}_1 , we look at its restriction to a domain V which is a perturbation of a $N_1^{2/3} \times N_1^{1/3} \log L$ rectangle. Since $N_1^{1/3} \log L = o(N_0^{1/3})$, the behavior of \mathfrak{L}_0 can be isolated from the analysis of \mathfrak{L}_1 , and as before the lower level-lines $\mathfrak{L}_2, \mathfrak{L}_3, \dots$ can be removed by FKG. This then allows us to show the desired lower bound for \mathfrak{L}_1 , and we can proceed inductively for \mathfrak{L}_n .

As in the upper bound on \mathfrak{L}_n (but with a reverse order of which level-lines are revealed, where here the induction reveals them from interior to exterior), one obtains that with high probability, conditional on $\mathfrak{L}_0, \dots, \mathfrak{L}_{n-1}$, the limit law of \mathfrak{L}_n stochastically dominates a Ferrari–Spohn diffusion. Consequently, the joint limit law of $\mathfrak{L}_0, \dots, \mathfrak{L}_{m-1}$ dominates independent Ferrari–Spohn laws.

Remark 1.8. Note that the upper bound had to reveal the level-lines from exterior to interior (revealing \mathfrak{L}_{n-1} would introduce a non-flat interacting boundary at distance $L^{1/3+o(1)}$ above \mathfrak{L}_n) whereas the lower bound had to reveal the level-lines from interior to exterior (revealing \mathfrak{L}_{n+1} would introduce a non-flat interacting boundary at distance $L^{1/3+o(1)}$ below \mathfrak{L}_n).

1.2. Proof outline and organization.

The paper is organized as follows.

In Section 2, we establish a polymer representation for the law of a level-line \mathfrak{L} using classical cluster expansion techniques. This requires some refinements of large deviation results of [35]. It turns out that the natural object to study here is not \mathfrak{L} , but the entire component of bonds γ dual to

height disagreement edges containing \mathfrak{L} , which we call a disagreement polymer. Thus, the polymer model of interest no longer falls under the family of “Ising polymers” for which much has been proven (e.g., the existence of a surface tension [17] and the depinning from a flat boundary [27]). Instead, the geometry of γ is allowed to be more complex than that of a contour, and its law has an additional term capturing the area tilt of finite components surrounded by γ .

In Section 3, we prove several foundational results for the aforementioned class of disagreement polymers, which are already known for “Ising polymers.” We begin by defining animals, cone-points, and irreducible components in order to establish a product structure to the polymer law, via Ornstein–Zernike analysis. Next, we extend results of [17, Sec. 4] to our model, proving the existence and key properties of the surface tension, as well as elementary depinning results. We then define the Wulff shape corresponding to the surface tension, and continue the Ornstein–Zernike theory to show the proper renormalization needed to turn the product structure into a random walk.

In Section 4, we show one side (growth of a droplet) of the macroscopic scaling limit of the $(H - n)$ level-line \mathfrak{L}_n , which should be given by a translation of Wulff shapes. (Hereafter, we prove results on \mathfrak{L}_n , but remember that the object we work with is γ .) The behavior in the $\mathbb{Z}\text{GFF}$ is markedly different from that in the SOS model: here the Wulff shape is only visible on the $O(L)$ scale for select values of L , and only for the top level-line \mathfrak{L}_0 (see Remark 4.5). A direct implication of this scaling limit is that \mathfrak{L}_n has a maximum deviation of $N_n^{1/3}(\log L)^C$ from the boundary of Λ_L , away from the corners. Following the strategy employed for the SOS model in [14], we look at the region \mathcal{D} known to be contained inside \mathfrak{L}_n and iteratively grow this region. The key is to show that in a rectangle V with $H - n, H - n - 1$ boundary conditions, \mathfrak{L}_n drops to a lower height. Then we can place such rectangles all around the boundary of the contained region \mathcal{D} to expand \mathcal{D} in every direction. A similar dropping lemma is needed in Section 5, but the difference is that here we need to handle boundary conditions at an angle. The lattice effect makes this a nontrivial change (being the reason there is even a Wulff shape to begin with), leading to a different analysis.

In Section 5, once we show that the height of \mathfrak{L}_n above the bottom side of $\partial\Lambda_L$ is at most $N_n^{1/3}(\log L)^C$, we prove that inside an $N_n^{2/3}(\log L)^C \times 2N_n^{2/3}(\log L)^C$ rectangle, the curve of \mathfrak{L}_n is stochastically dominated by an area tilted random walk in the middle $N_n^{2/3}$ interval. As mentioned above, this requires another dropping lemma, this time showing the stronger result that \mathfrak{L}_n drops from height $N_n^{1/3}(\log L)^C$ to $O(N_n^{1/3})$. (In turn, we must assume that the endpoints of \mathfrak{L}_n are at an angle of ≈ 0 .) Then, combining monotonicity considerations with previous results on area tilted random walks, we prove the sought upper bound on the limit law in Theorem 1.1.

In Section 6, we prove that inside a $TN_n^{2/3} \times N_n^{1/3}(\log L)^C$ rectangle, the curve \mathfrak{L}_n stochastically dominates an area tilted random walk in the middle $N_n^{2/3}$ interval. Since this random walk has the same increment law as the one studied in Section 5, it converges to the same Ferrari–Spohn diffusion, proving the lower bound on the limit law in Theorem 1.1.

In Section 7, we prove Theorem 1.5, extending our results to the $|\nabla\phi|^p$ model for $1 < p < \infty$. To that end, we prove a lower bound on the ratio $\widehat{\pi}_\infty^{(p)}(\phi_o = h)/\widehat{\pi}_\infty^{(p)}(\phi_o = h - 1)$, which was not addressed as part of the large deviation results from [35] when $p \neq 2$. This was not needed in that work, which only concerned the height histogram of those model, but is imperative here to obtain an asymptotically sharp bound on the law of the disagreement polymer.

1.3. Open problems and related works. In this work, we showed that the scaling limit of the m top level-lines, $\mathfrak{L}_0, \dots, \mathfrak{L}_{m-1}$, in the $\mathbb{Z}\text{GFF}$ above a floor at large enough β , is m independent Ferrari–Spohn diffusions. Moreover, this holds for the $|\nabla\phi|^p$ model for any fixed $p > 1$, which captures the complete universality class when varying p , as the behavior for $p = 1$ (SOS) is expected to be different. It remains an open problem to establish the limit law for the countably many level-lines $\mathfrak{L}_0, \mathfrak{L}_1, \dots$ in SOS, conjectured to be a Brownian line ensemble with a geometric area tilt; see [1, 9–11, 16, 24, 39] for a plethora of works studying this object.

For $p > 1$, several notable problems remain open. Theorems 1.1 and 1.5 address the law of the level-lines away from the corners. It is expected that at the corners, the level-lines should have fluctuations of $L^{1/2-o(1)}$ (as opposed to $L^{1/3-o(1)}$), with an associated scaling limit of independent Brownian motions. (This is open for all $p \geq 1$, i.e., including SOS.) In addition, our results exclude an exceptional set \mathcal{B} of values of L (whereas the SOS analysis of level-lines in [14] only needed to avoid a critical point of transition between top level-lines). It would be interesting to investigate the critical and near-critical behavior of the model for values of L marking the onset of a new level-line.

Lastly, we mention recent developments on the rough phase of the $|\nabla\phi|^p$ model (with no floor). At high enough temperature with uniform boundary conditions, convergence of the $\mathbb{Z}\text{GFF}$ to the Gaussian Free Field (GFF) was shown in [2,3]; this is expected, but remains open for all other $p \geq 1$. When the boundary conditions are tilted, e.g., taking the value $[\theta_1 x + \theta_2 y]$ at every boundary point (x, y) for some slope $\theta = (\theta_1, \theta_2) \neq 0$, then, as long as at least one of the θ_i 's is irrational, the surface is known [40] to delocalize at some unknown rate (“qualitative delocalization”) for all β and $p \geq 1$. At high enough temperature, this was extended also to rational slopes in [33] for SOS. More recently, logarithmic delocalization at high enough temperature was established for all $0 < p \leq 2$ in [37]. It is further believed that the rough phase induced by the slope should have a GFF scaling limit, mirroring the high temperature regime with flat boundary conditions. However, this remains a formidable open problem at large β , where even logarithmic fluctuations are not yet established. One recent result in this direction is for SOS tilted by an added long range potential [34], which indeed has a GFF scaling limit (for any slope with $\theta_1, \theta_2 > 0$ including when both are rational; this is in contrast to $|\nabla\phi|^p$ for all $p > 1$, including $\mathbb{Z}\text{GFF}$, which are rigid at large β when $\theta_1, \theta_2 \in \mathbb{Z}$).

2. LAW OF DISAGREEMENT POLYMERS

The goal of this section is to formulate the cluster expansion framework for the level-lines of the $\mathbb{Z}\text{GFF}$ configuration ϕ . In what follows, we will typically refer to dual-edges (in $(\mathbb{Z}^2)^*$) as bonds, to help distinguish them from the edges of \mathbb{Z}^2 . We say $u, v \in \mathbb{Z}^2$ are $*$ -adjacent if their $L^\infty(\mathbb{R}^2)$ distance is 1 (i.e., they are either adjacent in \mathbb{Z}^2 , or their bounding cells share a single corner). For any point $u \in \mathbb{R}^2$, we will denote the x and y coordinates by u_1, u_2 respectively. Further let ∂U be the boundary bonds of U , i.e., the set of bonds dual to uv with $u \in U$ and $v \notin U$, and let $\partial_v U$ be the external vertex boundary of U , i.e., every vertex $v \notin U$ adjacent to some $u \in U$.

Definition 2.1 (Disagreement polymer). Let ϕ be a \mathbb{Z} -valued height function on the vertices of a connected domain $V \subset \mathbb{Z}^2$. Associate to each bond $e \in (\mathbb{Z}^2)^*$, dual to some edge $(x, y) \in \mathbb{Z}^2$, the gradient $(\nabla\phi)_e := \phi_x - \phi_y$, where x is taken to be the NORTH vertex if (x, y) is vertical and the WEST vertex if (x, y) is horizontal. We say the bond e is a *disagreement bond* of ϕ if $(\nabla\phi)_e \neq 0$ (including when ϕ_x or ϕ_y are specified by the boundary conditions). A *disagreement polymer* γ is a (maximal) connected component of the disagreement bonds of ϕ , and we let \mathcal{P}_ϕ denote the disagreement polymers in ϕ . For such a polymer, let D_i be the connected components of $V \setminus \gamma$, noting that by the maximality of γ , within each D_i , all the vertices that are $*$ -adjacent to $V \setminus D_i$ must have the same height h_i in ϕ . Call the triple $(\gamma, \{D_i\}, \{h_i\})$ a *labeled disagreement polymer*. For brevity, we denote labeled disagreement polymers by γ (omitting the $\{h_i\}$ from the notation).

NB. In the graph on $(\mathbb{Z}^2)^*$ whose edges are γ , no vertex can have degree 1 (degrees 0, 2, 3, 4 are possible), hence the restriction that all $v \in D_i$ that are $*$ -adjacent to $V \setminus D_i$ have the same height. See the example in Fig. 4 illustrating this.

Definition 2.2 (Energy, length and decorations of a disagreement polymer). Let $(\gamma, \{D_i\}, \{h_i\})$ be a labeled disagreement polymer. Its *length* $\mathcal{N}(\gamma)$ and *energy* $\mathcal{E}_\beta(\gamma)$ are defined as

$$\mathcal{N}(\gamma) = \sum_{e \in \gamma} |(\nabla\phi)_e| \quad , \quad \mathcal{E}_\beta(\gamma) := \beta \sum_{e \in \gamma} |(\nabla\phi)_e|^2 .$$

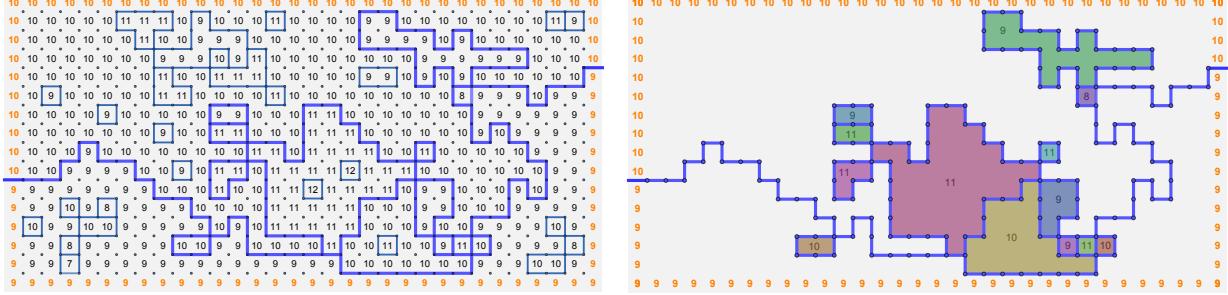


FIGURE 4. Left: A height function ϕ on a rectangle with $h, h-1$ boundary conditions at $h=10$, with all disagreement bonds in blue. The disagreement polymer γ is highlighted. Right: The corresponding labeled disagreement polymer $(\gamma, \{D_i\}, \{h_i\})$. Each region D_i is assigned a single height h_i which can be read off from ϕ on the left (its interior boundary).

To describe the law of the disagreement polymer, we will use a family of decoration functions Φ on subsets $W \subset \mathbb{Z}^2$ —each referred to as a cluster if it is also connected—satisfying the following properties (see the more general family of decoration functions in Definition 3.1).

- (i) If W is not connected, then $\Phi(W) = 0$.
- (ii) The function Φ is translation invariant, as well as invariant under a rotation by $\pi/2$ or reflection with respect to one of the axes.
- (iii) There exists a constant $C > 0$ such that for every W , we have $|\Phi(W)| \leq \exp(-(\beta - C)\mathbf{d}(W))$, where $\mathbf{d}(W)$ is the size of the smallest connected set of bonds in $(\mathbb{Z}^2)^*$ containing all the boundary bonds of W .

The main result of this section is the following, where here and in what follows, we use the notation $D_i^\circ = D_i \setminus \Delta_\gamma$ to denote $D_i \setminus \Delta_\gamma$ where

$$\Delta_\gamma := \{u \in V : \text{dist}(u, \gamma) \leq 1/\sqrt{2}\},$$

as the value of ϕ on $D_i \cap \Delta_\gamma$ (and elsewhere on $\partial_v D_i^\circ$) are specified to be the corresponding h_i .

Proposition 2.3 (Cluster expansion with a floor). *Fix $n \geq 0$. Let $V \subset \mathbb{Z}^2$ be a connected domain, and consider the \mathbb{Z} GFF model $\pi_{V;F}^\xi$ with a floor at 0 imposed only on a subset $F \subset V$, and boundary conditions ξ that are $H-n$ on a $*$ -connected path in $\partial_v V$ and $H-n-1$ elsewhere so that they induce a unique disagreement polymer $(\gamma, \{D_i\}, \{h_i\})$ in $V \cup \partial_v V$ that contains boundary disagreements². Then for $\beta \geq \beta_0$, the law of this unique disagreement polymer γ is given by*

$$\pi_{V;F}^\xi(\gamma) = \frac{1}{Z_{V;F}^\xi} \exp \left(-\mathcal{E}_\beta(\gamma) + \sum_{\substack{W \subset V \\ W \cap \Delta_\gamma \neq \emptyset}} \Phi(W) \right) \prod_{i \geq 0} \widehat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F), \quad (2.1)$$

for $Z_{V;F}^\xi = Z_{V;F}^\xi(\beta, n)$ and a decoration function $\Phi(W)$ as per Definition 2.2.

Moreover, if we further have $|F| \leq L(\log L)^\kappa$ and $|\partial F| \leq L^{1-\delta}$ for fixed $\kappa > 0$ and $0 < \delta < \frac{1}{3}$, denote by D_0 and D_1 the regions of γ containing the boundary vertices of V at heights $H-n$ and $H-n-1$, respectively, and let

$$E := \left\{ |\gamma| \leq L^{1-\delta}, |F| - |D_0 \cap F| - |D_1 \cap F| \leq L^{1-\delta} \right\},$$

²We say the disagreement polymer contains boundary disagreements if it contains a bond dual to uv for $u, v \notin V$.

then the following holds for all $\beta \geq \beta_0$ (not depending on κ, δ). The probability distribution given by

$$p_{V;F}^\xi(\gamma) := \frac{1}{\tilde{Z}_{V;F}^\xi} \exp \left(-\mathcal{E}_\beta(\gamma) + \frac{|D_0 \cap F|}{N_n} + \sum_{\substack{W \subset V \\ W \cap \Delta_\gamma \neq \emptyset}} \Phi(W) \right) \prod_{i \geq 2} \hat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F)$$

for $\gamma \in E$, with N_n from Eq. (1.5) and $\tilde{Z}_{V;F}^\xi$ a normalizer, satisfies that for every $\gamma \in E$,

$$\pi_{V;F}^\xi(\gamma \mid E) = (1 + o(1))p_{V;F}^\xi(\gamma). \quad (2.2)$$

We prove this result in Section 2.2, after we establish some properties of the measure $\hat{\pi}_\infty$.

2.1. Large deviations without a floor. The following was established in [35, Eq. (3.1)–(3.4)]) for the infinite volume \mathbb{Z} GFF measure $\hat{\pi}_\infty$: for every fixed large enough β and every $h \geq 2$,

$$e^{-c_0 \beta h / \log h} \leq \frac{\hat{\pi}_\infty(\phi_o = h)}{\hat{\pi}_\infty(\phi_o = h-1)} \leq e^{-c_1 \beta h / \log h}, \quad (2.3)$$

$$\hat{\pi}_\infty(\phi_o = h) = \exp \left[-2\pi\beta \frac{h^2}{\log h} + O\left(\frac{h^2}{\log^2 h}\right) \right], \quad (2.4)$$

$$\hat{\pi}_\infty(\phi_z = h \mid \phi_o = h) \leq e^{-c_2 \beta h / \log h}, \quad (2.5)$$

where $c_0, c_1, c_2 > 0$ are absolute constants.

Remark 2.4. A stronger upper bound of $\exp(-c_2 h^2 / \log^2 h)$ was stated in [35, Eq. (3.3)], even though the bound in Eq. (2.5) was the one later proved in [35, §3.3]. The weaker bound still suffices for the proof of the main theorem there: one needs only to modify the hypothesis in [35, Prop. 4.4] (where this result is used) from $h \geq \log \log L$ to $h \geq (\log \log L)^2$, which would reproduce the same estimates needed in that proof from the weaker Eq. (2.5). The proof applied [35, Prop. 4.4] in a box Λ_ℓ with $h = (\log \ell)^{1/2+o(1)}$, and hence the stronger hypothesis $h \geq (\log \log \ell)^2$ would be valid. NB. In [35], the aim was to show $\hat{\pi}_V^0(\min_{x \in V} \phi_x \geq -h) \leq \exp[-(1+o(1))I(h)]$ for the correct $I(h)$, while our analysis of the level-lines requires a much more precise bound of the form $(1+o(1))\exp[-I(h)]$.

For our proofs, it will be crucial to have the stronger upper bound of $\exp(-ch^2 / \log^2 h)$ on the probability appearing in the left hand of Eq. (2.5); we adapt the argument of [35] to obtain it, as well as estimates in an arbitrary region V containing a ball of a certain radius around the origin o .

Theorem 2.5 (adapting [35, Thm. 3.1]). *There exist constants $\beta_0 > 0$ and $c > c' > 0$ so that the following holds for every $\beta \geq \beta_0$ and integer $h \geq 2$. Let $V \subset \mathbb{Z}^2$ be a region containing $\mathcal{B}_r(o)$, the ball of radius $r = \lceil 2ch / \log h \rceil$ centered at the origin o , as well as $\mathcal{B}_{r+1}(z)$ for a site $z \in V$. Then*

$$\exp \left(-c\beta \frac{h}{\log h} \right) \leq \frac{\hat{\pi}_V^0(\phi_o = h)}{\hat{\pi}_V^0(\phi_o = h-1)} \leq \exp \left(-c'\beta \frac{h}{\log h} \right), \quad (2.6)$$

$$\exp \left(-2\pi\beta \frac{h^2}{\log h} - c\beta \frac{h^2}{\log^2 h} \right) \leq \hat{\pi}_V^0(\phi_o = h) \leq \exp \left(-2\pi\beta \frac{h^2}{\log h} + c\beta \frac{h^2}{\log^2 h} \right), \quad (2.7)$$

$$\hat{\pi}_V^0(\phi_z = h \mid \phi_o = h) \leq \exp \left(-c\beta \frac{h^2}{\log^2 h} \right). \quad (2.8)$$

Proof. The arguments of Eqs. (2.3) and (2.4) extend more or less verbatim to the setting given here of a more general domain V in Eqs. (2.6) and (2.7), provided that $V \supset \mathcal{B}_r(o)$ for the given r , and we begin by explaining this point. Let $R = \lfloor h / \log h \rfloor$.

The lower bound given in [35] on Eq. (2.3) was stated for $\hat{\pi}_\infty(\phi_o = h) / \hat{\pi}_\infty(\phi_o = h-1)$, carried out on $V = \mathcal{B}_L(o)$ for $L \gg R$, but in fact holds for $\hat{\pi}_V$ for any domain $V \supset \mathcal{B}_R(o)$ ³.

³The proof of this bound in [35] is concluded immediately after [35, Claim 3.5], using nothing only that $V \supset \mathcal{B}_R(o)$ until that point; it does appeal to [35, Claim 3.6], but the latter is already phrased for a general domain V .

Remark 2.6. To see why the requirement $V \supset \mathcal{B}_R(o)$ is the only one needed for the lower bound on Eq. (2.6), we briefly summarize that argument. The proof first reduced the lower bound to showing that all x neighbors of o would have $\phi_x \geq h - cR$, for a large enough c , as expected given that the (harmonic) optimum ϕ^* for the \mathbb{R} -valued Dirichlet problem has $h - \phi_x^* \asymp h/\log h$ at those sites: this step is valid in any domain V . Next, factoring out the contribution of ϕ^* , the problem was reduced, sequentially, to

- (i) the configuration $\sigma := \phi - \phi^*$ on $V \setminus \{o\}$ (already here the assumption $\mathcal{B}_R(o) \subset V$ is used, as one defines ϕ^* as the \mathbb{R} -valued solution on $\mathcal{B}_R(o)$, and 0 outside);
- (ii) then to its integer part $\bar{\sigma} := \lfloor \sigma \rfloor$ in $V \setminus \{o\}$;
- (iii) and finally to $\bar{\sigma}$ in a \mathbb{Z} GFF model on $V \setminus \{o\}$ where the interaction strength in the interior $\mathcal{B}_R(o)$ is modified to 1/2 vs. 1 in its exterior.

(This second series of reductions incurs a cost that is absorbed by further increasing c .) The dominant term in the lower bound is $\exp(-c\beta R)$, yielding the sought bound that appears here, for the more general V , in the left-hand of Eq. (2.6).

The upper bound of [35] on Eq. (2.3) holds for every $V \supset \mathcal{B}_R(o)$ as follows: arguing very similarly to the proof there⁴, if one chooses $r = \delta R$ for a small enough $\delta > 0$ (fixed independently of β), then

$$\frac{\widehat{\pi}_V(\phi_o = h)}{\widehat{\pi}_V(\phi_o = h-1)} \leq e^{-\frac{3}{4}\beta r} + e^{\varepsilon_\beta r} \frac{\widehat{\pi}_{\mathcal{B}_r(o)}(\phi_o = h)}{\widehat{\pi}_V(\phi_o = h-1)} \leq e^{-\frac{3}{4}\beta\delta R} + e^{-c_0\beta\log(\frac{1}{\delta})R^2} \frac{\widehat{\pi}_V(\phi_o = h)}{\widehat{\pi}_V(\phi_o = h-1)},$$

using that $\widehat{\pi}_{\mathcal{B}_r(o)}(\phi_o = h) \leq e^{-c_0\beta\log(\frac{1}{\delta})R^2} \widehat{\pi}_{\mathcal{B}_R(o)}(\phi_o = h)$ and that $\widehat{\pi}_{\mathcal{B}_R(o)}(\phi_o = h) \leq e^{\varepsilon_\beta R} \widehat{\pi}_V(\phi_o = h)$ using [35, Lem. 3.6] for the first inequality and [35, Cor. 3.9] for the second one⁵. Rearranging the above equation yields the sought upper bound on $\widehat{\pi}_V(\phi_o = h)/\widehat{\pi}_V(\phi_o = h-1)$ in Eq. (2.6).

It is for Eq. (2.4) where one needs to extend the radius of the ball contained in V from R to cR ; more precisely, one has the following more general form of [35, Cor. 3.9], via the exact same proof⁶:

Claim 2.7 (extension of [35, Cor. 3.9] via the same proof). *Let $c_0 > 0$ be an absolute constant satisfying, for all $h \geq 1$ and $V \supset \mathcal{B}_R(o)$, that $\widehat{\pi}_V(\phi_o = h)/\widehat{\pi}_V(\phi_o = h-1) \geq \exp(-c_0\beta h/\log h)$. Setting $r = \lceil 2c_0h/\log h \rceil$, for every $V \supset \mathcal{B}_r(o)$ one has*

$$e^{-\varepsilon_\beta r} \leq \frac{\widehat{\pi}_V(\phi_o = h)}{\widehat{\pi}_{\mathcal{B}_r(o)}(\phi_o = h)} \leq e^{\varepsilon_\beta r}.$$

The required bound in Eq. (2.7) now follows from [35, Lem. 3.10], which showed that

$$\left| \log \widehat{\pi}_{\mathcal{B}_r(o)}(\phi_o = h) - 2\widehat{\pi}\beta \frac{h^2}{\log r} \right| \leq c\beta \frac{h^2}{\log^2 r} + c\beta r^2,$$

where $c > 0$ is some absolute constant⁷.

It remains to establish Eq. (2.8), via a small modification of the argument of [35]. Letting

$$X = \max_{x \sim z} \phi_x \quad , \quad Y := \min_{x \sim z} \phi_x,$$

the proof of Eq. (2.5) in [35] considered the events $\{X \leq h\}$ and $\{Y \geq h - \delta\sqrt{h/\log h}\}$ to derive the sought estimate. Here we will instead consider $\{X \leq \delta_1 h/\log h\}$ and $\{Y \geq h - \delta_2 h/\log h\}$. Specifically, let $c' > 0$ be the constant in the upper bound form Eq. (2.6), and let

$$E_1 := \{X \leq h + \Delta\} \quad , \quad E_2 := \{Y \geq h - 5\Delta\} \quad \text{for} \quad \Delta := \left\lfloor (c'/200) \frac{h}{\log h} \right\rfloor.$$

⁴See the argument that appears immediately after [35, Lem. 3.10].

⁵Here one appeals only to the first part of [35, Cor. 3.9], which is valid for any $V \supset \mathcal{B}_R(o)$.

⁶The upper bound in that result used the fact $r \geq 2c_0R$ to replace $\widehat{\pi}_V(\phi_o = h-1)$ in [35, Eq. (3.12)] by $\widehat{\pi}_V(\phi_o = h)$; the lower bound holds for all r .

⁷As seen in the short proof of [35, Lem. 3.10], the constant c' associated with the error term $e^{c'r^2}$ in its statement is of the form $c' = c\beta$ for some absolute constant $c > 0$.

By iterating the upper bound in Eq. (2.6), using here that $\mathcal{B}_r(x) \subset V$ for each $x \sim z$, one has that

$$\begin{aligned} \widehat{\pi}_V^0(E_1^c \mid \phi_o = h) &\leq 4 \max_{x \sim z} \frac{\widehat{\pi}_V^0(\phi_x > h + \Delta)}{\widehat{\pi}_V^0(\phi_o = h)} \leq 4 \max_{x \sim z} \frac{\widehat{\pi}_V^0(\phi_x = h)}{\widehat{\pi}_V^0(\phi_o = h)} \exp\left(-c'\beta\Delta \frac{h}{\log h}\right) \\ &\leq \exp\left(-(c' - o(1))\beta \frac{h\Delta}{\log h}\right), \end{aligned} \quad (2.9)$$

with the second line following from Claim 2.7 to show that $\widehat{\pi}_V(\phi_x = h)/\widehat{\pi}_V(\phi_o = h) \leq \exp(\varepsilon_\beta r)$, which is negligible as $r \asymp \frac{h}{\log h} = o(\Delta \frac{h}{\log h})$. It thus suffices to show that, for some $c > 0$,

$$\widehat{\pi}_V^0(\phi_z = h \mid \phi_o = h, E_1) \leq \exp\left(-c\beta \frac{h^2}{\log^2 h}\right). \quad (2.10)$$

Next, on the events $E_1 \cap E_2^c$, we can reveal ϕ at the neighbors $\{x_i\}_{i=1}^4$ of z and then increase those values by monotonicity, to find that

$$\begin{aligned} \widehat{\pi}_V^0(\phi_z = h \mid \phi_o = h, E_1, E_2^c) &\leq \widehat{\pi}_V^0(\phi_z \geq h \mid \phi_{x_1} = h - 5\Delta, \phi_{x_i} = h + \Delta \text{ for } 2 \leq i \leq 4) \\ &\leq (1 + \varepsilon_\beta)e^{-\beta(3\Delta^2 + (5\Delta)^2)} = (1 + \varepsilon_\beta)e^{-28\beta\Delta^2}. \end{aligned}$$

Thus, bounding $\widehat{\pi}_V^0(\phi_z = h \mid \phi_o = h, E_1) \leq \widehat{\pi}_V^0(\phi_z = h \mid \phi_o = h, E_1, E_2^c) + \widehat{\pi}_V^0(E_2 \mid \phi_o = h, E_1)$, in order to show Eq. (2.10), it is enough to show that for some $c > 0$,

$$\widehat{\pi}_V^0(E_2 \mid \phi_o = h, E_1) \leq \exp\left(-c\beta \frac{h^2}{\log^2 h}\right). \quad (2.11)$$

Again by monotonicity,

$$\begin{aligned} \widehat{\pi}_V^0(\phi_z \geq h + \Delta \mid \phi_o = h, E_1, E_2) &\geq \widehat{\pi}_V^0(\phi_z \geq h + \Delta \mid \phi_{x_i} = h - 5\Delta \text{ for } 1 \leq i \leq 4) \\ &\geq (1 - \varepsilon_\beta)e^{-4\beta(6\Delta)^2} = (1 - \varepsilon_\beta)e^{-144\beta\Delta^2}. \end{aligned} \quad (2.12)$$

On the other hand,

$$\begin{aligned} \widehat{\pi}_V^0(\phi_z \geq h + \Delta \mid \phi_o = h, E_1, E_2) &\leq \frac{\widehat{\pi}_V^0(\phi_z \geq h + \Delta \mid \phi_o = h, E_1)}{\widehat{\pi}_V^0(E_2 \mid \phi_o = h, E_1)} \\ &\leq \frac{\widehat{\pi}_V^0(\phi_z \geq h + \Delta \mid \phi_o = h)}{\widehat{\pi}_V^0(E_2 \mid \phi_o = h, E_1)\widehat{\pi}_V^0(E_1 \mid \phi_o = h)} \leq \frac{\exp(-(c' - o(1))\beta \frac{h\Delta}{\log h})}{\widehat{\pi}_V^0(E_2 \mid \phi_o = h, E_1)}, \end{aligned}$$

where the last transition used Eq. (2.9) to show that $\widehat{\pi}_V^0(E_1 \mid \phi_o = h) = 1 - o(1)$, whereas the numerator was bounded from above first by $\widehat{\pi}_V^0(\phi_z \geq h + \Delta)/\widehat{\pi}_V^0(\phi_o = h)$ and then by iterating the upper bound in Eq. (2.6) and thereafter using Claim 2.7, as was done in Eq. (2.9). Combining this with Eq. (2.12), we find that

$$\widehat{\pi}_V^0(E_2 \mid \phi_o = h, E_1) \leq (1 + \varepsilon_\beta) \exp\left(\beta\left(144\Delta - (c' - o(1))\frac{h}{\log h}\right)\Delta\right) \leq \exp\left(-(c'/4)\beta \frac{h\Delta}{\log h}\right)$$

by the choice of Δ . This establishes Eq. (2.11) and completes the proof. \blacksquare

The following lemma shows that the large deviation estimate in Eq. (2.7) is also an upper bound for all points in V , in particular at sites x near the boundary of V (whereas Theorem 2.5 was only applicable for sites x at distance at least $ch/\log h$ from ∂V).

Lemma 2.8. *There exists $c_0 > 0$ such that, for β large and every $V \subset \mathbb{Z}^2$, $h \geq 2$, and $x \in V$,*

$$\widehat{\pi}_V^0(\phi_x = h) \leq \exp\left(-2\pi\beta \frac{h^2}{\log h} + c_0 \frac{h^2}{\log^2 h}\right). \quad (2.13)$$

Proof. Let $r = \lceil 2ch/\log h \rceil$ for the constant $c > 0$ from Theorem 2.5, and let $V' = V \cup \mathcal{B}_r(x)$. Then

$$\widehat{\pi}_{V'}^0(\phi_x \geq h) \geq \widehat{\pi}_{V'}^0(\phi_x \geq h, \phi|_{\partial V \setminus \partial V'} \geq 0) \geq \left(\inf_{\substack{\xi: \xi|_{\partial V'} = 0, \\ \xi|_{\partial V \setminus \partial V'} \geq 0}} \widehat{\pi}_V^\xi(\phi_x \geq h) \right) \widehat{\pi}_{V'}^0(\phi|_{\partial V \setminus \partial V'} \geq 0).$$

By FKG, this is minimized at $\xi \equiv 0$; namely,

$$\widehat{\pi}_{V'}^0(\phi_x \geq h) \geq \widehat{\pi}_V^0(\phi_x \geq h) \widehat{\pi}_{V'}^0(\phi|_{\partial V \setminus \partial V'} \geq 0) \geq \widehat{\pi}_V^0(\phi_x \geq h) \prod_{x \in \partial V \setminus \partial V'} \widehat{\pi}_{V'}^0(\phi_x \geq 0), \quad (2.14)$$

with the last inequality again following from FKG. A standard Peierls argument (see, e.g., [4]) shows that for any $h \geq 0$ and any V ,

$$\widehat{\pi}_V^0(\phi_x > h) \leq \varepsilon_\beta \widehat{\pi}_V^0(\phi_x = h),$$

where $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow \infty$ (and similarly, $\widehat{\pi}_V^0(\phi_x < -h) \leq \varepsilon_\beta \widehat{\pi}_V^0(\phi_x = -h)$ by symmetry). In particular, conditional on $\phi \equiv 0$ on a subset of $\partial V \setminus \partial V'$, the probability that $\phi_x = 0$ for an additional vertex x along that boundary is at least $\exp(-\varepsilon_\beta)$, and so

$$\prod_{x \in \partial V \setminus \partial V'} \widehat{\pi}_{V'}^0(\phi_x \geq 0) \geq e^{-\varepsilon_\beta |\partial V \setminus \partial V'|} \geq e^{-\varepsilon_\beta h^2/\log^2 h}, \quad (2.15)$$

where the last inequality holds (for a different ε_β) using that, since $\partial V \setminus \partial V' \subset \mathcal{B}_r(x)$, we can infer that $|\partial V \setminus \partial V'| \leq \pi r^2 \leq O(h^2/\log^2 h)$. Combining Eqs. (2.14) and (2.15), we find that

$$\widehat{\pi}_V^0(\phi_x = h) \leq e^{\varepsilon_\beta h^2/\log^2 h} \widehat{\pi}_{V'}^0(\phi_x = h), \quad (2.16)$$

whence the proof concludes via Theorem 2.5. \blacksquare

We can now bound the probability of the floor event. Our bound will be asymptotically tight, whereas the analogous estimate in [35, Prop. 4.4] (which considered a similar event: rather than $B = \bigcap_{x \in V} \{\phi_x \geq -h\}$ addressed here, it pertained B intersected with another event forbidding macroscopic nonzero paths) had a $(1 + o(1))$ term, not as a prefactor, but within the exponent in that probability. It is crucial that we have this more precise estimate for our cluster expansion expression in Section 2.2, and it comes at a cost of area and boundary constraints on the domain. As we describe below (see Remark 2.11), these constraints are more stringent in \mathbb{Z} GFF compared to the case of sos, where the corresponding estimate was applicable to an $L^{2/3+\varepsilon} \times L^{2/3+\varepsilon}$ box (whereas we can only afford to address a $L^{2/3}(\log L)^c \times L^{1/3}(\log L)^c$ box in the setting of \mathbb{Z} GFF).

Lemma 2.9. *Fix $0 < \delta < \frac{1}{3}$, $\kappa > 0$, and $n \geq 0$. Let $V \subset \mathbb{Z}^2$ be a connected region, let $F \subset V$ be a subset satisfying $|F| \leq L(\log L)^\kappa$ and $|\partial F| = O(L^{1-\delta})$, and set $h = H - n$. Then*

$$\widehat{\pi}_V^0(\phi_x \geq -h, \forall x \in F) = (1 + o(1)) \exp \left(-\widehat{\pi}_\infty(\phi_o < -h)|F| \right).$$

Proof. We begin with the lower bound. By FKG,

$$\begin{aligned} \widehat{\pi}_V^0(\phi_x \geq -h, \forall x \in F) &\geq \prod_{x \in F} \widehat{\pi}_V^0(\phi_x \geq -h) = \prod_{x \in F} (1 - \widehat{\pi}_V^0(\phi_x < -h)) \\ &\geq \exp \left(- \sum_{x \in F} \widehat{\pi}_V^0(\phi_x < -h) - \sum_{x \in F} \widehat{\pi}_V^0(\phi_x < -h)^2 \right), \end{aligned} \quad (2.17)$$

using $1 - s \geq \exp(-s - s^2)$ for $0 \leq s \leq \frac{1}{2}$, applied to $\widehat{\pi}_V^0(\phi_x < -h) = o(1)$. Moreover, by Eqs. (2.3) and (2.4) and the definition in Eq. (1.2) of H , we have $\widehat{\pi}_\infty(\phi_o < -h) \leq L^{-1+o(1)}$ (moving from $\widehat{\pi}_\infty(\phi_o = -(h+1))$ for $h = H$ to $h = H - n$ incurs a multiplicative cost of $\exp(nc_0\beta \frac{H}{\log H}) = L^{o(1)}$), and hence by Lemma 2.8 we have that for all x ,

$$\widehat{\pi}_V^0(\phi_o < -h) \leq L^{-1+o(1)}, \quad (2.18)$$

as moving from the upper bound $\widehat{\pi}_\infty(\phi_o < -h)$ to $\widehat{\pi}_V^0(\phi_o = -h)$ incurs an additional multiplicative cost of $\exp(c \frac{H^2}{\log^2 H}) = L^{o(1)}$. This implies that

$$\sum_{x \in F} \widehat{\pi}_V^0(\phi_x < -h)^2 \leq L^{-2+o(1)} |F| \leq L^{-1+o(1)} = o(1)$$

using our hypothesis on $|F|$ (with room to spare: the stringent requirement on $|F|$ would come from the upper bound). Equation (2.18) further implies that

$$\sum_{\substack{x \in F \\ \text{dist}(x, \partial F) \leq \log^2 L}} \widehat{\pi}_V^0(\phi_x < -h) \leq L^{-1+o(1)} |\partial F| \leq L^{-\delta+o(1)} = o(1) \quad (2.19)$$

by our assumption on $|\partial F|$. A straightforward consequence of the Peierls argument of [4] is the following decay of correlation property of $\widehat{\pi}$: for every $x \in F$ and R such that $\text{dist}(x, \partial V) > 2R$,

$$\left\| \widehat{\pi}_V^0(\phi|_{\mathcal{B}_R(x)} \in \cdot) - \widehat{\pi}_\infty(\phi|_{\mathcal{B}_R(o)} \in \cdot) \right\|_{\text{TV}} \leq e^{-c\beta R}. \quad (2.20)$$

(To see this, note that one can couple $\phi \sim \widehat{\pi}_V^0$ and $\phi' \sim \widehat{\pi}_\infty^0$ so that they agree on $\mathcal{B}_R(x)$ and $\mathcal{B}_R(o)$, respectively, provided there is no $*$ -connected path of sites P connecting $\partial B_{R/2}(x)$ to ∂V , along which every site y has $\phi_y \neq 0$ or $\phi'_y \neq 0$; in particular, either ϕ or ϕ' have more than $|P|/2$ nonzero sites along P . This occurs with probability at most $\exp(-c\beta R)$ by Peierls applied to both ϕ, ϕ' .) Thus, if x is such that $\text{dist}(x, \partial V) \geq \log^2 L$ (which holds if $\text{dist}(x, \partial F) \geq \log^2 L$ as $F \subset V$) then

$$|\widehat{\pi}_V^0(\phi_x < -h) - \widehat{\pi}_\infty(\phi_o < -h)| \leq e^{-c\beta \log^2 L},$$

and hence our bound on $|F|$ again yields

$$\sum_{\substack{x \in F \\ \text{dist}(x, \partial F) \geq \log^2 L}} \widehat{\pi}_V^0(\phi_x < -h) \leq |F|(\widehat{\pi}_\infty(\phi_o < -h) + e^{-c\beta \log^2 L}) \leq |F|\widehat{\pi}_\infty(\phi_o < -h) + o(1).$$

Overall, we conclude a lower bound of

$$\widehat{\pi}_V^0(\phi_x \geq -h, \forall x \in F) \geq (1 - o(1)) \exp(-\widehat{\pi}_\infty(\phi_o < -h)|F|). \quad (2.21)$$

We proceed to the upper bound, which is where the stronger assumptions on $|F|$ will be needed. The proof will follow a grid partitioning argument using Bonferroni's inequalities, as was done in [13, Prop. 7.7] and [14, Prop. A.1] for the sos model, in [21, Prop. 6.2] for the 3D Ising model, and in [35, Prop 4.4] for the \mathbb{Z} GFF (where it was applied more crudely, as the focus there was the height histogram of \mathbb{Z} GFF surface as opposed to the fluctuations of its level-lines).

Let

$$\mathfrak{u} := \delta/3, \quad \mathfrak{s} := \delta/4,$$

and partition \mathbb{Z}^2 into squares P_i of side-length $L^{\mathfrak{u}} + L^{\mathfrak{s}}$, and let $Q_i \subset P_i$ be the concentric squares of side-length $L^{\mathfrak{u}}$. We will refer to $P_i \setminus Q_i$ as the *shell* of the square Q_i . Observe that if

$$\begin{aligned} V_1 &= \bigcup \{P_i : P_i \cap F^c \neq \emptyset\}, \\ V_2 &= \bigcup \{(P_i \setminus Q_i) \cap F : P_i \cap V_1 = \emptyset\}, \\ V_3 &= F \setminus (V_1 \cup V_2), \end{aligned}$$

then $|V_1| \leq |\partial F|(\log L)^{2\mathfrak{u}} = O(L^{1-\delta/3})$ by our assumption on ∂F and choice of \mathfrak{u} , while $|V_2| \leq L^{\mathfrak{s}-\mathfrak{u}}|F|$ (each shell $P_i \setminus Q_i$, fully contained in F , has $|P_i|L^{\mathfrak{s}-\mathfrak{u}}$ sites), which is at most $L^{1-\delta/12+o(1)}$ by the assumption on F (and choice of $\mathfrak{u}, \mathfrak{s}$). Combined, we have

$$\widehat{\pi}_\infty(\phi_o < -h)(|V_1| + |V_2|) \leq L^{-\delta/12+o(1)} = o(1),$$

so it suffices to show that, for the increased event that restricts its attention only to V_3 we have

$$\widehat{\pi}_V^0 \left(\bigcap_{x \in V_3} \{\phi_x \geq -h\} \right) \leq (1 + o(1)) \exp \left(-\widehat{\pi}_\infty(\phi_o < -h) |V_3| \right). \quad (2.22)$$

We now aim to decrease the event under consideration via intersecting it with

$$\mathfrak{D}_s = \{\phi \text{ has a no disagreement polymers } \gamma \text{ with } \text{diam}(\gamma) \geq L^s/4\}.$$

Indeed, by the standard Peierls argument,

$$\widehat{\pi}_V^0(\mathfrak{D}_s^c) \leq |V| \exp \left(-(\beta - c)L^s/4 \right) = o \left(\exp \left(-\widehat{\pi}_\infty(\phi_o < -h) |V_3| \right) \right)$$

since $\widehat{\pi}_\infty(\phi_o < -h) |V_3| \leq L^{o(1)}$, and thus Eq. (2.22) will follow once we show that

$$\widehat{\pi}_V^0 \left(\mathfrak{D}_s \cap \bigcap_{x \in V_3} \{\phi_x \geq -h\} \right) \leq (1 + o(1)) \exp \left(-\widehat{\pi}_\infty(\phi_o < -h) |V_3| \right). \quad (2.23)$$

Denoting by $\{i_j\}_{j \geq 1}$ the indices of the Q_i 's that contribute to V_3 , we reveal $\phi|_{P_{i_j}}$ in step j , and denote the associated filtration by $(\mathcal{F}_j)_{j \geq 0}$. For brevity, write

$$Q'_{i_j} := F \cap Q_{i_j}.$$

We will argue that

$$\widehat{\pi}_V^0 \left(\mathfrak{D}_s \cap \bigcap_{x \in Q'_{i_j}} \{\phi_x \geq -h\} \mid \mathcal{F}_{j-1} \right) \leq \exp \left(- \left(1 - e^{-c\beta \frac{h^2}{\log^2 h}} \right) \widehat{\pi}_\infty(\phi_o < -h) |Q'_{i_j}| \right), \quad (2.24)$$

from which Eq. (2.23) will readily follow as $\sum_j |Q'_{i_j}| = |V_3| \leq |F| \leq L(\log L)^\kappa$ by assumption, so

$$\begin{aligned} e^{-c\beta \frac{h^2}{\log^2 h}} \widehat{\pi}_\infty(\phi_o < -h) \sum_j |Q'_{i_j}| &\leq e^{-c\beta \frac{h^2}{\log^2 h} + c_0 \beta n \frac{h}{\log h}} \widehat{\pi}_\infty(\phi_o < -H) |V_3| \\ &\leq e^{-(c\beta - o(1)) \frac{h^2}{\log^2 h}}, \end{aligned} \quad (2.25)$$

using the lower bound of Eq. (2.3) (or the one in Eq. (2.6)) in the first inequality (to move from $\widehat{\pi}_\infty(\phi_o < -h)$ to $\widehat{\pi}_\infty(\phi_o < -H)$ at a cost of $e^{c_0 \beta n \frac{h}{\log h}}$) and Eq. (1.2) and the hypothesis on $|F|$ in the second inequality, as $h^2/\log^2 h \asymp \log L/\log \log L$, thus $(\log L)^C = o(\exp(ch^2/\log^2 h))$ for any fixed $C, c > 0$.

To establish Eq. (2.24), denote by C the outermost $*$ -connected circuit of zeros in P_{i_j} , noting that for every $\phi \in \mathfrak{D}_s$, the distance of C from ∂P_{i_j} cannot exceed $L^s/4$. Letting U denote the interior of C , we see that $B_r(x) \subset U$ for every $x \in Q_{i_j}$ with $r \asymp L^s \gg \frac{h}{\log h}$, our assumption in Theorem 2.5. We now apply Bonferroni's inequalities to infer that

$$\widehat{\pi}_U^0 \left(\bigcap_{x \in Q'_{i_j}} \{\phi_x \geq -h\} \right) \leq 1 - \sum_{x \in Q'_{i_j}} \widehat{\pi}_U^0(\phi_x < -h) + \frac{1}{2} \sum_{\substack{x, y \in Q'_{i_j} \\ x \neq y}} \widehat{\pi}_U^0(\phi_x < -h, \phi_y < -h).$$

By the decay of correlations bound in Eq. (2.20), we can replace $\widehat{\pi}_U^0$ with $\widehat{\pi}_\infty$ in each of the sums at an additive error cost of

$$|Q'_{i_j}|^2 \exp(-c\beta L^s) = \chi_1 |Q'_{i_j}| \widehat{\pi}_\infty(\phi_o < -h) \quad \text{for } \chi_1 < \exp(-(c\beta - o(1))L^s) < L^{-100}.$$

Next, we again use the decorrelation estimate in Eq. (2.20) to infer that

$$\begin{aligned} \sum_{\substack{x,y \in Q'_{i_j} \\ \text{dist}(x,y) \geq \log^2 L}} \widehat{\pi}_\infty(\phi_x < -h, \phi_y < -h) &\leq \left(\sum_{x \in Q'_{i_j}} \widehat{\pi}_\infty(\phi_x < -h) \right)^2 + |Q'_{i_j}|^2 e^{-c\beta \log^2 L} \\ &\leq (\chi_2 + \chi_3) |Q'_{i_j}| \widehat{\pi}_\infty(\phi_o < -h), \end{aligned}$$

where χ_2, χ_3 correspond to the two terms on the right-hand of the first line, and satisfy

$$\chi_2 < L^{-1+2u+o(1)} < L^{-1+\delta}, \quad \chi_3 < \exp(-(c\beta - o(1)) \log^2 L) < L^{-100}.$$

Finally, and this will be the dominant term in our error, we address pairs $x, y \in Q'_{i_j}$ at distance at most $\log^2 L$ as follows. Noting that $\widehat{\pi}_\infty(\phi_x > h, \phi_y > h) \leq (1 + \varepsilon_\beta) \widehat{\pi}_\infty(\phi_x = h+1, \phi_y = h+1)$ by a standard Peierls argument (enumerating the two inner most $(h+2)$ -level-line loops that surround each of x, y (possibly this is the same loop), and subtracting 1 in each of their interiors if nonempty), it suffices to bound the sum of said probability. This is achieved by Eq. (2.8) as follows:

$$\begin{aligned} \sum_{\substack{x,y \in Q'_{i_j} \\ 0 < \text{dist}(x,y) \leq \log^2 L}} \widehat{\pi}_\infty(\phi_x = -h-1, \phi_y = -h-1) &= \sum_{x \in Q'_{i_j}} \widehat{\pi}_\infty(\phi_x = -h-1) \\ &\quad \times \sum_{\substack{y \in Q'_{i_j} \\ 0 < \text{dist}(x,y) \leq \log^2 L}} \widehat{\pi}_\infty(\phi_y = -h-1 \mid \phi_x = -h-1) \\ &\leq \chi_4 \widehat{\pi}_\infty(\phi_o < -h) |Q'_{i_j}|, \end{aligned}$$

where

$$\chi_4 \leq (\log L)^4 \exp\left(-c\beta \frac{h^2}{\log^2 h}\right).$$

Summing these, we see that for $\chi = \sum_{i=1}^4 \chi_i \leq \exp(-(c\beta - o(1)) \frac{h^2}{\log^2 h})$ (dominated by χ_4), we have

$$\widehat{\pi}_U^0 \left(\bigcap_{x \in Q'_{i_j}} \{\phi_x \geq -h\} \right) \leq 1 - (1 - \chi) \widehat{\pi}_\infty^0(\phi_o < -h) |Q'_{i_j}| \leq \exp\left(-(1 - \chi) \widehat{\pi}_\infty^0(\phi_o < -h) |Q'_{i_j}|\right),$$

thus establishing Eq. (2.24) and completing the proof. ■

Remark 2.10. The probability $\widehat{\pi}_\infty(\phi_y = h \mid \phi_x = h)$, which we controlled through (2.8), governed the error-term in the proof of Lemma 2.9. Had we instead used Eq. (2.5)—which features an exponent of $c\beta \frac{h}{\log h}$ as opposed to $c\beta \frac{h^2}{\log^2 h}$ from Eq. (2.8)—it would have competed with the exponent $c_0 \beta n \frac{h}{\log h}$ from Eq. (2.3) which appears in our estimate for $\widehat{\pi}_\infty(\phi_o > -h)$ already when considering $h = H - n$ for $n = 1$. We will need to consider $n = 1$, the second-from-top level-line, already to address the law of the top level-line, as we need to establish that there ample spacing between them (and of course, in Theorem 1.1 we proceed to further obtain the joint law of the m top level-lines).

Remark 2.11. The analysis of $\text{dist}(x, y) \leq \log^2 L$ in the final part of the upper bound generated the term $\chi_4 |F| \widehat{\pi}_\infty(\phi_o < -h) \leq p |F| / L$ for $p = \exp(-c \frac{\log L}{\log \log L}) = L^{-o(1)}$ derived from Eq. (2.8). Thus, we can only obtain an $o(1)$ additive error in the exponent if $|F| = L^{1+o(1)}$, and moreover we can handle $L(\log L)^c$ but not any arbitrary $L^{1+o(1)}$. The analogous setting in sos had $p \leq L^{-1/2}$ (due to the different nature of the large deviation problem), resulting in a much larger applicable domain of area $L^{4/3+o(1)}$ in [14, Prop. A.1] (see the hypothesis on the domain Λ in Eq. (A.2) there).

2.2. Cluster expansion and proof of Proposition 2.3. Cluster expansion for disagreement polymers under without a floor (i.e., under $\widehat{\pi}_V^\xi$) will follow from classical results, as formulated in the next proposition. We will thereafter use it in conjunction with Lemma 2.9 to derive Proposition 2.3.

Proposition 2.12 (Cluster expansion). *Let $V \subset \mathbb{Z}^2$ be a connected domain, and consider the \mathbb{Z} GFF model $\widehat{\pi}_V^\xi$ with boundary conditions ξ that are 1 on a $*$ -connected path in $\partial_v V$ and 0 elsewhere so that they induce a unique disagreement polymer $(\gamma, \{D_i\}, \{h_i\})$ in $V \cup \partial_v V$ that contains boundary disagreements. Then for $\beta \geq \beta_0$, the law of this unique disagreement polymer is given by*

$$\widehat{\pi}_V^\xi(\gamma) = \frac{1}{\widehat{Z}_V^\xi} \exp \left(-\mathcal{E}_\beta(\gamma) + \sum_{\substack{W \subset V \\ W \cap \Delta_\gamma \neq \emptyset}} \Phi(W; \gamma) \right) \quad (2.26)$$

for $\widehat{Z}_V^\xi = \widehat{Z}_V^\xi(\beta)$ and a decoration function $\Phi(W)$ as per Definition 2.2.

Proof. The proof will be a standard application of the framework of the Kotecký–Preiss [31] cluster expansion. Observe that, by Definitions 2.1 and 2.2, for any connected domain V we have

$$\widehat{\pi}_V^0(\phi) = \frac{1}{\widehat{Z}_V^0} \exp \left(- \sum_{\gamma \in \mathcal{P}_\phi} \mathcal{E}_\beta(\gamma) \right),$$

where the sum goes over the disagreement polymers arising from ϕ , each of the form $(\gamma, \{D_i\}, \{h_i\})$. Equivalently, letting \mathcal{P}_V be the set of all disagreement polymers that can arise in $\phi \sim \widehat{\pi}_V^0$ (or, equivalently, arising from any $\phi \sim \widehat{\pi}_V^k$ for boundary conditions all- k), we put

$$\widehat{Z}_{\mathcal{P}_V}^0 = \sum_{\substack{\{\gamma_j\} \subset \mathcal{P}_V \\ \gamma_j \text{ pairwise disjoint}}} \prod_j e^{-\mathcal{E}_\beta(\gamma_j)},$$

so that the partition function \widehat{Z}_V^0 from above is synonymous with $\widehat{Z}_{\mathcal{P}_V}^0$. Since $\mathcal{E}_\beta(\gamma) \geq \beta \mathcal{N}(\gamma)$,

$$\sum_{\gamma: \gamma \cap \gamma_0 \neq \emptyset} e^{(\beta - C_0) \mathcal{N}(\gamma)} e^{-\mathcal{E}_\beta(\gamma)} \leq |\gamma_0| \sum_{\gamma \ni \gamma_0} e^{-C_0 \mathcal{N}(\gamma)} \leq \mathcal{N}(\gamma_0)$$

for a large enough C_0 , satisfying the criterion of the main theorem of [31] for $a(\gamma) = \mathcal{N}(\gamma)$ and $\text{dist}(\gamma) = (\beta - C_0 - 1) \mathcal{N}(\gamma)$ in their notation. By that theorem, one has

$$\log \widehat{Z}_{\mathcal{P}_V}^0 = \sum_{\mathfrak{P} \subset \mathcal{P}_V} \Phi_0(\mathfrak{P}) \quad \text{where} \quad \Phi_0(\mathfrak{P}) = \sum_{\mathfrak{P}' \subset \mathfrak{P}} (-1)^{|\mathfrak{P}| - |\mathfrak{P}'|} \log \widehat{Z}_{\mathfrak{P}'}^0,$$

and the function Φ_0 satisfies $\Phi_0(\{\gamma_j\}) = 0$ if the graph formed by the edges of $\{\gamma_j\}$ is not connected, and $\sum_{\mathfrak{P}: \mathfrak{P} \cap \gamma_0 \neq \emptyset} |\Phi_0(\mathfrak{P})| e^{(\beta - C_0 - 1) \mathcal{N}(\mathfrak{P})} \leq \mathcal{N}(\gamma_0)$, where $\mathcal{N}(\mathfrak{P}) := \sum_j \mathcal{N}(\gamma_j)$ for $\mathfrak{P} = \{\gamma_j\}$. It follows that $|\Phi_0(\mathfrak{P})| \leq \exp(-(\beta - C_1) \mathcal{N}(\mathfrak{P}))$. One can then, as done in [17, Sec. 3.9] (see also [13, §A.2]), define for any subset of vertices $W \subset \mathbb{Z}^2$,

$$\Phi(W) = \sum \left\{ \Phi_0(\mathfrak{P}) : \mathfrak{P} = \{(\gamma_j, \{D_{j,i}\}, \{h_{j,i}\})\} \text{ such that } \bigcup D_{j,i} = W \right\},$$

whence (not just under 0 boundary conditions, but for any all- k boundary conditions)

$$\log \widehat{Z}_V^0 = \sum_{W \subset V} \Phi(W), \quad (2.27)$$

and the aforementioned bound on $\Phi_0(\mathfrak{P})$ implies that (after summing it over all possible $\{h_{j,i}\}$ to get a decay of $\exp(-(\beta - C) \sum |\gamma_j|)$ and noting that the boundary bonds ∂W are a subset of $\bigcup_j \gamma_j$),

$$|\Phi(W)| \leq \exp(-(\beta - C) \mathbf{d}(W)),$$

recalling that $\mathbf{d}(W)$ is the size of the smallest connected set of bonds containing all boundary bonds of W . The fact that Φ is invariant under translation ($\Phi(W) = \Phi(W + x)$) as well as

under rotation by $\pi/2$ or reflection with respect to one of the axes is apparent from its combinatorial construction (as the graphs formed by the γ_j 's are isomorphic to the corresponding translations/rotations/reflections).

Finally, let $\phi \sim \hat{\pi}_V^\xi$ with ξ consisting of an interval of 0 and an interval of 1 as per Proposition 2.12. We will move from Eq. (2.27) to Eq. (2.26) via observing that, by definition, if $(\gamma, \{D_i\}, \{h_i\}) \in \mathcal{P}_\phi$, then

$$\hat{\pi}_V^\xi(\gamma) = \exp(-\mathcal{E}_\beta(\gamma)) \frac{\prod_i \hat{Z}_{D_i^\circ}^{h_i}}{\hat{Z}_V^0}.$$

We combine this with the fact that, by Eq. (2.27),

$$\frac{\prod_i \hat{Z}_{D_i^\circ}^{h_i}}{\hat{Z}_V^0} = \exp\left(-\sum_{W \subset V} \Phi(W) + \sum_i \sum_{W \subset D_i^\circ} \Phi(W)\right) = \exp\left(\sum_{\substack{W \subset V \\ W \cap \Delta_\gamma \neq \emptyset}} \Phi(W)\right), \quad (2.28)$$

to conclude the proof. \blacksquare

Proof of Proposition 2.3. We begin with the standard identity (see the proof of Proposition 2.12)

$$\begin{aligned} \pi_{V;F}^\xi(\gamma) &= \exp(-\mathcal{E}_\beta(\gamma)) \frac{1}{Z_{V;F}^\xi} \prod_i Z_{D_i^\circ;F}^{h_i} \\ &= \frac{\hat{Z}_V^0}{Z_{V;F}^\xi} \exp(-\mathcal{E}_\beta(\gamma)) \frac{1}{\hat{Z}_V^0} \prod_i \hat{Z}_{D_i^\circ;F}^{h_i} \hat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F). \end{aligned}$$

Absorbing $\hat{Z}_V^0/Z_{V;F}^\xi$ into the partition function, then applying Eq. (2.28) to rewrite $(\prod_i \hat{Z}_{D_i^\circ}^{h_i})/\hat{Z}_V^0$ in terms of the decoration function Φ , establishes Eq. (2.1).

To derive Eq. (2.2), we let $\gamma \in E$, and wish to apply Lemma 2.9 to $\hat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F)$. Our hypothesis on F gives $|D_i^\circ \cap F| \leq |F| \leq L(\log L)^\kappa$. Next, $|\partial(D_i^\circ \cap F)| \leq |\partial(D_i \cap F)| + O(|\Delta_\gamma|)$, which is at most $|\partial(D_i \cap F)| + O(|\gamma|)$. We further claim that

$$\partial(D_i \cap F) \subset \gamma \cup \partial F;$$

indeed, every bond b in the left-hand is dual to some edge uv for $u \in D_i \cap F$ and $v \notin D_i \cap F$, which, if $v \in F^c$, is counted in ∂F , and otherwise $v \in D_i^c$ and must have $b \in \gamma$ (as all other bonds of ∂D_i have $v \in V^c \subset F^c$). Altogether, we conclude that

$$|\partial(D_i^\circ \cap F)| \leq |\partial F| + O(|\gamma|) = O(L^{1-\delta})$$

by the definition of E , satisfying the hypothesis of Lemma 2.9. Applying that lemma we obtain that, for $i = 0, 1$,

$$\begin{aligned} \hat{\pi}_{D_i^\circ}^{H-n-i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F) &= \hat{\pi}_{D_i^\circ}^0(\phi_x \geq -(H-n-i), \forall x \in D_i^\circ \cap F) \\ &= (1 + o(1)) \exp(-\hat{\pi}_\infty(\phi_o < -(H-n-i))|D_i \cap F|), \end{aligned}$$

using here that

$$\hat{\pi}_\infty(\phi_o < -(H-n-1))|(D_i \setminus D_i^\circ) \cap F| \leq \hat{\pi}_\infty(\phi_o < -(H-n-1))O(|\gamma|) = L^{-\delta+o(1)} = o(1).$$

Recalling that $|F| \leq |D_0 \cap F| + |D_1 \cap F| + L^{1-\delta}$, we see that

$$\hat{\pi}_\infty(\phi_o < -(H-n-1))|D_1 \cap F| = \hat{\pi}_\infty(\phi_o < -(H-n-1))(|F| - |D_0 \cap F| + O(L^{1-\delta})),$$

whereby we again have $\widehat{\pi}_\infty(\phi_o < -(H - n - 1))L^{1-\delta} = L^{-\delta+o(1)} = o(1)$, and

$$\begin{aligned} \left(\widehat{\pi}_\infty(\phi_o < -(H - n)) - \widehat{\pi}_\infty(\phi_o < -(H - n - 1)) \right) |D_0 \cap F| &= \widehat{\pi}_\infty(\phi_o = -(H - n)) |D_0 \cap F| \\ &= \frac{|D_0 \cap F|}{N_n} \end{aligned}$$

by the definition in Eq. (1.5). Thus, absorbing $\exp[-\widehat{\pi}_\infty(\phi_o < -(H - n - 1))|F|]$ into the partition function (being independent of γ), we find that

$$\pi_{V;F}^\xi(\gamma) \propto (1 + o(1)) \exp \left(-\mathcal{E}_\beta(\gamma) + \frac{|D_0 \cap F|}{N_n} + \sum_{\substack{W \subset V \\ W \cap \Delta_\gamma \neq \emptyset}} \Phi(W) \right) \prod_{i \geq 2} \widehat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F),$$

as required. \blacksquare

3. GEOMETRY OF THE DISAGREEMENT POLYMERS

Our eventual goal is to show that γ behaves like an area tilted random walk. Hence, we would like to show that upon removing the main area tilt $\frac{|D_0|}{N_n}$ from Eq. (2.2), γ falls under the Ornstein–Zernike setup. We will prove this in a more general polymer model setting.

To begin, we need to define the set of legal polymers from A to B . Fix any two vertices $A, B \in (\mathbb{Z}^2)^*$ and a simply connected domain V such that A, B are on ∂V . As ∂V is the boundary in \mathbb{R}^2 of a simply connected domain, consider any conformal map that sends ∂V to the unit circle centered at the origin, mapping A to the point $(0, -1)$ and B to the point $(0, 1)$. Let ξ be the boundary condition which is h along the interval of ∂V that maps to the arc of the unit circle in the upper half plane, and $h - 1$ along the rest of ∂V , where h can be chosen as desired. For every height function on V with boundary condition ξ , there is a unique labeled disagreement polymer γ which contains the boundary disagreements. Define $\mathcal{P}_V(A, B)$ as the set of all such possible disagreement polymers in this setting (note that the choice of h is irrelevant in defining this set, all that matters here is that the boundary heights differ by 1). We can then extend this definition to domains V with infinite volume by defining $\mathcal{P}_V(A, B) = \bigcup_{V' \subset V} \mathcal{P}_{V'}(A, B)$, where the union is over all simply connected $V' \subset V$ which have finite volume.

Next we will define the polymer weights with respect to interaction functions $\Phi(W; \gamma)$.

Definition 3.1. Let $\Phi(W; \gamma)$ denote a function on subsets W of \mathbb{Z}^2 and disagreement polymers γ , which satisfies the following properties:

- (i) If W is not connected, then $\Phi(W; \gamma) = 0$ for all γ .
- (ii) For every γ , the function $W \mapsto \Phi(W; \gamma)$ depends only on $W \cap \Delta_\gamma$.
- (iii) For every W, γ and $x \in \mathbb{Z}^2$, one has $\Phi(W; \gamma) = \Phi(W + x; \gamma + x)$. Moreover, $\Phi(W, \gamma)$ is invariant under a rotation of (W, γ) by a multiple of $\pi/2$ or reflection with respect to one of the axes.
- (iv) There exists a constant $C > 0$ such that for every W, γ , we have $|\Phi(W; \gamma)| \leq \exp(-(\beta - C)\mathbf{d}(W))$, where $\mathbf{d}(W)$ is the size of the smallest connected set of bonds in $(\mathbb{Z}^2)^*$ containing all the boundary bonds of W .

Note that the above describes a more general class of functions compared to Definition 2.2, so that modifications of $\Phi(W; \gamma)$ (denoted as $\Phi_U(W; \gamma)$ and $\Phi'_U(W; \gamma)$ in this section) also satisfy the above properties. With this in mind, we define the weight of γ interacting with a domain U with boundary conditions at height $H - n$ and $H - n - 1$ as

$$\widehat{q}_U^n(\gamma) = \exp \left(-\mathcal{E}_\beta(\gamma) + \sum_{\substack{W \subset U \\ W \cap \Delta_\gamma \neq \emptyset}} \Phi(W; \gamma) \mathbb{1}_{\{W \subset U\}} \right) \prod_{i \geq 2} \widehat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ). \quad (3.1)$$

For brevity, define the shorthand notation

$$\Phi_U(W; \gamma) = \Phi(W; \gamma) \mathbb{1}_{\{W \subset U\}}.$$

Remark 3.2. It is clear that if $\Phi(W; \gamma)$ satisfies the properties of Definition 3.1, then so does $\Phi_U(W; \gamma)$. As usual, the results of this paper hold more generally for any choice of Φ_U that satisfies the decay condition Item (iv) above and such that $\Phi_U(W; \gamma) = \Phi(W; \gamma)$ for any $W \subset U$.

Note that the weight $\hat{q}_U^n(\gamma)$ makes sense even when γ is the disagreement polymer for a domain $V \neq U$. Moreover, the definition of $\hat{q}_U^n(\gamma)$ does not depend on this reference domain V , as $\{D_i\}_{i \geq 2}$ are the finite areas encapsulated by γ . Hence, it makes sense to consider polymer partition functions of the form

$$\hat{Z}_{V,U}^n(A, B) := \sum_{\gamma \in \mathcal{P}_V(A, B)} \hat{q}_U^n(\gamma).$$

If \mathcal{E} is an event about γ , then we define

$$\hat{Z}_{V,U}^n(A, B \mid \mathcal{E}) := \sum_{\gamma \in \mathcal{P}_V(A, B) \cap \mathcal{E}} \hat{q}_U^n(\gamma).$$

When the points A, B are clear from context, we will drop them from the notation so we instead have $\hat{Z}_{V,U}^n$ and $\hat{Z}_{V,U}^n(\mathcal{E})$.

3.1. Non-negative decoration functions and product structure.

Following standard treatment, we now adjust $\mathcal{E}_\beta(\gamma)$ to make the decoration functions $\Phi_U(W; \gamma)$ non-negative (see, e.g., [27, Sec. 3.1] or [7, Sec. 2.4]). Let $\nabla_\gamma = \bigcup_{b=(y, y+e_i) \in \gamma} \{b, b+e_i, b-e_i\}$. Then we can consider the decoration functions given by

$$\Phi'_U(W; \gamma) = |W \cap \nabla_\gamma| e^{-(\beta - C)\mathbf{d}(W)} + \Phi_U(W; \gamma), \quad (3.2)$$

for the constant C from Item (iv) from Definition 3.1 above. It is again clear that if $\Phi(W; \gamma)$ satisfies the properties in Definition 3.1, then so does $\Phi'_U(W; \gamma)$. Moreover, by the decay bound in Item (iv) applied to $\Phi_U(W; \gamma)$, we have that $\Phi'_U(W; \gamma) \geq 0$. Let $c(\beta) = \sum_{W \cap \nabla_\gamma \neq \emptyset} e^{-(\beta - C)\mathbf{d}(W)}$, noting that $c(\beta) \rightarrow 0$ as $\beta \rightarrow \infty$. Then, we obtain that $\sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi_U(W; \gamma) = -3c(\beta)|\gamma| + \sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi'_U(W; \gamma)$. Hence, we can write

$$\hat{q}_U^n(\gamma) = \exp \left(-\mathcal{E}_\beta(\gamma) - 3c(\beta)|\gamma| + \sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi'_U(W; \gamma) \right) \prod_{\substack{i \geq 1 \\ h_i < j-1}} \hat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ).$$

For simplicity of notation, we now define

$$\mathcal{E}_\beta^*(\gamma) = \mathcal{E}_\beta(\gamma) + 3c(\beta)|\gamma| - \log \left(\prod_{\substack{i \geq 1 \\ h_i < j-1}} \hat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ) \right), \quad (3.3)$$

so that

$$\hat{q}_U^n(\gamma) = \exp \left(-\mathcal{E}_\beta^*(\gamma) + \sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi'_U(W; \gamma) \right). \quad (3.4)$$

In many parts of the paper, we will only care about the total interaction. For simplicity, we define as shorthand the notation

$$\mathfrak{I}_U(\gamma) = \sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi'_U(W; \gamma), \quad (3.5)$$

so that we can equivalently write

$$\hat{q}_U^n(\gamma) = \exp \left(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_U(\gamma) \right).$$

Recall that the boundary condition n is needed to determine either the heights h_i , if viewing γ as labeled disagreement bonds, or the energy $\mathcal{E}_\beta(\gamma)$, if viewing γ as a triple $(\gamma, \{D_i\}, \{h_i\})$, so that the term \mathcal{E}_β^* implicitly depends on n . It will also be convenient to define

$$\beta' = \beta - c(\beta).$$

Now define $\Psi_U(W, \gamma) = (\exp(\Phi'_U((W; \gamma)) - 1) \mathbb{1}_{\{W \cap \nabla_\gamma \neq \emptyset\}})$. We have now ensured that $\Psi_U(W, \gamma) \geq 0$, so that the weights can eventually be interpreted as probabilities. For this, observe that we can write

$$\exp \left(\sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi'_U(W; \gamma) \right) = \prod_{W \cap \nabla_\gamma \neq \emptyset} \left((e^{\Phi'_U(W; \gamma)} - 1) + 1 \right) = \sum_{\underline{W} = \{W_i\}} \prod_i \Psi_U(W_i; \gamma),$$

where the final sum is over all possible finite collections of components \underline{W} .

Definition 3.3. Define an animal as a pair $\Gamma = [\gamma, \underline{W}]$, where γ is a disagreement polymer of a height function in some domain V , and \underline{W} is a finite collection of connected components of vertices. We can assign each animal with the weight

$$\hat{q}_U^n(\Gamma) = \exp(-\mathcal{E}_\beta^*(\gamma)) \prod_{W \in \underline{W}} \Psi_U(W; \gamma).$$

Observe that we can again sum over all possible finite collections \underline{W} to write

$$\hat{q}_U^n(\gamma) = \sum_{\Gamma = [\gamma, \underline{W}]} \hat{q}_U^n(\Gamma). \quad (3.6)$$

Hence, we will apply the previous notation of $\mathcal{P}_V(A, B)$ and $\hat{Z}_{V,U}^n(A, B)$ to animals. More precisely, we say $\Gamma \in \mathcal{P}_V(A, B)$ to mean that $\Gamma = [\gamma, \underline{W}]$ for some $\gamma \in \mathcal{P}_V(A, B)$ and any finite collection of components \underline{W} , and $\hat{Z}_{V,U}^n(A, B)$ is also a partition function for animals through Eq. (3.6).

Definition 3.4. We say a point $m = (m_1, m_2)$ is a cut-point of Γ if $m \in \Gamma$ and the intersection of Γ and the vertical line $x = m_1$ is the point m . That is, γ only crosses the line once, and so in particular no components of \underline{W} or finite regions D_i cross the line.

Suppose $\Gamma = \Gamma_1 \circ \Gamma_2$ is the decomposition of Γ before and after a cut point m . More precisely, if $\Gamma = [\gamma, \underline{W}]$, then we can write $\Gamma_1 = [\gamma_1, \underline{W}_1]$ and $\Gamma_2 = [\gamma_2, \underline{W}_2]$ where $\gamma = \gamma_1 \circ \gamma_2$ for $\gamma_1 \in \mathcal{P}_V(A, m)$ and $\gamma_2 \in \mathcal{P}_V(m, B)$, and $\underline{W} = \underline{W}_1 \cup \underline{W}_2$ where all components of \underline{W}_1 lie to the left of m and all components of \underline{W}_2 lie to the right of m . By the definition of the law \hat{q}_U^n , we immediately have

$$\hat{q}_U^n(\Gamma) = \hat{q}_U^n(\Gamma_1) \hat{q}_U^n(\Gamma_2). \quad (3.7)$$

Remark 3.5. The only property of $\mathcal{E}_\beta^*(\gamma)$ that will be used is that

$$\mathcal{E}_\beta^*(\gamma) \geq (\beta - c(\beta)) \mathcal{N}(\gamma) =: \beta' \mathcal{N}(\gamma). \quad (3.8)$$

(Recall from Definition 2.2 that $\mathcal{N}(\gamma)$ is the number of bonds in γ , counting the absolute value of the gradient along the bonds.) Moreover, in bounding the number of disagreement polymers γ that satisfy some criterion, we will only use the fact that each γ is a connected set of dual bonds, where each bond is labeled with an integer whose absolute value represents the multiplicity of that bond. In particular, we will not use any of the additional geometric restrictions that come with the fact that γ is the disagreement polymer of a height function ϕ ; this is only needed to ensure that the regions D_i and heights h_i are well-defined. However, to keep notation consistent, we will still denote the integer label on a bond e by $(\nabla\phi)_e$.

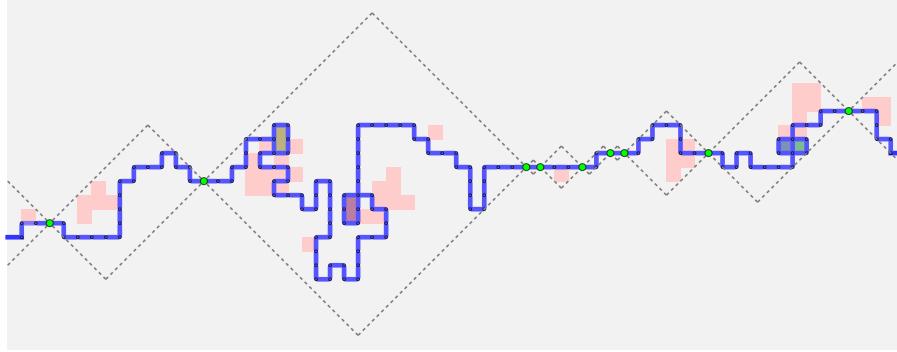


FIGURE 5. An animal $\Gamma = [\gamma, \underline{W}]$, with cone-points in green. The forward and backward cones emanating from the cone-points form diamonds which encapsulate Γ . The disagreement polymer γ is in blue, and the components \underline{W} are in pink.

3.2. Cone-points and irreducible components.

We now define cone-points as a subset of the cut-points. See Fig. 5 for an illustration of the diamonds formed by the cone-points.

Definition 3.6 (Cone-points). The δ -forward cone from u is the set $\mathcal{Y}_\delta^\blacktriangleleft(u) := u + \{(x, y) \in \mathbb{Z}^2 : |y| \leq \delta x\}$. Similarly, the backward cone from u is the set $\mathcal{Y}_\delta^\blacktriangleright(u) = u - \{(x, y) \in \mathbb{Z}^2 : |y| \leq \delta x\}$. When δ is omitted from the notation, we assume $\delta = 1$. When u is omitted from the notation, we assume $u = (0, 0)$. For a disagreement polymer γ , u is a cone-point of γ if $\gamma \subset \mathcal{Y}^\blacktriangleright(u) \cap \mathcal{Y}^\blacktriangleleft(u)$. Similarly, for an animal $\Gamma = [\gamma, \underline{W}]$, u is a cone-point for Γ if $\Gamma \subset \mathcal{Y}^\blacktriangleright(u) \cap \mathcal{Y}^\blacktriangleleft(u)$.

Definition 3.7 (Irreducible components). An animal $\Gamma \in \mathcal{P}_V(u, v)$ is called left-irreducible if it has no cone-points, and $\Gamma \subset \mathcal{Y}^\blacktriangleright(v)$. Similarly, Γ is right-irreducible if it has no cone-points and $\Gamma \subset \mathcal{Y}^\blacktriangleleft(u)$. We say Γ is irreducible if it is both right and left irreducible. The set of all left-irreducible, right-irreducible, and irreducible animals with starting point at o^* is denoted A_L, A_R, A respectively. (So each of these sets consist of Γ in $\bigcup_B \mathcal{P}_V(o^*, B)$.)

A few remarks are in order. First note that the only difference between A being a cone-point for γ vs. Γ is the additional criteria on the clusters in \underline{W} . Note also that by definition, every cone-point is a cut-point. Suppose now that Γ has at least two cone-points. Then, we can decompose Γ according to all of its cone-points, so that

$$\Gamma = \Gamma^{(L)} \circ \Gamma^{(1)} \circ \dots \circ \Gamma^{(k)} \circ \Gamma^{(R)}$$

where $\Gamma^{(L)} \in A_L$, each $\Gamma^{(i)} \in A$, and $\Gamma^{(R)} \in A_R$. Repeatedly applying Eq. (3.7), we obtain the product structure of the weights:

$$\widehat{q}_U^n(\Gamma) = \widehat{q}_U^n(\Gamma^{(L)}) \widehat{q}_U^n(\Gamma^{(R)}) \prod_{i=1}^k \widehat{q}_U^n(\Gamma^{(i)}). \quad (3.9)$$

We will now show that a typical animal Γ has lots of cone-points. First, suppose S is a simple (non self-intersecting) path of dual bonds, and is given the weight $e^{-\beta|S|}$. Define also the partition function

$$Z_V^{SW}(A, B) = \sum_{\substack{S: A \mapsto B \\ S \in V}} e^{-\beta|S|},$$

where the sum is over all simple paths from A to B which stay inside V . We will focus our attention on domains for which a typical simple path S drawn from the measure given by the above weights will have linear length and number of cone-points.

Definition 3.8. Fix $\varepsilon, \delta \in (0, 1)$, and A, B with $B \in \mathcal{Y}_\delta^\blacktriangle(A) \setminus \{A\}$. We call V an ε -nice domain with respect to A, B if there exists $\beta_0, v_0, \delta_0, c > 0$ such that uniformly over $\beta \geq \beta_0$ and $r \geq 1 + \varepsilon$, we have

$$Z_V^{\text{SW}}(A, B \mid |S| \geq r\|A - B\|_1) \leq ce^{-v_0\beta r\|A - B\|_1} Z_V^{\text{SW}}(A, B), \quad (3.10)$$

$$Z_V^{\text{SW}}(A, B \mid |\text{Cpts}(S)| < 2\delta_0\|A - B\|_1) \leq ce^{-v_0\beta\|A - B\|_1} Z_V^{\text{SW}}(A, B). \quad (3.11)$$

Throughout this paper, we will refer back to the following cigar shape.

Definition 3.9 (Cigar shape). Fix A, B with angle denoted $\theta_{A,B}$ and $B_1 - A_1 = M_{A,B}$. Define the curves $\mathcal{C}_{A,B}^\pm$ by

$$\mathcal{C}^\pm(t) = \tan(\theta_{A,B})(t - A_1) + \left(\frac{(t - A_1)(M_{A,B} - t + A_1)}{M_{A,B}} \right)^{1/2} (\log L)^2.$$

Define the cigar shape $\mathcal{C} = \mathcal{C}(\overline{AB})$ as the region in between the curves $\mathcal{C}_{A,B}^+$ and $\mathcal{C}_{A,B}^-$.

Remark. The domains that we will consider in this paper include \mathbb{Z}^2 , the upper half plane \mathbb{H}_+ , and domains containing $\mathcal{C}(\overline{AB})$ or $\mathcal{C}(\overline{AB}) \cap \mathbb{H}_+$. In each setting, Eq. (3.10) is easily satisfied by a Peierls argument mapping S to a minimal length path from A to B . The existence of a linear number of cone-points was established on the domain \mathbb{Z}^2 in [8, Section 2.7], but the proof method is robust and has since been applied to reach the same conclusion for other domains, such as \mathbb{H}_+ in [27, Lemma 3]. By the product structure of Eq. (3.7), this implies that for the domains \mathbb{Z}^2 and \mathbb{H}_+ , the simple paths behave like random walks. As $\mathcal{C}(\overline{AB})$ is defined to contain a typical random walk from A to B , it is then easy to extend the results to domains containing $\mathcal{C}(\overline{AB})$ or $\mathcal{C}(\overline{AB}) \cap \mathbb{H}_+$ as well.

To go from simple paths to disagreement polymers, it will be useful to define the upper and lower envelopes of a disagreement polymer.

Definition 3.10. Suppose $\gamma \in \mathcal{P}_V(A, B)$. The upper envelope $\text{UE}(\gamma)$ is the highest (by lexicographical ordering) simple path from A to B that is a subset of the bonds of γ (so in particular, the labels of the bonds are ignored). We can analogously define the lower envelope $\text{LE}(\gamma)$.

The next lemma states that for nice domains, a typical animal Γ has linear length and number of cone-points.

Lemma 3.11. Fix $\varepsilon, \delta \in (0, 1)$, and A, B with $B \in \mathcal{Y}_\delta^\blacktriangle(A) \setminus \{A\}$. Let V be an ε -nice domain, and U any domain. There exists $\beta_0, v_0, \delta_0, c > 0$ such that uniformly over n , $\beta \geq \beta_0$, and $r \geq 1 + 2\varepsilon$, we have

$$\widehat{Z}_{V,U}^n(A, B \mid |\gamma| \geq r\|A - B\|_1) \leq ce^{-v_0\beta r\|A - B\|_1} \widehat{Z}_{V,U}^n(A, B), \quad (3.12)$$

$$\widehat{Z}_{V,U}^n(A, B \mid |\text{Cpts}(\Gamma)| < 2\delta_0\|A - B\|_1) \leq ce^{-v_0\beta\|A - B\|_1} \widehat{Z}_{V,U}^n(A, B). \quad (3.13)$$

Proof. The proof of Eq. (3.12) is standard, we simply compare γ to γ_0 , a shortest path from 0 to y where all the edge labels are 1. Note that since V is ε -nice, the length of γ_0 is at most $(1 + \varepsilon)\|A - B\|$. We have by Eq. (3.4) and Item (iv) of Definition 3.1 that

$$\widehat{q}_U^n(\gamma) \leq \widehat{q}_U^n(\gamma_0) \exp(-\mathcal{E}_\beta^*(\gamma) + \beta'|\gamma_0| + e^{-(\beta' - C)}(|\gamma| + |\gamma_0|)). \quad (3.14)$$

Now, the number of rooted connected sets of bonds of size k is at most C^k . For each such set of bonds γ , we then obtain the following upper bound by allowing each $(\nabla\phi)_e$ to take on any value in $\mathbb{Z} \setminus \{0\}$, and applying Eq. (3.8):

$$\sum_{\substack{\gamma \in \mathcal{P}_V(A, B), \\ |\gamma|=k}} \exp(-\mathcal{E}_\beta^*(\gamma)) \leq C^k \prod_{e \in \gamma} \sum_{j=1}^{\infty} 2 \exp(-\beta' j) \leq C^k e^{-\beta' k}. \quad (3.15)$$

Hence, we have

$$\begin{aligned}
\widehat{Z}_{V,U}^n(A, B \mid |\gamma| \geq r\|A - B\|_1) &\leq \sum_{k \geq r\|A - B\|_1} \sum_{\substack{\gamma \in \mathcal{P}_V(A, B), \\ |\gamma|=k}} \widehat{q}_U^n(\gamma) \\
&\leq \widehat{q}_U^n(\gamma_0) \sum_{k \geq r\|A - B\|_1} \exp(-\beta' k + \beta'(1 + \varepsilon)\|A - B\|_1 + Ck) \\
&\leq \widehat{q}_U^n(\gamma_0) C \exp(-(\beta' - C)r\|A - B\|_1 + \beta'(1 + \varepsilon)\|A - B\|_1) \\
&\leq C e^{-v_0\beta' r\|A - B\|_1} \widehat{Z}_{V,U}^n(A, B)
\end{aligned}$$

as long as $v_0 < 1 - \frac{1+\varepsilon}{r} - \frac{C}{\beta'}$.

To show Eq. (3.13), we will first show the same inequality holds when considering cone-points of γ instead of Γ . First note that by Item (iv) we have

$$\left| \log \frac{\exp(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_U(\gamma))}{\exp(-\mathcal{E}_\beta^*(\gamma))} \right| \leq e^{-(\beta - C)} |\gamma|. \quad (3.16)$$

As $e^{-(\beta - C)} |\gamma|$ is smaller than $v_0 \beta$ for sufficiently large β , it suffices to prove the bound replacing $\widehat{q}_U^n(\gamma)$ with $\tilde{q}^n(\gamma) := e^{-\mathcal{E}_\beta^*(\gamma)}$. That is, if we define $\tilde{Z}_{V,U}^n(A, B) = \sum_{\gamma \in \mathcal{P}_V(A, B)} \tilde{q}^n(\gamma)$, we will show that

$$\tilde{Z}_{V,U}^n(A, B \mid |\text{Cpts}(\gamma)| < 2\delta_0\|A - B\|_1) \leq c e^{-v_0\beta\|A - B\|_1} \tilde{Z}_{V,U}^n(A, B). \quad (3.17)$$

Now, for every $\gamma \in \mathcal{P}_V(A, B)$, let $S = \text{UE}(\gamma)$. Observe that since each S is a simple path, the energy $\mathcal{E}_\beta^*(\gamma)$ is simply β times the length of the path, $\beta|\gamma|$. (In particular, there are no domains D_i enclosed by γ that contribute an area term to the energy, and so there is no dependence on n anymore). Hence, we have

$$\sum_{S: S = \text{UE}(\gamma), \gamma \in \mathcal{P}_V(A, B)} \tilde{q}^n(S) = Z_V^{\text{SW}}(A, B). \quad (3.18)$$

Moreover, for each γ such that $\text{UE}(\gamma) = S$, we can view the set of bonds in γ as a collection $\{S, (B_e)_{e \in S}\}$ where B_e are connected components of dual bonds indexed by the first bond $e \in S$ that the component is connected to. For every S , denote by $\mathcal{P}_V(S)$ the set of labeled disagreement polymers γ such that $\text{UE}(\gamma) = S$. Using Eq. (3.8), we can upper bound the sum over $\gamma \in \mathcal{P}_V(S)$ by enumerating over collections $\{S, (B_e)_{e \in S}\}$ and then enumerating over adding integer labels on γ :

$$\begin{aligned}
\sum_{\gamma \in \mathcal{P}_V(S)} \tilde{q}^n(\gamma) &\leq \prod_{e \in S} \left(1 + \sum_{k \geq 2} e^{-\beta'(k-1)} + \sum_{k \leq -1} e^{-\beta'(|k|-1)} \right) \left(1 + \sum_{l=1}^{\infty} C^l \left(\sum_{j=1}^{\infty} 2e^{-\beta'j} \right)^l \right) \tilde{q}^n(S) \\
&\leq (1 + C e^{-\beta})^{|S|} \tilde{q}^n(S) \\
&\leq \exp(C e^{-\beta} |S|) \tilde{q}^n(S).
\end{aligned}$$

The sums over k captures the possibilities that at each bond $e \in \gamma$, the gradient $(\nabla\phi)_e$ is either ≥ 2 or ≤ -1 instead of 1. The sum over l includes the possible sizes of the connected component B_e , and there are C^l possible components of size l for some constant C . For each such component B_e , each bond of B_e can also have a height disagreement of j along that edge for $|j| \geq 1$. The above with Eq. (3.18) implies that if \mathcal{E} is measurable with respect to $\text{UE}(\gamma)$, then

$$\tilde{Z}_{V,U}^n(A, B \mid \mathcal{E}, |\text{UE}(\gamma)| \leq 1.1\|A - B\|_1) \leq \exp(C e^{-\beta}\|A - B\|_1) Z_V^{\text{SW}}(A, B \mid \mathcal{E}, |\text{UE}(\gamma)| \leq 1.1\|A - B\|_1). \quad (3.19)$$

Combining Eqs. (3.11) and (3.19) thus implies that

$$\begin{aligned} \tilde{Z}_{V,U}^n(A, B \mid \text{Cpts}(\text{UE}(\gamma))) &< 2\delta_0 \|A - B\|_1, |\text{UE}(\gamma)| \leq 1.1 \|A - B\|_1 \\ &\leq e^{-c\beta \|A - B\|_1} Z_V^{\text{SW}}(A, B) \leq e^{-c\beta \|A - B\|_1} \tilde{Z}_{V,U}^n(A, B), \end{aligned}$$

where the last inequality follows as $Z_V^{\text{SW}}(A, B)$ is just a restriction of $\tilde{Z}_{V,U}^n(A, B)$ to simple paths.

Moreover, since $|\text{UE}(\gamma)| \leq |\gamma|$, we can also bound the event that $\text{UE}(\gamma)$ is too long. We have from Eq. (3.12) and Eq. (3.16) that

$$\tilde{Z}_{V,U}^n(A, B \mid |\text{UE}(\gamma)| > 1.1 \|A - B\|_1) \leq e^{-c\beta \|A - B\|_1} \tilde{Z}_{V,U}^n(A, B). \quad (3.20)$$

Hence, to prove Eq. (3.17), it suffices to show

$$\begin{aligned} \tilde{Z}_{V,U}^n(A, B \mid \text{Cpts}(\gamma) < \frac{3}{2}\delta_0 \|A - B\|_1, \text{Cpts}(\text{UE}(\gamma)) \geq 2\delta_0 \|A - B\|_1, |\text{UE}(\gamma)| \leq 1.1 \|A - B\|_1) \\ &\leq e^{-c\beta \|A - B\|_1} \tilde{Z}_{V,U}^n(A, B). \end{aligned} \quad (3.21)$$

However, observe that in the computation above Eq. (3.19), each component B_e adds on a multiplicative factor of $O(e^{-\beta|B_e|})$ to the weight of $\tilde{q}^n(S)$ (more precisely, this factor is in the interval $[e^{-\beta|C_e|}, (\sum_{j=1}^{\infty} 2e^{-\beta j})^{|C_e|}]$, depending on the admissible assignments of height differences along the edges of B_e). In [7, Propositions 2.12], it was proved that for an animal $\Gamma = [\gamma, \mathbb{W}]$, many cone-points of γ are also cone-points of the whole animal Γ , precisely in the sense of Eq. (3.21). The proof was a combinatorial computation using the fact that each cluster W of the animal contributes a factor of $e^{-(\beta-C)\mathbf{d}(W)}$ to the weight, and hence the exact same computation proves Eq. (3.21) with components B_e taking the role of clusters. This finally establishes Eq. (3.17). From here we can show that many cone-points of γ imply many cone-points of Γ to conclude Eq. (3.13) again using the computation of [7, Propositions 2.12] (this time actually using it to enumerate over clusters in the animal), as it did not matter there that γ had the simpler geometry of a simple path. ■

3.3. Surface tension existence and properties. Fix n (which determines the heights h_i in the weight $\tilde{q}_U^n(\gamma)$), and fix a unit vector \vec{n} with angle θ . Recall the definitions of $\mathcal{P}_{\mathbb{Z}^2}$ and $\hat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n$ from the beginning of Section 3. Let N be such that the point $N\vec{n}$ lies on the lattice. Define the surface tension as

$$\tau_{\beta,n}(\theta) := \tau_{\beta,n}(\vec{n}) := - \lim_{N \rightarrow \infty} \frac{1}{\|N\vec{n}\|_1} \log \hat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(\mathbf{o}^*, N\vec{n}) \quad (3.22)$$

where the limit is over N such that $N\vec{n}$ lies on the lattice. We will use both the notation of \vec{n} and θ . As usual, if the above limit exists, we can extend $\tau_{\beta,n}$ to a function over all of \mathbb{R}^2 by homogeneity. We will show the following proposition about the surface tension.

Proposition 3.12. *The above limit exists and satisfies the following properties:*

- (i) *The convergence of the limit is uniform over all unit vectors \vec{n} .*
- (ii) *$\tau_{\beta,n}$ is analytic as a function from \mathbb{R}^2 to \mathbb{R}^2*
- (iii) *$\tau_{\beta,n}$ is strictly convex, i.e., for any two vectors \mathbf{u}, \mathbf{v} pointing in different directions,*

$$\tau_{\beta,n}(\mathbf{u}) + \tau_{\beta,n}(\mathbf{v}) > \tau_{\beta,n}(\mathbf{u} + \mathbf{v}).$$

- (iv) *$\tau_{\beta,n}$ is symmetric under rotations by $\pi/4$, reflections across the x and y axis and the diagonals $y = \pm x$.*

The statement of Item (iv) follows by construction: it is clear that each listed symmetry defines a bijection from $\mathcal{P}_{\mathbb{Z}^2}(0, N\vec{n})$ to $\mathcal{P}_{\mathbb{Z}^2}(0, N\vec{n}')$, where \vec{n}' is the image of \vec{n} under the symmetry. The energy, $\mathcal{E}_\beta^*(\gamma)$, is preserved under the symmetry, since $\mathcal{N}(\gamma)$ remain unchanged, the encapsulated domains D_i are just rotated/reflected so that their corresponding area terms are unchanged,

and finally the functions $\Phi(W; \gamma)$ are also preserved under symmetries as stated in Item (iii) of Definition 3.1.

The statements of Items (i) to (iii) (and many more properties) regarding the surface tension of a polymer model have been well studied in [17, Ch. 4] in the case where the shape of γ is a contour and $\mathcal{E}_\beta^*(\gamma) = \beta'|\gamma|$. (See [17, Sec. 4.12] for analyticity and uniform convergence, and [17, Sec. 4.21] for strict convexity.) We provide a brief review of the beginning proof ideas there in order to show what adjustments need to be made in our setting. (The careful reader will notice slight differences due to conventional inconsistencies, such as a choice of β vs. 2β in the law of γ — these are inconsequential and hence ignored.)

Remark 3.13. The key point is that the polymer model in [17] is shown to be a perturbation of minimal length paths between the start and end points. Even though our model has more complicated features — γ is allowed to be any connected component of bonds, and there is also an additional penalty in the area term for finite regions enclosed by γ — any portion of γ that is not locally behaving like a minimal length path is sufficiently penalized due to Eq. (3.8), so that our setting is also a perturbation of minimal length paths.

We begin by recalling [17, Secs. 4.1–4.4], which set up the following definitions.⁸ Define the point-to-line polymers

$$\mathcal{P}_{\mathbb{Z}^2}(N) := \bigcup_{\vec{n}} \mathcal{P}_{\mathbb{Z}^2}(\mathbf{o}^*, t_N^{\vec{n}}).$$

Let $h(\gamma)$ be the difference in height between the start and end of γ . That is, if $\gamma \in \mathcal{P}_V(A, B)$, then $h(\gamma) := B_2 - A_2$, the difference between the y -coordinates of B and A . Define also the partition function with a complex parameter z as

$$\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(N, z) = \sum_{\gamma \in \mathcal{P}_{\mathbb{Z}^2}(N)} \exp \left(-\mathcal{E}_\beta^*(\gamma) + \frac{1}{2}\beta' h(\gamma)z + \sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi'_{\mathbb{Z}^2}(W; \gamma) \right)$$

(Ultimately, we only care about $z = 0$, but including this parameter is needed for studying the characteristic polynomial of a random variable in later sections of [17].) Recalling the definitions of cut-points and animals $\Gamma = [\gamma, \underline{W}]$ from Section 3.1 as well as the notation $\Psi(W; \gamma) = (\exp(\Phi'((W; \gamma)) - 1) \mathbb{1}_{\{W \cap \nabla_\gamma \neq \emptyset\}})$, we can define the animal weights $\widehat{q}_z^n(\Gamma)$ for $\Gamma = [\gamma, \underline{W}]$

$$\widehat{q}_z^n(\Gamma) := \exp(-\mathcal{E}_\beta^*(\gamma) + \frac{1}{2}\beta' h(\gamma)z) \prod_{W \in \underline{W}} \Psi(W).$$

As before, we can then obtain a product structure for the weights,

$$\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(N, z) = \sum_{\Gamma_0, \dots, \Gamma_l} \prod_{i=1}^l \widehat{q}_z^n(\Gamma_i),$$

where $\{\Gamma_i\}_{i=1}^l$ is the decomposition by cut points of some animal whose polymer γ is in $\mathcal{P}_{\mathbb{Z}^2}(N)$. We assume that each resulting Γ_i cannot be further decomposed by cut points, and we call such a Γ_i as wild (not to be confused with the previously defined irreducible components, which refer specifically to cone-points).

Next, in [17, Sec. 4.5], the authors consider only Γ capturing the $\beta \rightarrow \infty$ behavior. That is, they consider restricting to the set of $\Gamma = [\gamma, \underline{W}]$ where γ has minimal length in the sense that there are no two horizontal bonds with the same x -coordinate, and $\underline{W} = \emptyset$. This is denoted by $\mathcal{I}_{N, \vec{n}}^\infty$, a notation that we will leave unchanged to emphasize that even though we start with a more complicated set of polymers γ than in [17], once we impose the above restrictions, we are reduced

⁸Dictionary between the notation of this paper and [17], modulo that this paper considers a more general set of polymers: $\mathcal{P}_{\mathbb{Z}^2}(\mathbf{o}^*, t_N^{\vec{n}}) = \mathcal{I}_{N, \vec{n}}$, $\mathcal{P}_{\mathbb{Z}^2}(N) = \mathcal{I}_N$, the complex parameter z is denoted H , $\widehat{q}_z^n(\Gamma_i) = \Psi_H(\hat{\xi}_i)$, $\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(N, z) = \Xi(N, H)$.

to the same set of Γ considered there (and with the same weight as well). The case of no fixed endpoint is also considered, with $\mathcal{I}_N^\infty := \bigcup_{\bar{n}} \mathcal{I}_{N,\bar{n}}^\infty$. The partition function summing over just animals in \mathcal{I}_N^∞ is then directly computed ([17, Eq. 4.5.7-8]) as

$$\sum_{\Gamma \in \mathcal{I}_N^\infty} \hat{q}_z^n(\Gamma) = \left(\frac{\sinh(\beta')}{\cosh(\beta') - \cosh(z\beta')} \right)^N =: (Q_z)^N.$$

Now for each wild animal Γ_i in the decomposition of some Γ , define $J(\Gamma_i)$ as the projection onto the x -coordinate of the starting and ending cut-points of Γ_i . In [17, Sections 4.6–4.7], the focus is on the sum over all possible wild animals projecting onto a fixed interval I ,

$$\sum_{\Gamma_i: J(\Gamma_i) = I} \hat{q}_z^n(\Gamma_i).$$

To be precise, the sum should be interpreted either as restricting to Γ_i rooted at the origin, since $\hat{q}_z^n(\Gamma_i)$ is invariant under vertical shifts of Γ_i . The key bound in [17, Lemma 4.7] shows that this decays exponentially with $(\beta - C)|I|$ relative to the total weight of (not necessarily wild) animals with $J(\Gamma) = I$ and $\Gamma \in \mathcal{I}_N^\infty$. This is a mathematical statement of the perturbation described in Remark 3.13. The proof there relies on the fact that the polymers are contours, and thus will need to be modified to fit our setting. We do this in Lemma 3.14.

In [17, Sec. 4.8], the key bound is used with the cluster expansion machinery of [31] to obtain a formula for $\log \hat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(N, z)$ as a sum of some abstract functions $\Phi(I)$ over all intervals $I \subset [0, N]$. From here onwards, the authors work primarily with the functions Φ and the properties of Φ proven in [31], which otherwise loses the information of the original polymer model. The first exception is in [17, Sec. 4.14] where the polymer weights are reintroduced to add extra parameters to the weights. However, the same cluster expansion machinery is immediately used, justified also by Lemma 3.14 (see the equation below [17, Eq. 4.14.18]), after which the authors work with the resulting $\Phi(I)$ functions again. The second exception is in the proof of the sharp triangle inequality [17, Sec. 4.21]. It is straightforward to check that the proof there requires only the results proved previously in the chapter, and is not impacted by our different polymer setting. Hence, to obtain the results of [17, Chapter 4] in our setting, it suffices to reprove the key bound given in [17, Lemma 4.7], which we now state more explicitly.

Lemma 3.14. *Fix $\delta > 0$. Let z, z_r for $r = 1, \dots, N$ be complex numbers such that $|\Re z| < 2 - \delta/\beta'$. Let s be a real number in $[0, \delta/3]$. For any interval $I \subset [0, N]$ with integer endpoints, define*

$$\hat{X}_{N,s}(I) := \left(\prod_{r \in I} Q_{z_r} \right)^{-1} \sum_{\Gamma: J(\Gamma) = I} \hat{q}_z^n(\Gamma) \exp(s|I|).$$

Then, there exists some β_0 and constant c such that for all N , $\beta \geq \beta_0$, z, z_r , and all such intervals I , we have

$$|\hat{X}_{N,s}(I)| \leq \exp(-(\beta - \beta_0)(|I| - 1)).$$

Proof. Let $\Gamma = (\gamma, W_1, \dots, W_i)$, and let $J(\Gamma) = I = [m', m'']$. We can write $\mathcal{N}(\gamma) = \mathcal{N}_h(\gamma) + \mathcal{N}_v(\gamma)$ where $\mathcal{N}_h(\gamma), \mathcal{N}_v(\gamma)$ denote the contribution to $\mathcal{N}(\gamma)$ from horizontal and vertical bonds of γ , respectively. Moreover, we define $\mathcal{N}'_h(\gamma) := \mathcal{N}_h(\gamma) - (|I| - 1)$ and $\mathcal{N}'_v(\gamma) := \mathcal{N}_v(\gamma) - h(\gamma)$. We begin with the observation that for every integer $m' < m < m''$, either γ intersects the vertical line $x = m$ more than once, or there is some cluster W that intersects the line $x = m$. This implies that

$$|J(\gamma)| - 1 \leq \mathcal{N}'_h(\gamma) + \sum_{W \in \Gamma} d(W). \tag{3.23}$$

(Note that this is off from [17, Eq. 4.7.6] by a factor of $1/2$, because if γ is required to be a contour, then the only way γ can intersect a vertical line more than once is to intersect it three times. The only impact is that the decay bound in the lemma is off by a factor of 2 , which is of no consequence.)

By the assumption in Eq. (3.8), for an upper bound we can replace $\mathcal{E}_\beta^*(\gamma)$ with $\beta' \mathcal{N}(\gamma)$ in $\hat{q}_z^n(\Gamma)$. With this simplification we can now follow the calculations exactly as up to [17, Eq. 4.7.18] to obtain

$$|\hat{X}_{N,\gamma}(I)| \leq e^{-(\beta-\beta_1)(|I|-1)} \sum_{\gamma: J(\gamma)=I} \exp\left(-\frac{\delta}{2}|h(\gamma)| - \beta'(\mathcal{N}(\gamma) - |h(\gamma)|) - \beta_2 \mathcal{N}_h(\gamma)\right), \quad (3.24)$$

where β_1, β_2 are just constants that can be taken arbitrarily large depending on δ , but not dependent on β' . (Note that above [17, Eq. 4.7.13], the additional bound $|h(\gamma)| \leq \mathcal{N}(\gamma)$ is used to upper bound $-\frac{\delta}{2}|h(\gamma)| - \beta'(\mathcal{N}(\gamma) - |h(\gamma)|) \leq -\frac{\delta}{2}\mathcal{N}(\gamma)$, which is in the final expression obtained in [17, Eq. 4.7.18], the analog of Eq. (3.24) above. We will need the sharper bound as we have written.)

It now remains to show that the sum

$$\Xi := \sum_{\gamma: J(\gamma)=I} \exp\left(-\frac{\delta}{2}|h(\gamma)| - \beta'(\mathcal{N}(\gamma) - |h(\gamma)|) - \beta_2 \mathcal{N}_h(\gamma)\right)$$

is bounded by a constant independent of β' . First observe that in order for γ to climb height $h(\gamma)$, we need at least $|h(\gamma)|$ vertical bonds e such that $\text{sgn}((\nabla\phi)_e) = \text{sgn}(h(\gamma))$. Hence, we partition the bonds of γ into the three sets

$$\begin{aligned} B_1 &= \{e \in \gamma : \text{sgn}((\nabla\phi)_e) = \text{sgn}(h(\gamma)), e \text{ is horizontal}\} \\ B_2 &= \{e \in \gamma : \text{sgn}((\nabla\phi)_e) \neq \text{sgn}(h(\gamma)), e \text{ is horizontal}\} \\ B_3 &= \{e \in \gamma : e \text{ is vertical}\}, \end{aligned}$$

and obtain the upper bound

$$\Xi \leq \sum_{\gamma: J(\gamma)=I} \prod_{e \in B_1} e^{-\frac{\delta}{2} - \beta'(|(\nabla\phi)_e|-1)} \prod_{e \in B_2} e^{-\beta' |(\nabla\phi)_e|} \prod_{e \in B_3} e^{-\beta_2 |(\nabla\phi)_e|}. \quad (3.25)$$

We next enumerate over the bonds by partitioning γ into fragments. For every γ such that $J(\gamma) = I$, we consider slicing γ at all vertical lines $x = x_i$ for integer x_i , so that we are cutting all horizontal bonds in half. We end up with a collection of fragments F , each consisting of a (possibly empty) vertical segment γ , and at least one (but possibly more) horizontal half-bond attached to it. Each bond (or half-bond) also inherits the label $(\nabla\phi)_e$ from γ .

For a fragment F , we can denote its height by $h(F)$, and the number of vertical and horizontal bonds by $\mathcal{N}_v(F)$ and $\mathcal{N}_h(F)$ respectively (accounting for the labels as in Definition 2.2) where now $\mathcal{N}_h(F)$ may be a half integer. We additionally define a marked fragment as a pair (F, e) of a fragment and one of its horizontal half-bonds. The marked half-bonds will tell us later how to reconstruct γ given these fragments. Finally, if R is any connected (in \mathbb{R}^2) set of vertical bonds and horizontal half bonds, we call a horizontal half-bond of R open if it is not part of a whole horizontal bond in R . We will consider an ordering of open half-bonds of R by associating to each open half-bond its point at which another half-bond could join to it to create a whole bond, and then applying a lexicographic ordering to such points.

We now construct an injection from γ to a list of marked fragments $\mathfrak{F} = \mathfrak{F}(\gamma)$, to be indexed as $\mathfrak{F}[i]$. Starting at $i = 1$, set $\mathfrak{F}[1]$ to be (F, \emptyset) where F is the fragment containing the unique half-bond of γ intersecting the left-most vertical line $x = m'$ (we allow $e = \emptyset$ for this starting fragment only). At each step i , define R_i to be the subset of γ consisting of the fragments which have been added to \mathfrak{F} by the end of step i , so that R_1 is the fragment $\mathfrak{F}[1]$ (rooted at the origin, where γ is assumed to be rooted). Now let $i = 2$. Let e be the minimal open horizontal half-bond of R_{i-1} . Let F be the fragment in $\gamma \setminus R_{i-1}$ that contains the half-bond e' such that e' and e combine to make a whole bond in γ . Set $\mathfrak{F}[i] = (F, e')$. Then define R_i as instructed above, set $i = i + 1$, and repeat until we have reached $R_N = \gamma$, where N is the number of fragments in γ .

Now we show that this map is injective. Suppose we start with \mathfrak{F} and we know that $\mathfrak{F} = \mathfrak{F}(\gamma)$ for some γ , and that \mathfrak{F} was built through a sequence of R_i as described above (but of course we do

not know what γ or the R_i are). We know the number of fragments in γ is $N := |\mathfrak{F}|$. We start by placing the fragment of $\mathfrak{F}[1]$ so that its leftmost horizontal bond starts at the origin, and call this T_1 . By construction, $T_1 = R_1$. Now set $i = 2$. We know $T_{i-1} = R_{i-1}$. Let e be the minimal open half-bond of T_{i-1} . Calling $\mathfrak{F}[i] = (F, e')$, let T_i be the result of adding F to T_{i-1} so that e' and e join together to make a whole vertical bond. By construction, R_i also results from attaching (F, e') to the minimal open half-bond of R_{i-1} via e' . Since $R_{i-1} = T_{i-1}$, then $R_i = T_i$. Now set $i = i + 1$ and repeat until we have exhausted the whole list \mathfrak{F} , so we have reached T_N . Since $T_N = R_N = \gamma$, we are done as we have constructed γ only from knowing \mathfrak{F} .

Thus, we can upper bound the sum over γ by a sum over marked fragments. Consider F with height $h(F) = k$. We next bound the contribution to Eq. (3.25) of the bonds in such a fragment. We begin with the horizontal half-bonds. There are $2k + 2$ choices of where to put the marked half-bond. Accounting for the possible label of this half bond, its total contribution to the weight is at most

$$(2k + 2) \sum_{j=1}^{\infty} 2e^{-j\beta_2/2}.$$

Similarly, each of the other $2k + 1$ locations can either have a labeled half-bond or not, for a total contribution of

$$(1 + \sum_{j=1}^{\infty} 2e^{-j\beta_2/2})^{2k+1}.$$

For the k vertical bonds, we split up into the two cases of B_1 and B_2 . Assume without loss of generality that $\text{sgn}(h(\gamma)) = 1$, as the -1 case is analogous and the result is the same. For $e \in B_1$, we get a contribution of $\sum_{j \geq 1} e^{-\frac{\delta}{2} - \beta'(j-1)}$. For $e \in B_2$, we get a contribution of $\sum_{j \leq -1} e^{-\beta'|j|}$. Hence, we get a total contribution of

$$(\sum_{j \geq 1} e^{-\frac{\delta}{2} - \beta'(j-1)} + \sum_{j \leq -1} e^{-\beta'|j|})^k.$$

Altogether, we have that

$$\begin{aligned} & \sum_{k=0}^{\infty} \sum_{F: h(F)=k} \prod_{e \in B_1 \cap F} e^{-\frac{\delta}{2} - \beta'(|(\nabla\phi)_e| - 1)} \prod_{e \in B_2 \cap F} e^{-\beta' |(\nabla\phi)_e|} \prod_{e \in B_3 \cap F} e^{-\beta_2 |(\nabla\phi)_e|} \\ & \leq \sum_{k=0}^{\infty} \left((2k + 2) \sum_{j=1}^{\infty} 2e^{-j\beta_2/2} \right) \left(1 + \sum_{j=1}^{\infty} 2e^{-j\beta_2/2} \right)^{2k+1} \left(\sum_{j \geq 1} e^{-\frac{\delta}{2} - \beta'(j-1)} + \sum_{j \leq -1} e^{-\beta'|j|} \right)^k \\ & \leq \sum_{k=0}^{\infty} C e^{-\beta_2/2} k e^{c k e^{-\beta_2/2}} e^{-k\delta/3} \end{aligned} \tag{3.26}$$

for sufficiently large β' compared to δ . We can then choose $\beta_2 = \beta_2(\delta)$ and a constant $C = C(\delta)$ so that $k e^{-k\delta/3 + c k e^{-\beta_2/2}} < C e^{-k\delta/4}$ for all $k \geq 0$. Then the above sum is upper bounded by

$$\sum_{k=0}^{\infty} C e^{-\beta_2/2} k e^{c k e^{-\beta_2/2}} e^{-k\delta/3} \leq C' e^{-\beta_2/2} \frac{1}{1 - e^{-\delta/4}} \leq C'' e^{-\beta_2}. \tag{3.27}$$

Finally, enumerating over γ as a collection of marked fragments (and ignoring any other compatibility conditions) and noting that there must be at least $|I|$ fragments, we have by Eqs. (3.25)

to (3.27) that

$$\begin{aligned}
\Xi &\leq \sum_{N \geq |I|} \sum_{|\mathfrak{F}|=N} \prod_{i=1}^N \prod_{e \in B_1 \cap \mathfrak{F}[i]} e^{-\frac{\delta}{2} - \beta'(|(\nabla \phi)_e| - 1)} \prod_{e \in B_2 \cap \mathfrak{F}[i]} e^{-\beta' |(\nabla \phi)_e|} \prod_{e \in B_3 \cap \mathfrak{F}[i]} e^{-\beta_2 |(\nabla \phi)_e|} \\
&\leq \sum_{N \geq |I|} \left(\sum_F \prod_{e \in B_1 \cap F} e^{-\frac{\delta}{2} - \beta'(|(\nabla \phi)_e| - 1)} \prod_{e \in B_2 \cap F} e^{-\beta' |(\nabla \phi)_e|} \prod_{e \in B_3 \cap F} e^{-\beta_2 |(\nabla \phi)_e|} \right)^N \\
&\leq \sum_{N \geq |I|} (C'' e^{-\beta_2})^N \\
&\leq e^{-(\beta_2 - C)|I|},
\end{aligned}$$

which is a β independent constant. \blacksquare

As a result of the above discussion, we can now conclude the remaining properties of Proposition 3.12, as well as the following results from [17, Chapter 4]:

Proposition 3.15. *The following statements hold for any finite n and sufficiently large β . Let $A = (A_1, A_2), B = (B_1, B_2)$ be any two points, and assume that A is to the left of B . Let $\theta_{A,B}, M_{A,B}, \ell_{A,B}$ denote the angle, horizontal length, and length of the line segment \overline{AB} respectively, and assume that $|\theta_{A,B}|$ is bounded away from $\pi/2$. Let $\mathcal{S} = \mathcal{S}(A, B)$ denote the infinite vertical strip with A, B on its boundary, and let $W \supset \mathcal{S}$.*

(i) [17, Eq. 4.15.4] Let $\bar{h}_t(\gamma)$ be the maximum height of γ above \overline{AB} at the line $x = A_1 + t$. There exist constants $C, c > 0$ such that for any A, B as above,

$$\widehat{Z}_{W, \mathbb{Z}^2}^n(A, B \mid \bar{h}_t(\gamma) \geq j) \leq C \sqrt{M_{A,B}} e^{-c(j \wedge j^2/t)} \widehat{Z}_{W, \mathbb{Z}^2}^n(A, B).$$

Let $\ell_t(\gamma)$ be the number of bonds in γ to the left of $x = A_1 + t$. For some constant $K(\theta_{A,B})$, we have

$$\widehat{Z}_{W, \mathbb{Z}^2}^n(A, B \mid \ell_t(\gamma) \geq j + K(\theta_{A,B})t) \leq C \sqrt{M_{A,B}} e^{-c(j \wedge j^2/t)} \widehat{Z}_{W, \mathbb{Z}^2}^n(A, B).$$

(ii) [17, §4.12.3] There exists $C > 0$ such that for any A, B as above,

$$\begin{aligned}
\left| \log \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A, B) + \tau_{\beta, n} \ell_{A,B} + \frac{1}{2} \log M_{A,B} \right| &\leq C, \text{ and} \\
\left| \log \widehat{Z}_{\mathcal{S}, \mathbb{Z}^2}^n(A, B) + \tau_{\beta, n} \ell_{A,B} + \frac{1}{2} \log M_{A,B} \right| &\leq C.
\end{aligned}$$

Remark. Note that Item (i) above was originally written only for $W = \mathbb{Z}^2$. However, an immediate consequence of Item (ii) is that for any $W \supset \mathcal{S}$,

$$\left| \log \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A, B) - \log \widehat{Z}_{W, \mathbb{Z}^2}^n(A, B) \right| \leq C, \quad (3.28)$$

and this implies the generalization.

Now assume additionally that $1 \leq \ell_{A,B} \ll L$. We next prove two useful lemmas that allow us to switch between domains and modify interactions, as long as we contain the cigar shape capturing typical random walk fluctuations (recall Definition 3.9).

Lemma 3.16. *Let $V \supset \mathcal{C}(\overline{AB})$. Then,*

$$|\log \widehat{Z}_{V, \mathbb{Z}^2}^n(A, B) - \log \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A, B)| \leq C.$$

Consequently, both the large deviation bounds in Item (i) and the surface tension result in Item (ii) of Proposition 3.15 hold for $\widehat{Z}_{V, \mathbb{Z}^2}^n(A, B)$.

Proof. It suffices to consider $V = \mathcal{C}(\overline{AB})$. It is clear that $\widehat{Z}_{\mathcal{C}(\overline{AB}), \mathbb{Z}^2}^n(A, B) \leq \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A, B)$. To show the other side, we take a union bound over Item (i) of Proposition 3.15 to obtain

$$\widehat{Z}_{\mathcal{S}, \mathbb{Z}^2}^n(A, B \mid \gamma \not\subset \mathcal{C}(\overline{AB})) \leq M_{A, B} e^{-c(\log L)^2} \widehat{Z}_{\mathcal{S}, \mathbb{Z}^2}^n(A, B) = o(1) \widehat{Z}_{\mathcal{S}, \mathbb{Z}^2}^n(A, B).$$

This implies that

$$\widehat{Z}_{\mathcal{C}(\overline{AB}), \mathbb{Z}^2}^n(A, B) \geq (1 - o(1)) \widehat{Z}_{\mathcal{S}, \mathbb{Z}^2}^n(A, B),$$

and then we can apply Eq. (3.28) to conclude. \blacksquare

Lemma 3.17. *Let $V \supset \mathcal{C}(\overline{AB})$. Let W contain the infinite vertical strip with sides $x = A_1 + (\log L)^2$ and $x = B_1 - (\log L)^2$. Then, we have*

$$|\log \widehat{Z}_{V, W}^n(A, B) - \log \widehat{Z}_{V, \mathbb{Z}^2}^n(A, B)| \leq C(\log L)^2.$$

Proof. Let $\mathfrak{T}_{A, B}$ be the set of points to the left of $x = A_1 + 2(\log L)^2$ and to the right of $x = B_1 - 2(\log L)^2$. Since W contains the strip, we have

$$|\log \widehat{q}_W^n(\gamma) - \log \widehat{q}_{\mathbb{Z}^2}^n(\gamma)| \leq C e^{-\beta} |\gamma \cap \mathfrak{T}_{A, B}| + e^{-c(\log L)^2} |\gamma \cap \mathfrak{T}_{A, B}^c|.$$

Let $\widehat{\mathbf{P}}$, $\widehat{\mathbf{E}}$ be the probability and expectation for the polymer model with weight $\widehat{q}_{\mathbb{Z}^2}^n$ and partition function $\widehat{Z}_{V, \mathbb{Z}^2}^n$. Then we have

$$\frac{\widehat{Z}_{V, W}^n(A, B)}{\widehat{Z}_{V, \mathbb{Z}^2}^n(A, B)} \leq \widehat{\mathbf{E}}[\exp(C e^{-\beta} |\gamma \cap \mathfrak{T}_{A, B}| + e^{-c(\log L)^2} |\gamma \cap \mathfrak{T}_{A, B}^c|)].$$

An easy Peierls argument shows that

$$\widehat{\mathbf{P}}(|\gamma| \geq 2\ell_{A, B} + j) \leq e^{-(\beta - C)j},$$

while Lemma 3.16 shows that $|\gamma \cap \mathfrak{T}_{A, B}|$ has an exponential tail past $O((\log L)^2)$ that beats the gain of $\exp(C e^{-\beta} |\gamma \cap \mathfrak{T}_{A, B}|)$ for sufficiently large β , so that we have

$$\frac{\widehat{Z}_{V, W}^n(A, B)}{\widehat{Z}_{V, \mathbb{Z}^2}^n(A, B)} \leq e^{C(\log L)^2}.$$

For the lower bound, we have

$$\begin{aligned} \frac{\widehat{Z}_{V, W}^n(A, B)}{\widehat{Z}_{V, \mathbb{Z}^2}^n(A, B)} &\geq \widehat{\mathbf{E}}[\exp(-C e^{-\beta} |\gamma \cap \mathfrak{T}_{A, B}| - e^{-c(\log L)^2} |\gamma \cap \mathfrak{T}_{A, B}^c|)] \\ &\geq e^{-C(\log L)^2} \widehat{\mathbf{P}}(|\gamma| \leq 3\ell_{A, B}, |\gamma \cap \mathfrak{T}_{A, B}| \leq 2(\log L)^2) \\ &\geq (1 - o(1)) e^{-C(\log L)^2}. \end{aligned} \quad \blacksquare$$

We end with a payoff of this subsection - one of the main benefits of comparing two partition functions is the ability to show an event has low probability in one polymer model by bounding its probability in a simpler polymer model. This is stated more generally as follows.

Lemma 3.18. *Let V_1, V_2, U_1, U_2 be any domains such that $V_1 \subset V_2$ and $U_1 \subset U_2$. Suppose we have the bounds on the partition functions*

$$|\log \widehat{Z}_{V_1, U_1}^n(A, B) - \log \widehat{Z}_{V_1, U_2}^n(A, B)| \leq f_1(A, B), \text{ and} \quad (3.29)$$

$$|\log \widehat{Z}_{V_1, U_2}^n(A, B) - \log \widehat{Z}_{V_2, U_2}^n(A, B)| \leq f_2(A, B), \quad (3.30)$$

where we allow the more general form of Φ_U as in Remark 3.2. Then, for any subset $\mathcal{E} \subset \mathcal{P}_{V_1}(A, B)$ such that

$$\widehat{Z}_{V_2, U_2}^n(A, B \mid \mathcal{E}) \leq p \widehat{Z}_{V_1, U_2}^n(A, B),$$

we also have

$$\widehat{Z}_{V_1, U_1}^n(A, B \mid \mathcal{E}) \leq e^{2f_1(A, B) + \frac{1}{2}f_2(A, B)} \sqrt{p} \widehat{Z}_{V_1, U_1}^n(A, B).$$

Proof. Let \mathbf{P}_{V_2, U_2}^n be the probability measure given by

$$\mathbf{P}_{V_2, U_2}^n(\gamma) = \frac{\widehat{q}_{U_2}(\gamma)}{\widehat{Z}_{V_2, U_2}^n(A, B)},$$

and let \mathbf{E}_{V_2, U_2}^n be expectation under \mathbf{P}_{V_2, U_2}^n . Let $\mathbf{P}_{V_1, U_2}^n, \mathbf{P}_{V_1, U_1}^n$ be defined similarly. Define $\Delta \mathfrak{I}(\gamma) = \mathfrak{I}_{U_1}(\gamma) - \mathfrak{I}_{U_2}(\gamma)$. Then, we have

$$\mathbf{E}_{V_1, U_2}^n[\exp(\Delta \mathfrak{I}(\gamma))] = \frac{\widehat{Z}_{V_1, U_1}^n(A, B)}{\widehat{Z}_{V_1, U_2}^n(A, B)} \quad \text{and} \quad \mathbf{P}_{V_1, U_1}^n(\mathcal{E}) = \frac{\mathbf{E}_{V_1, U_2}^n[\mathbb{1}_{\mathcal{E}} \exp(\Delta \mathfrak{I}(\gamma))]}{\mathbf{E}_{V_1, U_2}^n[\exp(\Delta \mathfrak{I}(\gamma))]}.$$

The assumption Eq. (3.29) implies that $\mathbf{E}_{V_1, U_2}^n[\exp(\Delta \mathfrak{I}(\gamma))] \geq e^{-f_1(A, B)}$. Applying the bound with respect to $\widetilde{\Phi}_{U_1} := 2\Phi_{U_1} - 2\Phi_{U_2}$ (justified since $U_1 \subset U_2$), we also have $\mathbf{E}_{V_1, U_2}^n[\exp(2\Delta \mathfrak{I}(\gamma))] \leq e^{f_1(A, B)}$. By Cauchy-Schwarz for the first inequality and Eq. (3.30) for the second, this implies that

$$\mathbf{P}_{V_1, U_1}^n(\mathcal{E}) \leq e^{2f_1(A, B)} \sqrt{\mathbf{P}_{V_1, U_2}^n(\mathcal{E})} \leq e^{2f_1(A, B) + \frac{1}{2}f_2(A, B)} \sqrt{\mathbf{P}_{V_2, U_2}^n(\mathcal{E})},$$

concluding the proof. \blacksquare

3.4. Wulff Shape. In this subsection, we recall a few facts about the Wulff shape, which will allow for a more fine-tuned analysis of the polymers via Proposition 3.19. For any function τ , we can define a corresponding Wulff shape

$$\mathcal{W} = \mathcal{W}(\tau) := \bigcap_{y \in \mathbb{R}^2} \{h \in \mathbb{R}^2 : h \cdot y \leq \tau(y)\}. \quad (3.31)$$

Now for τ being the surface tension $\tau_{\beta, n}$ defined in Eq. (3.22), by definition we have

$$\log \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(o^*, y) = -\tau_{\beta, n}(y)(1 + o_{\|y\|_1}(1)).$$

Thus, the sum $\sum_{y \in \mathbb{Z}^2} e^{h \cdot y} \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(o^*, y)$ converges if and only if $h \cdot y < \tau_{\beta, n}(y)$ for all $y \in \mathbb{Z}^2$ such that $\|y\|_1$ is sufficiently large. But then since $\tau_{\beta, n}$ was defined by homogeneity and is also a continuous function, this is equivalent to requiring that $h \cdot y < \tau_{\beta, n}(y)$ for all $y \in \mathbb{R}^2$. Hence, an equivalent definition for the Wulff shape is

$$\mathcal{W} := \overline{\left\{h \in \mathbb{R}^2 : \sum_{y \in \mathbb{Z}^2} e^{h \cdot y} \mathcal{G}(y) < \infty\right\}}.$$

Now, for any $h \in \mathbb{R}^2$, we can define the weight

$$W^h(\Gamma) = e^{h \cdot X(\Gamma)} \widehat{q}_{\mathbb{Z}^2}^n(\Gamma)$$

where $X(\Gamma)$ is the difference between the ending and starting point of Γ (i.e., if $\Gamma \in \mathcal{P}_{\mathbb{Z}^2}(A, B)$, then $X(\Gamma) := B - A$). A special value of h will be $h_y^n = \nabla \tau_{\beta, n}(y)$, so that by the homogeneity of $\tau_{\beta, n}$, we have $h_y^n \cdot y = \tau_{\beta, n}(y)$. Note that h_y^n is only dependent on the angle of y and not the length. We next state an important result which allows us to view the weights of irreducible components as probabilities, when normalized by $e^{h_y^n \cdot X(\Gamma)}$. Recall the definitions of the sets A, A_L, A_R from Definition 3.7.

Proposition 3.19. *For any $\delta \in (0, 1)$, there exists β_0 such that for all $\beta > \beta_0$, any $y \in \mathcal{Y}_\delta(o^*) \setminus \{o^*\}$, we have*

$$\sum_{\Gamma \in A} \mathbb{P}^{h_y^n}(\Gamma) := \sum_{\Gamma \in A} e^{h_y^n \cdot X(\Gamma)} \widehat{q}_{\mathbb{Z}^2}^n(\Gamma) = 1.$$

Moreover, the expectation is collinear to y , i.e. there exists some constant $\alpha = \alpha(\beta, y) > 0$ such that

$$\mathbb{E}^{h_y^n}[X(\Gamma)] = \alpha y.$$

Lastly, there exists a constant $\nu_g > 0$ such that for $k \geq 1$,

$$\sum_{\Gamma \in \mathcal{A}_L \cup \mathcal{A}_R} \mathbb{P}^{h_y^n}(\Gamma) 1_{\{|\Gamma| > k\}} \leq C e^{-\nu_g \beta k}.$$

This last exponential decay property holds more generally, replacing h_y^n with any $h \in \mathcal{W}$, or replacing $\mathbb{P}^{h_y^n}(\Gamma)$ with $e^{h_y^n \cdot X(\Gamma)} \hat{q}_U^n(\Gamma)$ for any domain U .

Versions of this proposition have appeared before in slightly different settings [7, 27, 29]. The closest to our setting is the treatment in [7], where the definition of cones is exactly the same. The only differences are in the weights $\hat{q}_{\mathbb{Z}^2}^n(\Gamma)$ and the set of admissible Γ , both of which are irrelevant to the proof of [7, Prop. 2.14, Eq. 5.18], which only uses the properties of the surface tension in Eq. (3.22) and the existence of cone-points in Lemma 3.11. Hence, the result also holds in our disagreement polymer setting.

We can now conclude by defining the variance σ_n^2 which appears in Theorem 1.1.

Definition 3.20 (Variance of the effective 2D RW). For the particular choice of $y = (1, 0)$, define σ_n^2 as the variance of the y -coordinate $X(\Gamma)_2$ under the measure $\mathbb{P}^{h_y^n}$ from Proposition 3.19.

4. INITIAL UPPER BOUND ON THE DISPLACEMENT OF A LEVEL-LINE

The goal of this section is to prove Theorem 4.4, stating that w.h.p., the $(H-n)$ level-line contains a translation of Wulff shapes, which in turn provides an initial upper bound on the displacement of the level-lines from the sides of Λ_L . We will use the following shorthand notation:

Definition 4.1. For a subset $A \subseteq \mathbb{Z}^2$, let $\mathcal{E}_n(A)$ be the event that the $H-n$ level-line contains A .

We begin with some notation surrounding the Wulff shape. Recalling the definition of the Wulff shape $\mathcal{W}(\tau)$ from Eq. (3.31), let $\mathcal{W}_1(\tau)$ be $\mathcal{W}(\tau)$ rescaled to have unit area. Define

$$w_1(\tau) = \int_{\partial \mathcal{W}_1(\tau)} \tau(\theta_s) ds,$$

where θ_s is the direction of the normal with respect to $\partial \mathcal{W}_1(\tau)$ at s . Our target shape for the level-lines will be a translation of Wulff shapes.

Definition 4.2. Let $\mathcal{L}^n(\ell, r)$ be the set obtained by first taking the union of all translates of $\ell \mathcal{W}_1(\tau_{\beta, n})$ inside the unit square, and then dilating the shape by a factor of $(1+r)$.

Fix δ as any small constant (e.g., $\delta = 1/10$), and define

$$\begin{aligned} \ell_n &= \ell_n(\delta) := w_1(\tau_{\beta, n}) N_n / 2(1-\delta)L, \\ \kappa_{n,b} &:= N_n^{1/3} (\log L)^b / L. \end{aligned}$$

Remark 4.3. Observe that since $N_0 \leq L/5\beta$, then $\ell_0 \leq w_{1,\beta}/10(1-\delta)\beta$. Moreover, $w_{1,\beta} \approx 4\beta$ (up to a multiplicative $(1 \pm \varepsilon)$). Hence, $\ell_n \leq \ell_0 \leq 1/2$.

We next state the main theorem of this section.

Theorem 4.4. Consider the \mathbb{Z} GFF on an $L \times L$ box with zero b.c. for non-exceptional values of L (as in Theorem 1.1). Recall that $H = H(L) := \max\{h : \hat{\pi}_\infty(\phi_o \geq h) \geq 5\beta/L\}$ and $N_n = N_n(L) := 1/\hat{\pi}_\infty(\phi_o \geq H-n)$. Then, w.h.p., we have

$$\mathcal{E}_n(L \mathcal{L}^n(\ell_n(1+L^{-1/2}), -(H-n)\kappa_{n,15})).$$

In particular, if x is a point on the bottom boundary of Λ at distance at least, say, $L/10$ from the corners, then

- (a) the vertical distance of the $(H - n)$ level-line from x is at most $N_n^{1/3}(\log L)^{16}$;
- (b) by the upper bound of Eq. (2.6), the analogous distance of the $(H - n - 1)$ level-line is at most

$$N_{n+1}^{1/3}(\log L)^{16} \leq N_n^{1/3} \exp\left(-c\sqrt{\frac{\beta \log L}{\log \log L}}\right).$$

Remark 4.5. The shape $LL^n(\ell_0(1 + L^{-1/2}), -(H - n)\kappa_{n,15})$ is flat away from the corners, but around the corners there will be an arc of the Wulff shape. When $N_0 = O(L)$, then the size of this arc is $O(L)$ and the Wulff shape is visible (note that the shape and hence the H level-line will still occupy a $1 - \varepsilon_\beta$ fraction of the sites however). However, for most values of L (i.e., with the exception of a zero logarithmic density set, similar to Remark 1.3), $\ell_0 = o(1)$ so that $LL^n(\ell_0(1 + L^{-1/2}), -(H - n)\kappa_{n,15})$ is flat all the way up to distance $o(L)$ away from the corners, and there is no visible Wulff shape at the $O(L)$ scale. When we consider $n \geq 1$, we always have $\ell_n = o(1)$, so the latter picture is always the case. This is in contrast with the SOS picture where the Wulff shape is always visible, even for the $(H - n)$ level-line for finite n .

Remark 4.6. While our focus is the \mathbb{Z} GFF, and later (in Section 7) we also extend Theorem 4.4 to the $|\nabla\phi|^p$ model for all $p > 1$, we note that Item (a) of the theorem applies also to the case $p = 1$: Using [14, Lemma 5.9] in lieu of Proposition 2.3, the proof of Theorem 4.4 extends to SOS to show that, for $H = \lfloor \frac{1}{4\beta} \log L \rfloor$ and any fixed $n \geq 0$, the vertical distance of its $(H - n)$ level-line is at most $L^{1/3}(\log L)^{16}$ with high probability (refining the L^ε in [14, Thm. 6.2] into a $\text{polylog}(L)$).

4.1. Growth gadget. In this section, we will prove that in a region around a line segment at the correct scale, the level-line will drop below a certain point, motivated by the treatment in [14, Section 5.2]. This will eventually be used (see Lemma 4.20) to show that the level-line drops far enough below where a Wulff shape would be, which will allow us to iteratively grow the region which we know the level-line contains.

We start with a lemma stating that for a rectangle in the middle of Λ_L , we can enforce the desired $H - n, H - n - 1$ boundary conditions at the cost of moving from the rectangle to a wiggly domain approximating the rectangle. The proof will be postponed to Appendix B.

Lemma 4.7. Fix n , and let R be an $\ell_1 \times \ell_2$ rectangle, where $\ell_1, \ell_2 \gg (\log L)^6$. Let $\Lambda_{9L/10} \subset V \subset \Lambda_L$, and $0 \leq k \leq H - n - 1$. Suppose that, w.h.p. under π_V^k , R is contained in the interior of \mathfrak{L}_{n-1} and that there are two points A, B on the sides of ∂R with the top arc of ∂R from A to B in the interior of \mathfrak{L}_n , such that if d is the distance from $\{A, B\}$ to the top or bottom of R , then $\sqrt{\ell_1}(\log L)^2 \leq d$. Then there exists a π_V^k -measurable distribution on connected regions $Q \subset R$ with marked boundary conditions ξ satisfying Items 1 to 4 below, such that the following holds. If \mathcal{A}_1 is the area of the interior of \mathfrak{L}_n intersected with Q , and \mathcal{A}_2 is the area above the $(H - n)$ level-line in Q under π_Q^ξ , then $\mathcal{A}_2 \subset \mathcal{A}_1$ w.h.p. under π_V^k .

- (1) Q is simply connected,
- (2) $\text{dist}(\partial Q, \partial R) \leq (\log L)^2$,
- (3) There exists $A', B' \in \partial Q$ such that $\mathcal{C}(\overline{A'B'}) \subset Q$, $\mathcal{Y}^\blacktriangle(A')$ does not intersect the left side of Q , $\mathcal{Y}^\blacktriangleright(B')$ does not intersect the right side of Q , and $\max\{d(A, A'), d(B, B')\} \leq 2(\log L)^5$.
- (4) The boundary conditions ξ assigns height $H - n$ on the top arc from A' to B' , and $H - n - 1$ on the bottom arc.

The main theorem of this subsection is the following, showing that the $H - n$ level-line drops below a certain height. Fix $a > 4$ and $b \geq 3a$. Motivated by the above lemma, consider a $N_n^{2/3}(\log L)^a \times 2N_n^{2/3}(\log L)^a$ rectangle R with A, B distance at $CN_n^{1/3}(\log L)^b$ above the bottom of R for some $C > 1$. (The requirements on ℓ_1, ℓ_2, d are easily satisfied.) Let $Q = Q_{a,b}$ satisfy the

above properties in Lemma 4.7. Assume for simplicity that the slope of \overline{AB} is $\theta = \theta_{A,B} \in [0, \pi/4]$ and the midpoint of \overline{AB} is the origin. For future reference, we define here also F as the intersection of Q with the parallelogram that shares sides with R , has top and bottom sides parallel to \overline{AB} , with height $N_n^{1/3}(\log L)^b$ and centered at the origin.

Theorem 4.8. *The following holds uniformly over all possible Q as above. Let \mathfrak{L}_n be the (unique) $(H - n)$ level-line induced by ξ . Then, with π_Q^ξ -probability $1 - e^{-c(\log L)^2}$, \mathfrak{L}_n lies below the point $X = (0, Y + \sigma(\log L)^a + (\log L)^2)$, where*

$$Y = -\frac{N_n^{1/3}(\log L)^{2a}}{8(\tau_{\beta,n}(\theta) + \tau_{\beta,n}''(\theta)) \cos(\theta)^3}, \quad (4.1)$$

$$\sigma^2 = \frac{N_n^{2/3}(\log L)^a}{4(\tau_{\beta,n}(\theta) + \tau_{\beta,n}''(\theta)) \cos(\theta)^3}.$$

Remark 4.9. In Section 4.2, we will apply the above for the choice of $a = 5, b = 15$. As the error probability is a $e^{-c(\log L)^2}$, any union bound over a polynomial number of applications of Theorem 4.8 will still have an error probability of the same form.

In contrast to the polymer model studied in Section 3, we now need to reintroduce the primary area term. Let V be a finite simply connected set, and consider $\gamma \in \mathcal{P}_V(A, B)$ for any $A, B \in V$. Let $F \subset V$, and assume for simplicity that $\overline{AB} \subset F$. Define $\mathcal{A}_F(\gamma)$ as the signed area of the region above $\text{UE}(\gamma)$ with respect to the line segment \overline{AB} , after restricting to F . That is, $\mathcal{A}_F(\gamma)$ is equal to $|D_0 \cap F|$ minus the area above the line segment \overline{AB} in F . (Note that although $|D_0|$ depends on the reference domain V , the normalized area $\mathcal{A}_F(\gamma)$ does not.) Define the polymer weight with boundary conditions $H - n$ and $H - n - 1$, area tilt μ on F , and domain of interaction U , by

$$q_{U;F,\mu}^n(\gamma) = \exp \left(-\mathcal{E}_\beta^*(\gamma) + \frac{\mu}{N_n} \mathcal{A}_F(\gamma) + \sum_{W \cap \nabla_\gamma \neq \emptyset} \Phi'_U(W; \gamma) \right) \quad (4.2)$$

$$=: \exp \left(-\mathcal{E}_\beta^*(\gamma) + \frac{\mu}{N_n} \mathcal{A}_F(\gamma) + \mathfrak{I}_U(\gamma) \right),$$

where $\mathcal{E}_\beta^*(\gamma)$ and $\Phi'_U(W; \gamma)$ are as defined in Eqs. (3.2) and (3.3). For convenience, recall that

$$\mathcal{E}_\beta^*(\gamma) = \mathcal{E}_\beta(\gamma) + 3c(\beta)|\gamma| - \log \left(\prod_{\substack{i \geq 1 \\ h_i < j-1}} \widehat{\pi}_{D_i^\circ}^{h_i}(\phi_x \geq 0, \forall x \in D_i^\circ) \right).$$

As usual, we can define the partition function

$$Z_{V,U;F,\mu}^n(A, B) = \sum_{\gamma \in \mathcal{P}_V(A, B)} q_{U;F,\mu}^n(\gamma),$$

with the notation $Z_{V,U;F,\mu}^n(A, B \mid \mathcal{E})$ denoting the restriction of the sum to $\gamma \in \mathcal{E}$, for any event \mathcal{E} . Note that when $\mu = 0$, we recover the definitions of \widehat{q}_U^n and $\widehat{Z}_{V,U}^n$ from Section 3. Of course, $q_{U;F,\mu}^n(\gamma)$ is just a renormalization of the cluster expansion weight $p_{V;F}^\xi(\gamma)$, a fact we will frequently apply without reference.

We next prove a few preliminary lemmas. In order to prove Theorem 4.8, we will in several instances want to show that at a specific point, γ has the geometry of a simple path.

Definition 4.10. For a bond $b \in \gamma$, call $D_b = D_b(\gamma)$ the connected component of all finite regions in $\mathbb{Z}^2 \setminus \gamma$ containing a region that has b as part of its boundary. If there is no such finite region adjacent to b , then $D_b = \emptyset$.

Lemma 4.11. *Fix any V, U, F , and any points $A, B \in V$. Let $b(\gamma)$ be the first horizontal bond where the upper envelope $\text{UE}(\gamma)$ hits the vertical line $x = m$ for a fixed $m \in [A_1, B_1]$. Then there exists a constant $C > 0$ such that for any $k \geq 0$, we have*

$$Z_{V,U;F,\mu}^n(A, B \mid |\partial D_{b(\gamma)}| \geq k) \leq e^{-(\beta-C)k/2} Z_{V,U;F,\mu}^n(A, B).$$

In particular, we have

$$Z_{V,U;F,\mu}^n(A, B \mid D_{b(\gamma)} = \emptyset) \geq (1 - \varepsilon_\beta) Z_{V,U;F,\mu}^n(A, B).$$

Proof. By construction, D_b is the union of some D_i for $i \geq 2$ such that one arc of ∂D_b is a subset of $\text{UE}(\gamma)$ and the remaining arc is a subset of $\text{LE}(\gamma)$. Let t_1, t_2 be the first and last points on $\text{UE}(\gamma)$ that are intersected by ∂D_b . Let $\Delta(D_{b(\gamma)})$ be the length of the shortest path from t_1 to t_2 that does not exit $D_{b(\gamma)}$. In particular, we have $|\partial D_{b(\gamma)}| \geq 2\Delta(D_{b(\gamma)})$, and so

$$|\partial D_{b(\gamma)}| - \Delta(D_{b(\gamma)}) \geq \frac{1}{2}|\partial D_{b(\gamma)}|.$$

Let \mathbf{f} be the map which replaces the portion of γ between t_1 and t_2 by a path attaining length $\Delta(D_{b(\gamma)})$ (which can be chosen arbitrarily — say by the minimal lexicographic one that does not exit $D_{b(\gamma)}$). Note that this map can only lower $\text{UE}(\gamma)$ so that $\mathcal{A}_F(\gamma) \leq \mathcal{A}_F(\mathbf{f}(\gamma))$. Moreover, the map only removes some regions D_i for $i \geq 2$ and does not change any others. Thus, by the definition of $\mathcal{E}_\beta^*(\gamma)$ and the decay property of Φ_U , we have for some $C > 0$ that

$$q_{U;F,\mu}^n(\gamma) \leq q_{U;F,\mu}^n(\mathbf{f}(\gamma)) e^{-(\beta-C)\frac{1}{2}|\partial D_{b(\gamma)}|}. \quad (4.3)$$

Moreover, for a given ω in the image of \mathbf{f} , the number of preimages γ such that $\mathbf{f}(\gamma) = \omega$ and $|\partial D_{b(\gamma)}(\gamma)| = k$ is bounded above by the number of connected components of bonds of size k times the number of possible points t_1 . The former is at most s^k for some universal constant s , and the latter of which is at most k^2 (since we can look at the first intersection of ω with $x = m$, and t_1 is at most distance k from there). Hence, we obtain

$$\begin{aligned} Z_{V,U;F,\mu}^n(A, B \mid |\partial D_{b(\gamma)}| \geq k) &= \sum_{\gamma: |\partial D_{b(\gamma)}| \geq k} q_{U;F,\mu}^n(\gamma) \leq \sum_{j \geq k} \sum_{\omega \in \text{Image}(\mathbf{f})} \sum_{\substack{\gamma \in \mathbf{f}^{-1}(\omega): \\ |\partial D_{b(\gamma)}| = j}} q_{U;F,\mu}^n(\gamma) \quad (4.4) \\ &\leq \sum_{j \geq k} \sum_{\omega \in \text{Image}(\mathbf{f})} \sum_{\substack{\gamma \in \mathbf{f}^{-1}(\omega): \\ |\partial D_{b(\gamma)}| = j}} q_{U;F,\mu}^n(\mathbf{f}(\gamma)) e^{-(\beta-C)\frac{1}{2}j} \\ &\leq \sum_{j \geq k} e^{-(\beta-C')\frac{1}{2}j} Z_{V,U;F,\mu}^n(A, B) \\ &\leq e^{-(\beta-C'')\frac{1}{2}k} Z_{V,U;F,\mu}^n(A, B). \end{aligned}$$

■

Remark 4.12. Since the proof of Lemma 4.11 uses a Peierls map type argument, the results can be strengthened if we have more information about the image of the map \mathbf{f} . More specifically, suppose we know that for all $\gamma \in \mathcal{E}_1$, we have $\mathbf{f}(\gamma) \in \mathcal{E}_2$. Then, in Eq. (4.4) we can start instead with the sum $\sum_{\gamma: |\partial D_{b(\gamma)}| \geq k, \gamma \in \mathcal{E}_1} q_{U;F,\mu}^n(\gamma)$ and replace the sum $\sum_{\omega \in \text{Image}(\mathbf{f})}$ by $\sum_{\omega \in \text{Image}(\mathbf{f}) \cap \mathcal{E}_2}$. What was previously an upper bound of

$$\sum_{\omega \in \text{Image}(\mathbf{f})} q_{U;F,\mu}^n(\mathbf{f}(\gamma)) \leq Z_{V,U;F,\mu}^n(A, B)$$

now becomes

$$\sum_{\omega \in \text{Image}(\mathbf{f}) \cap \mathcal{E}_2} q_{U;F,\mu}^n(\mathbf{f}(\gamma)) \leq Z_{V,U;F,\mu}^n(A, B \mid \mathcal{E}_2),$$

and hence we get

$$Z_{V,U;F,\mu}^n(A, B \mid |\partial D_{b(\gamma)}| \geq k, \mathcal{E}_1) \leq e^{-(\beta-C)k/2} Z_{V,U;F,\mu}^n(A, B \mid \mathcal{E}_2).$$

We next show that we are, with high probability, in a setting where we have cluster expansion.

Lemma 4.13. *Let γ be the disagreement polymer from A to B in the setting above Theorem 4.8, recalling also Q, ξ, F there. Let \mathcal{G}_1 be the event that $|F| - |D_0 \cap F| - |D_1 \cap F| \leq L^{5/6}$ and $|\gamma| \leq 2N_n^{2/3}(\log L)^a$. Then, $\pi_{Q;F}^\xi(\mathcal{G}_1^c) \leq e^{-cL^{1/24}}$. Moreover, if \mathcal{G}^\perp is the event that \mathfrak{L}_n stays above a given horizontal line \mathcal{H} above the bottom of Q , then $\pi_{Q;F}^\xi(\mathcal{G}_1^c \mid \mathcal{G}^\perp) \leq e^{-cL^{1/24}}$.*

Proof. By Proposition 2.3 and comparing with the definition of \hat{q}_Q^n , we have

$$\pi_{Q;F}^\xi(\gamma) \propto \hat{q}_Q^n(\gamma) \prod_{i=0,1} \hat{\pi}_{D_i^\circ}^{H-n-i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F),$$

with the same statement holding for $\pi_{Q;F}^\xi(\gamma \mid \mathcal{G}^\perp)$ except restricting to $\gamma \in \mathcal{G}^\perp$. To obtain a rough lower bound on the probabilities, a standard computation using FKG and monotonicity in the boundary conditions obtains

$$\hat{\pi}_{D_i^\circ}^{H-n-i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F) \geq \frac{1}{2} \exp \left(- \sum_{x \in D_i^\circ \cap F} \hat{\pi}_{D_i^\circ}^0(\phi_x < -(H-n-1)) \right).$$

Using that $\hat{\pi}_{D_i^\circ}^0(\phi_x < -(H-n-1)) \leq L^{-1+o(1)}$ and the fact that $|F| \leq N_n(\log L)^{a+b}$, we obtain

$$\prod_{i=0,1} \hat{\pi}_{D_i^\circ}^{H-n-i}(\phi_x \geq 0, \forall x \in D_i^\circ \cap F) \geq \frac{1}{2} e^{-N_n(\log L)^{a+b} L^{-1+o(1)}} \geq e^{-L^{o(1)}}.$$

Hence, changing the weights to $\hat{q}_Q^n(\gamma)$ can only tilt the measure by a multiplicative factor of $e^{L^{o(1)}}$, so it suffices to prove that

$$\hat{Z}_{Q,Q}^n(A, B \mid \mathcal{G}_1^c) \leq e^{-cL^{1/24}} \hat{Z}_{Q,Q}^n(A, B), \text{ and} \quad (4.5)$$

$$\hat{Z}_{Q,Q}^n(A, B \mid \mathcal{G}_1^c, \mathcal{G}^\perp) \leq e^{-cL^{1/24}} \hat{Z}_{Q,Q}^n(A, B \mid \mathcal{G}^\perp). \quad (4.6)$$

Now in the polymer model, the bound on the length

$$\hat{Z}_{Q,Q}^n(A, B \mid |\gamma| > 2N_n^{2/3}(\log L)^a) \leq e^{-cN_n^{2/3}(\log L)^a} \hat{Z}_{Q,Q}^n(A, B)$$

follows by an easy Peierls argument (take the map that sends γ to a minimal length path from A to B). The same holds when intersecting with \mathcal{G}^\perp , as mapping to the minimal length path stays within the set \mathcal{G}^\perp .

Bounding the area of $|F| - |D_0 \cap F| - |D_1 \cap F| = \bigcup_{i \geq 2} D_i \cap F$ follows essentially the same proof of Lemma 4.11 but simpler since there is no area term anymore. By the above, we can assume that $|\gamma| \leq 2N_n^{2/3}(\log L)^a$. Thus, if $|F| - |D_0 \cap F| - |D_1 \cap F| \geq L^{5/6}$, then by the pigeonhole principle we must have $|D_b| \geq L^{1/12}$ for some $b \in \gamma$ which is adjacent to a vertex in F . So, we fix any bond b adjacent to a vertex in F , and let γ be such that $b \in \gamma$. As in Lemma 4.11, let t_1, t_2 be the first and last points on $\text{UE}(\gamma)$ that are intersected by D_b , and let $\Delta(D_b)$ be the length of the shortest path from t_1 to t_2 that does not exit D_b so that

$$|\partial D_b| - \Delta(D_b) \geq \frac{1}{2} |\partial D_b|.$$

Let \mathbf{f}_b be the map which replaces the portion of γ between t_1 and t_2 by a path attaining length $\Delta(D_b)$. By the definition of $\mathcal{E}_\beta^*(\gamma)$ and the decay property of Φ_Q , we have

$$\hat{q}_Q^n(\gamma) \leq \hat{q}_Q^n(\mathbf{f}_b(\gamma)) e^{-(\beta-C)\frac{1}{2}|\partial D_b|}. \quad (4.7)$$

Moreover, for a given ω in the image of \mathbf{f}_b , the number of preimages γ such that $\mathbf{f}_b(\gamma) = \omega$ and $|\partial D_b(\gamma)| = k$ is bounded above by the number of connected components of bonds of size k , which is at most s^k for some universal constant s . Finally, observe that by the isoperimetric inequality, $|D_b| \geq L^{1/12}$ implies $|\partial D_b| \geq 4L^{1/24}$. Hence, we obtain

$$\begin{aligned} \sum_{\substack{\gamma: b \in \gamma, \\ |D_b| \geq L^{1/12}}} \widehat{q}_Q^n(\gamma) &\leq \sum_{k \geq 4L^{1/24}} \sum_{\omega \in \text{Image}(\mathbf{f}_b)} \sum_{\substack{\gamma \in \mathbf{f}_b^{-1}(\omega): \\ |\partial D_b(\gamma)| = k}} \widehat{q}_Q^n(\gamma) \\ &\leq \sum_{k \geq 4L^{1/24}} \sum_{\omega \in \text{Image}(\mathbf{f}_b)} \sum_{\substack{\gamma \in \mathbf{f}_b^{-1}(\omega): \\ |\partial D_b(\gamma)| = k}} \widehat{q}_Q^n(\mathbf{f}_b(\gamma)) e^{-(\beta - C) \frac{1}{2} k} \\ &\leq \sum_{k \geq 4L^{1/24}} e^{-(\beta - C') \frac{1}{2} k} \widehat{Z}_{Q,Q}^n(A, B) \\ &\leq e^{-(\beta - C'') 2L^{1/24}} \widehat{Z}_{Q,Q}^n(A, B). \end{aligned}$$

We can then take a union bound over the at most $O(N_n(\log L)^{a+b})$ bonds adjacent to a vertex in F to conclude the proof of Eq. (4.5).

By the logic in Remark 4.12, we also obtain Eq. (4.6) as long as we can show that if $\gamma \in \mathcal{G}^\sqcup$, then $\mathbf{f}_b(\gamma) \in \mathcal{G}^\sqcup$, or equivalently that for every D_b , there exists a shortest path from t_1 to t_2 that stays above \mathcal{H} . Without loss of generality, assume \mathcal{H} is at a half-integer height, and suppose for contradiction that there is no such path. Then, take a shortest path P , and let u, v be two consecutive points of P on \mathcal{H} which mark a drop of P below \mathcal{H} (i.e., in between u, v , P lies strictly below \mathcal{H}). Observe that the region sandwiched between P and the arc of $\text{UE}(\gamma)$ between t_1, t_2 is contained in D_b . Since $\gamma \in \mathcal{G}^\sqcup$, $\text{UE}(\gamma)$ must be at or above \mathcal{H} , and in particular at or above the line segment \overline{uv} . Moreover, the arc of P from u to v must be strictly below \overline{uv} . Hence, \overline{uv} is in the sandwiched region, and is therefore in D_b . This means that the path P' which replaces the arc of P between u and v by \overline{uv} is in D_b . Since $|P'| < |P|$, this is a contradiction. ■

The above lemma shows we can restrict our attention to a set of “good” γ for which we have cluster expansion on. We will also want to then consider the partition function with respect to these cluster expansion weights, and it will be convenient to have the partition function sum over all possible $\gamma \subset Q$, and not just the “good” γ from Lemma 4.13. For the same reason as above, this difference is negligible. That is, for Q, F, A, B as above Theorem 4.8 and for any constant $\mu > 0$, we claim that

$$Z_{Q,Q;F,\mu}^n(A, B \mid \mathcal{G}_1^c) \leq e^{-cL^{1/24}} Z_{Q,Q;F,\mu}^n(A, B), \text{ and} \quad (4.8)$$

$$Z_{Q,Q;F,\mu}^n(A, B \mid \mathcal{G}_1^c, \mathcal{G}^\sqcup) \leq e^{-cL^{1/24}} Z_{Q,Q;F,\mu}^n(A, B \mid \mathcal{G}^\sqcup). \quad (4.9)$$

Indeed, the effect of the area term can only tilt the measure by a factor of $O(\exp(-(\log L)^{a+b}))$. Without the area term however, we reduce to showing that

$$\widehat{Z}_{Q,Q}^n(A, B \mid \mathcal{G}_1^c) \leq e^{-cL^{1/24}} \widehat{Z}_{Q,Q}^n(A, B), \text{ and}$$

$$\widehat{Z}_{Q,Q}^n(A, B \mid \mathcal{G}_1^c, \mathcal{G}^\sqcup) \leq e^{-cL^{1/24}} \widehat{Z}_{Q,Q}^n(A, B \mid \mathcal{G}^\sqcup),$$

which were proven above in Eqs. (4.5) and (4.6).

Lemma 4.14. *Let γ be the disagreement polymer from A to B in the setting above Theorem 4.8. Let \mathcal{G}_2 be the event that for all $b \in \gamma$, $|\partial D_b| \leq (\log L)^2$. Then, $\pi_{Q;F}^\xi(\mathcal{G}_2^c) \leq e^{-c(\log L)^2}$. Moreover, if \mathcal{G}^\sqcup is the event that \mathfrak{L}_n stays above a given horizontal line \mathcal{H} above the bottom of Q , then $\pi_{Q;F}^\xi(\mathcal{G}_2^c \mid \mathcal{G}^\sqcup) \leq e^{-c(\log L)^2}$.*

Proof. By Lemma 4.13, it suffices to upper bound $\pi_{Q;F}^\xi(\mathcal{G}_2^c, \mathcal{G}_1)$. This in turn is upper bounded by $\pi_{Q;F}^\xi(\mathcal{G}_2^c \mid \mathcal{G}_1)$, whence we can use Proposition 2.3 for the equality and Eq. (4.8) for the inequality to write

$$\pi_{Q;F}^\xi(\mathcal{G}_2^c \mid \mathcal{G}_1) = (1 + o(1)) \frac{Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}_2^c, \mathcal{G}_1)}{Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}_1)} \leq (1 + o(1)) \frac{Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}_2^c)}{Z_{Q,Q;F,1}^n(A, B)}.$$

For a fixed bond $b \in Q$, we have

$$Z_{Q,Q;F,1}^n(A, B \mid |\partial D_b| > (\log L)^2) \leq e^{-c(\log L)^2} Z_{Q,Q;F,1}^n(A, B)$$

by applying yet another map argument which is a minor modification of the ones in Lemmas 4.11 and 4.13 (in particular, we would apply the map \mathbf{f}_b from Lemma 4.13, but use the bound Eq. (4.3) from Lemma 4.11 since there is an area term). The proof now concludes after taking a union bound over all $b \in Q$. The case conditioning on \mathcal{G}^\square follows similarly as \mathbf{f}_b preserves the event \mathcal{G}^\square , just intersect every partition function above with the event \mathcal{G}^\square and use Eq. (4.9) instead of Eq. (4.8). We record for later use the end result that

$$Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}_2^c, \mathcal{G}^\square) \leq e^{-c(\log L)^2} Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}^\square). \quad (4.10)$$

■

Our last item before the proving the theorem is to provide a bound on partition functions in terms of the functions $\mathcal{G}_\mu^n(\ell, \theta)$, defined as

$$\mathcal{G}_\mu^n(\ell, \theta) = -\tau_{\beta,n}(\theta)\ell + \frac{\ell^3\mu^2}{24(\tau_{\beta,n}(\theta) + \tau_{\beta,n}''(\theta))N_n^2}. \quad (4.11)$$

Let $\delta = 1/10$. Let \mathcal{R}_k be the set of pairs of points (A', B') in Q with distance $\ell_{A', B'} \leq 2^{k(1-\delta)}N_n^{2/3}$, angle $\theta_{A', B'}$, and horizontal distance $M_{A', B'}$. Without loss of generality, assume A' is to the left of B' . Finally, let $\mathfrak{T}_{A', B'}$ be the set of points either to the left of the vertical line $x = A'_1 + 2(\log L)^2$ or to the right of the vertical line $x = B'_1 - 2(\log L)^2$. Fix any $\mu \geq 0$, and define the weights

$$w^n(\gamma) = \exp(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_{\mathbb{Z}^2}(\gamma) + e^{-\beta}|\gamma \cap \mathfrak{T}_{A', B'}| + \frac{\mu}{N_n} \mathcal{A}_F(\gamma)),$$

where the energy $\mathcal{E}_\beta^*(\gamma)$ is defined as in Eq. (3.3) with respect to boundary conditions $H - n$ and $H - n - 1$. The new term $e^{-\beta}|\gamma \cap \mathfrak{T}_{A', B'}|$ in the weight $w^n(\gamma)$ should be thought of as a buffer term that allows us to switch between different interaction functions $\mathfrak{I}_U(\gamma)$ (see Corollary 4.16). It will also be important for later use to note that the exact dimensions of Q and F (in particular, the choice of a, b above Theorem 4.8) play no role in the following proposition.

Proposition 4.15. *Let $Z_{A', B'}^n = \sum_{\gamma \in \mathcal{P}_{\mathbb{Z}^2}(A', B')} w^n(\gamma)$. Then, for all $k \leq \frac{a}{1-\delta} \log_2 \log L$, uniformly over all $A', B' \in \mathcal{R}_k$, we have*

$$Z_{A', B'}^n \leq z_k e^{\mathcal{G}_\mu^n(\ell_{A', B'}, \theta_{A', B'})}$$

where $z_1 = e^{c(\log L)^2}$, $z_k = z_1^{2^{k-1}}$.

Proof. Following [14, Section 5.3], this will be a proof by induction. Let $\widehat{\mathbf{P}}_{A', B'}$ denote the probability measure on γ given by

$$\widehat{\mathbf{P}}_{A', B'}(\gamma) = \frac{\widehat{q}_{\mathbb{Z}^2}^n(\gamma)}{\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A', B')},$$

and let $\widehat{\mathbf{E}}_{A', B'}$ denote the expectation. We begin with the base case taking $k = 1$. Fix A', B' distance $\ell_{A', B'} \leq 2N_n^{2/3}$ apart. Let $h_{\max}(\gamma)$ be the maximum height reached by γ with respect to the line segment $\overline{A'B'}$. By a union bound over Item (i) of Proposition 3.15, we obtain

$$\widehat{\mathbf{P}}_{A', B'}(h_{\max}(\gamma) \geq j) \leq CM_{A', B'}^{3/2} e^{-c(j \wedge j^2/M_{A', B'})}. \quad (4.12)$$

By an easy Peierls argument mapping γ to a minimal length path from A' to B' , we also have

$$\widehat{\mathbf{P}}_{A',B'}(|\gamma| \geq 2\ell_{A',B'} + j) \leq e^{-(\beta-C)j}.$$

Now, by Cauchy-Schwarz, we write

$$\begin{aligned} Z_{A',B'}^n &= \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A', B') \widehat{\mathbf{E}}_{A',B'} \left[e^{\frac{\mu}{N_n} \mathcal{A}_F(\gamma) + e^{-\beta} |\gamma \cap \mathfrak{T}_{A',B'}|} \right] \\ &\leq \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A', B') \widehat{\mathbf{E}}_{A',B'} \left[e^{\frac{2\mu}{N_n} \mathcal{A}_F(\gamma)} \right] \widehat{\mathbf{E}}_{A',B'} \left[e^{2e^{-\beta} |\gamma \cap \mathfrak{T}_{A',B'}|} \right]. \end{aligned}$$

By the fact that $\mathcal{A}_F(\gamma) \leq |\gamma|(h_{\max}(\gamma) \wedge N_n^{1/3}(\log L)^b)$, we have

$$\widehat{\mathbf{E}}_{A',B'} \left[e^{\frac{2\mu}{N_n} \mathcal{A}_F(\gamma)} \right] \leq \widehat{\mathbf{E}}_{A',B'} \left[e^{|\gamma| o(1)} \mathbb{1}_{|\gamma| \geq 5N_n^{2/3}} \right] + \widehat{\mathbf{E}}_{A',B'} \left[e^{\frac{10\mu}{N_n^{1/3}} h_{\max}(\gamma)} \right].$$

The first term on the right can then be bounded by a constant by the tail bound on $|\gamma|$. The second term can be bounded by $CM_{A',B'}^{1/2}$ by the tail bound on $h_{\max}(\gamma)$.

Similarly, the tail bound on $|\gamma \cap \mathfrak{T}_{A',B'}|$ given by Item (i) of Proposition 3.15 implies that $\widehat{\mathbf{E}}_{A',B'}[e^{2e^{-\beta} |\gamma \cap \mathfrak{T}_{A',B'}|}] \leq e^{c(\log L)^2}$. Moreover, by Item (ii) of Proposition 3.15, we have

$$\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A', B') \leq C e^{-\tau_{\beta,n}(\theta_{A',B'}) \ell_{A',B'}}.$$

Putting the above together implies the claim for $k = 1$.

Now we show the induction step. Let $(A', B') \in \mathcal{R}_{k+1}$. Let $b(\gamma)$ be the first horizontal bond where $\mathsf{UE}(\gamma)$ hits the middle vertical line between A' and B' . Let $C = (C_1, C_2)$ be the left endpoint of $b(\gamma)$. Define $\Delta C_2 = C_2 - \frac{A'_2 + B'_2}{2}$, and first consider the case where $|\Delta C_2| \geq N_n^{1/3}(\log L)^{3a}$. We have

$$Z_{A',B'}^n(|\Delta C_2| \geq N_n^{1/3}(\log L)^{3a}) = \widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A', B') \widehat{\mathbf{E}}_{A',B'} \left[\mathbb{1}_{|\Delta C_2| \geq N_n^{1/3}(\log L)^{3a}} e^{\frac{\mu}{N_n} \mathcal{A}_F(\gamma) + e^{-\beta} |\gamma \cap \mathfrak{T}_{A',B'}|} \right].$$

Hence, by another application of Cauchy-Schwarz, it suffices to bound the $\widehat{\mathbf{P}}_{A',B'}$ probability that $|\Delta C_2| \geq N_n^{1/3}(\log L)^{3a}$, which is at most $e^{-c(\log L)^{5a}}$ by Eq. (4.12). Hence, we altogether have

$$Z_{A',B'}^n(|\Delta C_2| \geq N_n^{1/3}(\log L)^{3a}) \leq e^{-\tau_{\beta,n}(\theta_{A',B'}) \ell_{A',B'} - c(\log L)^{5a}}. \quad (4.13)$$

For the case that $|\Delta C_2| \leq N_n^{1/3}(\log L)^{3a}$, we recall the event $D_{b(\gamma)}$ from Lemma 4.11 where we have (noting Remark 4.12)

$$Z_{A',B'}^n(D_{b(\gamma)} = \emptyset, |\Delta C_2| \leq N_n^{1/3}(\log L)^{3a}) \geq (1 - \varepsilon_\beta) Z_{A',B'}^n(|\Delta C_2| \leq N_n^{1/3}(\log L)^{3a}).$$

On the event $D_{b(\gamma)} = \emptyset$, we can split up $\gamma = \gamma_1 \circ \gamma_2$ as the segments of γ before and after C . Write

$$\mathcal{A}_F(\gamma) = \mathcal{A}_F^0 + \mathcal{A}_F(\gamma_1) + \mathcal{A}_F(\gamma_2),$$

where $\mathcal{A}_F(\gamma_1)$ is the signed area above γ_1 w.r.t. the line segment $\overline{A'C}$, $\mathcal{A}_F(\gamma_2)$ is the signed area above γ_2 w.r.t. the line segment $\overline{CB'}$, and \mathcal{A}_F^0 is the signed area of the triangle $A'CB'$ intersected with F (signed negative when C is above $\overline{A'B'}$). Note that $|\mathcal{A}_F^0| \leq \frac{\ell_{A',B'}}{2} |\Delta C_2| \cos(\theta_{A',B'})$. Now, define the correction term

$$\Delta \mathfrak{I}_{\mathbb{Z}^2}(\gamma_1, \gamma_2) := \mathfrak{I}_{\mathbb{Z}^2}(\gamma_1) + \mathfrak{I}_{\mathbb{Z}^2}(\gamma_2) - \mathfrak{I}_{\mathbb{Z}^2}(\gamma). \quad (4.14)$$

By the decay properties of Φ in Definition 3.1, we have that $|\Delta \mathfrak{I}_{\mathbb{Z}^2}(\gamma_1, \gamma_2)| \leq e^{-(\beta-C)} |\gamma_2 \cap \mathfrak{T}_{C,B'}|$. Hence, we have the bound

$$Z_{A',B'}^n(D_{b(\gamma)} = \emptyset, |\Delta C_2| \leq N_n^{1/3}(\log L)^{3a}) \leq \sum_{\gamma: |\Delta C_2| \leq N_n^{1/3}(\log L)^{3a}} e^{\frac{1}{N_n} |\mathcal{A}_F^0|} w^n(\gamma_1) w^n(\gamma_2).$$

Since both (A', C) and (C, B') are in \mathcal{R}_k , by the induction hypothesis, the above display is bounded by

$$z_k^2 \sum_{|\Delta C_2| \leq N_n^{1/3}(\log L)^{3a}} \exp \left(\frac{\mu \ell_{A', B'} |\Delta C_2| \cos(\theta_{A', B'}) + \mathcal{G}_\mu^n(\ell_{A', C}, \theta_{A', C}) + \mathcal{G}_\mu^n(\ell_{C, B'}, \theta_{C, B'})}{2N_n} \right).$$

By a Taylor expansion in θ followed by a Gaussian summation and the convexity of $\tau_{\beta, n}$ (see [14, Prop. 5.11] for details), this is in turn bounded above by

$$C(\beta) \sqrt{\ell_{A', B'}} z_k^2 \exp(\mathcal{G}_\mu^n(\ell_{A', B'}, \theta_{A', B'}))$$

for some constant $C(\beta)$. In total, we have proved that

$$Z_{A', B'}^n(|\Delta C_2| \leq N_n^{1/3}(\log L)^{3a}) \leq (1 + \varepsilon_\beta) C(\beta) \sqrt{\ell_{A', B'}} z_k^2 \exp(\mathcal{G}_\mu^n(\ell_{A', B'}, \theta_{A', B'})),$$

which together with Eq. (4.13), after increasing the constant c in the definition of z_1 to absorb the $C(\beta) \sqrt{\ell_{A', B'}}$ term above, concludes the induction step. \blacksquare

Corollary 4.16. *For any A', B' in Q , we have*

$$Z_{Q, \mathbb{Z}^2; F, \mu}^n(A', B') = \sum_{\gamma \in \mathcal{P}_Q(A', B')} q_{\mathbb{Z}^2; F, \mu}^n(\gamma) \leq \exp(\mathcal{G}_\mu^n(\ell_{A', B'}, \theta_{A', B'}) + O((\log L)^{2a})).$$

Let W be any region containing the infinite vertical strip between the lines $x = A'_1 + (\log L)^2$ and $x = B'_1 - (\log L)^2$. The same bound holds if we replace the weight $q_{\mathbb{Z}^2; F, \mu}^n(\gamma)$ with $q_{W; F, \mu}^n(\gamma)$, or with

$$\exp \left(-\mathcal{E}_\beta^*(\gamma) + \frac{\mu}{N_n} \mathcal{A}_F(\gamma) + \mathcal{I}_W(\gamma) - \Delta \mathcal{I}_W(\gamma, \gamma_{\text{left}}) \right)$$

uniformly over all $\gamma_{\text{left}} \in \mathcal{P}_Q(A, A')$, with $\Delta \mathcal{I}_W(\gamma, \gamma_{\text{left}})$ defined as in Eq. (4.14).

Proof. Note that because of the decay properties of the interaction terms, switching the weights from $\mathcal{I}_{\mathbb{Z}^2}(\gamma)$ to $\mathcal{I}_W(\gamma)$ or to $\mathcal{I}_W(\gamma) - \Delta \mathcal{I}_W(\gamma, \gamma_{\text{left}})$ comes with a factor of at most $e^{-(\beta - C)} |\gamma \cap \mathfrak{T}_{A', B'}|$. Thus, the corollary follows from Proposition 4.15 at $k = \frac{a}{1-\delta} \log_2 \log L$ and the trivial bound $Q \subset \mathbb{Z}^2$. \blacksquare

We are now finally ready to prove Theorem 4.8.

Proof of Theorem 4.8. For simplicity, we may assume that in the construction of Q we have $A' = A$ and $B' = B$, as the shift by $(\log L)^5$ is negligible compared to the margin of error provided in the proof of the theorem. Define the event \mathfrak{U} as the event that \mathfrak{L}_n intersects the line $x = 0$ at a point higher than $X = (0, Y + \sigma(\log L)^a + (\log L)^2)$ (recall the definitions of Y, σ from Eq. (4.1)). Recall the goal of the theorem is to upper bound $\pi_{Q, Q}^\xi(\mathfrak{U})$. Since \mathfrak{U} is a decreasing event in ϕ , by FKG we have $\pi_{Q, Q}^\xi(\mathfrak{U}) \leq \pi_{Q; F}^\xi(\mathfrak{U})$. Then by FKG again we can condition on the decreasing event \mathcal{G}^\sqcup that \mathfrak{L}_n stays above a horizontal line \mathcal{H} distanced $2(\log L)^2$ above the highest point on the bottom of Q , so that $\pi_{Q; F}^\xi(\mathfrak{U}) \leq \pi_{Q; F}^\xi(\mathfrak{U} \mid \mathcal{G}^\sqcup)$.

By Lemmas 4.13 and 4.14, we have $\pi_{Q; F}^\xi(\mathcal{G}_1^c \mid \mathcal{G}^\sqcup) \leq e^{-cL^{1/24}}$ and $\pi_{Q; F}^\xi(\mathcal{G}_2^c \mid \mathcal{G}^\sqcup) \leq e^{-c(\log L)^2}$. We are then left with $\pi_{Q; F}^\xi(\mathfrak{U} \mid \mathcal{G}_1, \mathcal{G}_2 \mid \mathcal{G}^\sqcup)$, which we upper bound by $\pi_{Q; F}^\xi(\mathfrak{U} \mid \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\sqcup)$. Since we are now conditioning on the event \mathcal{G}_1 , and also since \mathfrak{L} and hence $\mathfrak{U}, \mathcal{G}^\sqcup$ can be read from γ , we can use Proposition 2.3 to write the above as the ratio

$$\pi_{Q; F}^\xi(\mathfrak{U} \mid \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\sqcup) = (1 + o(1)) \frac{Z_{Q, Q; F, 1}^n(A, B \mid \mathfrak{U}, \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\sqcup)}{Z_{Q, Q; F, 1}^n(A, B \mid \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\sqcup)}. \quad (4.15)$$

We first upper bound the numerator of Eq. (4.15). Let S be an infinite vertical strip with sides containing the sides of Q . Let W be the extension of Q to S at the top and bottom of Q . (As we will see, the exact choice of S and point of extension are irrelevant.) We show we can move from Q

interactions to W interactions. Indeed, \mathcal{G}_1 implies γ stays a distance $(\log L)^2$ from the top of Q (as such an excursion would force γ to have length $3N_n^{2/3}(\log L)^a$, coming from $2N_n^{2/3}(\log L)^a$ vertical bonds and $N_n^{1/3}(\log L)^a$ horizontal bonds). Moreover, $\mathcal{G}_2 \cap \mathcal{G}^\perp$ implies γ stays distance $(\log L)^2$ away from the bottom of Q . Hence, by the decay properties of Φ and the polynomial bound on the length of γ , we can replace the weights $q_{Q;F,1}^n(\gamma)$ with $q_{W;F,1}^n(\gamma)$ at the cost of $(1 + o(1))$ and then disregard the events $\mathcal{G}_1, \mathcal{G}^\perp$ so that

$$Z_{Q,Q;F,1}^n(A, B \mid \mathfrak{U}, \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\perp) = (1 + o(1))Z_{Q,W;F,1}^n(A, B \mid \mathfrak{U}, \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\perp) \leq Z_{Q,W;F,1}^n(A, B \mid \mathfrak{U}, \mathcal{G}_2). \quad (4.16)$$

Next, let $b(\gamma)$ be the first horizontal bond where $\text{UE}(\gamma)$ hits the line $x = 0$. Let $\underline{\mathfrak{U}}$ be that \mathfrak{U}_n intersects $x = 0$ above $X - (0, (\log L)^2)$. Then, for the map \mathbf{f} in Lemma 4.11, we have $\gamma \in \mathfrak{U} \cap \mathcal{G}_2$ implies $\mathbf{f}(\gamma) \in \underline{\mathfrak{U}} \cap D_{b(\gamma)} = \emptyset$, so that Remark 4.12 gives us

$$Z_{Q,W;F,1}^n(A, B \mid |\partial D_{b(\gamma)}| > 0, \mathfrak{U}, \mathcal{G}_2) < \varepsilon_\beta Z_{Q,W;F,1}^n(A, B \mid D_{b(\gamma)} = \emptyset, \underline{\mathfrak{U}}).$$

The implies that

$$\begin{aligned} Z_{Q,W;F,1}^n(A, B \mid \mathfrak{U}, \mathcal{G}_2) &\leq Z_{Q,W;F,1}^n(A, B \mid D_{b(\gamma)} = \emptyset, \mathfrak{U}) + \varepsilon_\beta Z_{Q,W;F,1}^n(A, B \mid D_{b(\gamma)} = \emptyset, \underline{\mathfrak{U}}) \\ &\leq (1 + \varepsilon_\beta) Z_{Q,W;F,1}^n(A, B \mid D_{b(\gamma)} = \emptyset, \mathfrak{U}), \end{aligned} \quad (4.17)$$

where the first inequality uses the previous display, and the second inequality is because $\mathfrak{U} \subset \underline{\mathfrak{U}}$.

Now let $C = (-\frac{1}{2}, C_2)$ be the left endpoint of $b(\gamma)$. Consider the disjoint union $\{D_{b(\gamma)} = \emptyset\} \cap \underline{\mathfrak{U}} = \underline{\mathfrak{U}}_1 \cup \underline{\mathfrak{U}}_2$, where $\underline{\mathfrak{U}}_1, \underline{\mathfrak{U}}_2$ additionally intersect with the events that C_2 is smaller or bigger than $\widehat{Y} = Y + \frac{1}{2}\sigma(\log L)^a$, respectively.

We first bound $Z_{Q,W;F,1}^n(A, B \mid \underline{\mathfrak{U}}_1)$. Let $\mathbf{E}_\mu, \mathbf{P}_\mu$ be the expectation and probability defined for the polymer measure with weights $q_{W;F,\mu}^n$ and partition function $Z_{Q,W;F,\mu}^n(A, B)$. Then, we have

$$\begin{aligned} Z_{Q,W;F,1}^n(A, B \mid \underline{\mathfrak{U}}_1) &= \widehat{Z}_{Q,W}^n(A, B) \mathbf{E}_0[\mathbb{1}_{\underline{\mathfrak{U}}_1} \exp(\frac{1}{N_n} \mathcal{A}_F(\Gamma))] \\ &\leq C e^{-\tau_{\beta,n}(\theta)\ell + O((\log L)^{2a})} \sqrt{Z_{Q,W;F,2}^n(A, B) / \widehat{Z}_{Q,W}^n(A, B)} \sqrt{\mathbf{P}_0(\underline{\mathfrak{U}}_1)}, \end{aligned}$$

where the equality is by definition, and then we use Cauchy-Schwarz and Corollary 4.16 at $\mu = 0$ to get the inequality. We can also use Corollary 4.16 to upper bound $Z_{Q,W;F,2}^n(A, B)$ and Lemma 3.17 to lower bound $\widehat{Z}_{Q,W}^n(A, B)$, obtaining

$$Z_{Q,W;F,2}^n(A, B) / \widehat{Z}_{Q,W}^n(A, B) \leq e^{C(\log L)^{3a} + C(\log L)^2}.$$

Moreover, the event $\underline{\mathfrak{U}}_1$ implies that γ hits the line $x = 0$ once below \widehat{Y} and once above Y , which means there are two points on the vertical line distanced $\frac{1}{2}\sigma(\log L)^a \geq C(\beta, \theta)N_n^{1/3}(\log L)^{3a/2}$ apart. In the $\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n$ measure, this has probability $e^{-cN_n^{1/3}}$ by Item (i) of Proposition 3.15. By Lemma 3.18 with Lemmas 3.16 and 3.17 as the inputs, we have the same for our \mathbf{P}_0 measure:

$$\mathbf{P}_0(\underline{\mathfrak{U}}_1) \leq e^{-cN_n^{1/3}}.$$

Altogether, we obtain the upper bound

$$Z_{Q,W;F,1}^n(A, B \mid \underline{\mathfrak{U}}_1) \leq C e^{-\tau_{\beta,n}(\theta)\ell} e^{-cN_n^{1/3}}. \quad (4.18)$$

Next we bound $Z_{Q,W;F,1}^n(A, B \mid \underline{\mathfrak{U}}_2)$. We will sum over possible $C_2 \geq \widehat{Y}$. Let $\gamma = \gamma_1 \circ \gamma_2$ be the decomposition of γ into components before and after C . Note that this decomposition is now well-defined because we are on the event $D_{b(\gamma)} = \emptyset$. Note also that even though C is not necessarily a cut-point, the fact that $D_{b(\gamma)} = \emptyset$ still implies that the energy naturally decomposes as $\mathcal{E}_\beta^*(\gamma) = \mathcal{E}_\beta^*(\gamma_1) + \mathcal{E}_\beta^*(\gamma_2)$. As in the proof of Proposition 4.15, we will need the correction term

$$\Delta \mathfrak{I}_W(\gamma_1, \gamma_2) = \mathfrak{I}_W(\gamma_1) + \mathfrak{I}_W(\gamma_2) - \mathfrak{I}_W(\gamma_1 \circ \gamma_2).$$

We can also write $\mathcal{A}_F(\gamma) = \mathcal{A}_F^0 + \mathcal{A}_F(\gamma_1) + \mathcal{A}_F(\gamma_2)$, where $\mathcal{A}_F(\gamma_1)$ is the signed area above γ_1 with respect to the line segment \overline{AC} , $\mathcal{A}_F(\gamma_2)$ is defined analogously with respect to \overline{CB} , and \mathcal{A}_F^0 is the signed area of the triangle ACB intersected with F (signed as positive when C is below \overline{AB} and negative when C is above it). We obtain

$$\begin{aligned} Z_{Q,W;F,1}^n(A, B \mid \mathfrak{U}_2) &\leq \sum_{y \geq \widehat{Y}} \sum_{\gamma: C_2=y} e^{\frac{\mathcal{A}_F^0}{N_n}} e^{-\mathcal{E}_\beta^*(\gamma_1) + \mathfrak{I}_W(\gamma_1) + \frac{\mathcal{A}_F(\gamma_1)}{N_n}} e^{-\mathcal{E}_\beta^*(\gamma_2) + \mathfrak{I}_W(\gamma_2) - \Delta \mathfrak{I}_W(\gamma_1, \gamma_2) + \frac{\mathcal{A}_F(\gamma_2)}{N_n}} \\ &=: \sum_{y \leq \widehat{Y}} e^{\frac{1}{N_n} \mathcal{A}_F^0} \sum_{\gamma_1: C_2=y} e^{-\mathcal{E}_\beta^*(\gamma_1) + \mathfrak{I}_W(\gamma_1) + \frac{1}{N_n} \mathcal{A}_F(\gamma_1)} Z_{\gamma_1, y}. \end{aligned}$$

But now we can apply Corollary 4.16 twice to bound the above by

$$e^{C(\log L)^{2a}} \sum_{y \leq \widehat{Y}} \exp\left(\frac{1}{N_n} \mathcal{A}_F^0 + \mathcal{G}_1^n(\ell_{AC}, \theta_{AC}) + \mathcal{G}_2^n(\ell_{CB}, \theta_{CB})\right).$$

We can further decompose this sum into $\Sigma_1 + \Sigma_2$, where Σ_1 sums over values of $y \geq N_n^{1/3}(\log L)^{3a}$ and Σ_2 sums over the remaining values of $N_n^{1/3}(\log L)^{3a} \geq y \geq \widehat{Y}$. The case of Σ_1 is easy to bound since we have a large negative signed area from \mathcal{A}_F^0 . Indeed, by definition we have

$$\mathcal{G}_\mu^n(\ell, \theta) = -\tau_{\beta, n}(\theta)\ell + O((\log L)^{3a}). \quad (4.19)$$

However, the area term is $\frac{1}{N_n} \mathcal{A}_F^0 \leq -c(\log L)^{4a}$ for some $c > 0$, using here that $b \geq 3a$ so that the triangle ACB is contained in F when $y = N_n^{1/3}(\log L)^{3a}$. Hence, by the convexity of τ in Item (iii) of Proposition 3.12, and Eq. (4.19), we have

$$\Sigma_1 \leq \exp(-\tau_{\beta, n}(\theta)\ell - c(\log L)^{4a}) = \exp(\mathcal{G}_1^n(\ell, \theta) - c(\log L)^{4a}). \quad (4.20)$$

To bound Σ_2 , we wish to do a Taylor expansion on the angles in all the surface tension terms. Through such a computation, combined with the fact that $\mathcal{A}_F^0 = \frac{1}{2} N_n^{2/3}(\log L)^a y$ (since in this regime of y , the triangle ACB is contained in F) and $\ell \cos(\theta) = N_n^{2/3}(\log L)^a$, we eventually obtain

$$\Sigma_2 \leq (1 + o(1)) e^{\mathcal{G}_1^n(\ell, \theta)} \sum_{y \in [\widehat{Y}, N_n^{1/3}(\log L)^{3a}]} e^{-\frac{(y-\widehat{Y})^2}{2\sigma^2}}, \quad (4.21)$$

where we recall the definitions of Y and σ in Eq. (4.1). (See the computations leading up to [14, Eq. 5.11], where the same computation was done, for details. A further elaboration of the steps there can be found in, e.g., [22, Sec. 3.1].) We can now interpret the sum as the probability that a Gaussian with mean Y and variance σ^2 lies in the interval $[\widehat{Y}, N_n^{1/3}(\log L)^{3a}]$. Recall we have defined $\widehat{Y} = Y + \frac{1}{2}\sigma(\log L)^a$, so that this probability is $\exp(-c(\log L)^{2a})$ for some $c > 0$, resulting in the upper bound

$$\Sigma_2 \leq \exp(\mathcal{G}_1^n(\ell, \theta) - c(\log L)^{2a}). \quad (4.22)$$

Combining Eqs. (4.16) to (4.18), (4.20) and (4.22), we obtain as our final upper bound for the numerator of Eq. (4.15)

$$Z_{Q,Q;F,1}^n(A, B \mid \mathfrak{U}, \mathcal{G}_2) \leq \exp(\mathcal{G}_1^n(\ell, \theta) - c'(\log L)^{2a}). \quad (4.23)$$

We now turn to getting a lower bound for the denominator, which we recall is

$$Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\perp).$$

First note that by Eqs. (4.9) and (4.10), we have

$$Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\perp) \geq (1 - o(1)) Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}^\perp). \quad (4.24)$$

The lower bound in [14] consisted mostly of using the proof of [36, Lemma A.6]. We summarize these ideas together here and show the crucial modification needed in our setting (though for details that are unchanged, we will refer the reader to [36]). Begin by defining the optimal curve

$$\gamma_{\text{opt}}(x) = \overline{AB}(x) + \frac{(x - A_1)(B_1 - x)}{2N_n(\tau_{\beta,n}(\theta) + \tau_{\beta,n}''(\theta)) \cos^3(\theta)}. \quad (4.25)$$

Then, we can take a linear approximation \mathcal{L} of γ_{opt} consisting of $\log L$ many line segments. This linear approximation is such that the width of the line segments start at $O(N_n^{2/3}(\log L)^a)$ for the segment containing the middle of γ_{opt} , and decreases by a factor of two with each segment going left/right from there (as was done, e.g., in [36, Lemma A.6]⁹; we reproduce that argument here for completeness). Now let T be any line segment. Denote its endpoints by A_T, B_T , its length by ℓ_T , its angle by θ_T , and its horizontal length by M_T . Recall the definition of the cigar shape $\mathcal{C}(T)$ from Definition 3.9. For brevity, define the union of the cigars as

$$\mathcal{C} = \bigcup_{T \in \mathcal{L}} \mathcal{C}(T).$$

Note that the cigars are constructed such that all γ contained in \mathcal{C} are automatically in \mathcal{G}^\sqcup . We can then lower bound Eq. (4.24) by restricting to polymers which lie inside the union of these cigars,

$$\sum_{\gamma \in \mathcal{P}_Q(A, B) \cap \mathcal{G}^\sqcup} \exp(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_Q(\gamma) + \frac{1}{N_n} \mathcal{A}_F(\gamma)) \geq \sum_{\gamma \in \mathcal{P}_Q(A, B) \cap \mathcal{C}} \exp(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_Q(\gamma) + \frac{1}{N_n} \mathcal{A}_F(\gamma)). \quad (4.26)$$

Now forget about the main area term $\mathcal{A}_F(\gamma)$ for now. The shapes of the cigars imply that we can ignore interactions between the cigars at a cost of $e^{C(\log L)^2}$. Having done so, we can approximate by a product of polymer partition functions:

$$\begin{aligned} \sum_{\gamma \in \mathcal{P}_Q(A, B) \cap \mathcal{C}} \exp(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_Q(\gamma)) &\geq e^{-C(\log L)^2} \prod_{T \in \mathcal{L}} \sum_{\gamma_T \in \mathcal{P}_{\mathcal{C}(T)}(A_T, B_T)} \exp(-\mathcal{E}_\beta^*(\gamma_T) + \mathfrak{I}_Q(\gamma_T)) \\ &= e^{-C(\log L)^2} \prod_{T \in \mathcal{L}} \widehat{Z}_{\mathcal{C}(T), Q}^n(A_T, B_T). \end{aligned} \quad (4.27)$$

Now with the exception of the first and last few cigars close to the sides of Q , all of $\mathcal{C}(T)$ will be distance $\geq (\log L)^2$ from ∂Q so that

$$|\log \widehat{Z}_{\mathcal{C}(T), Q}^n(A_T, B_T) - \log \widehat{Z}_{\mathcal{C}(T), \mathbb{Z}^2}^n| \leq o(1).$$

For the cigars within distance $O((\log L)^2)$ from the sides of Q , we can still switch Q interactions to \mathbb{Z}^2 interactions at a total cost of another $e^{C(\log L)^2}$. By Lemma 3.16, we have

$$\log \widehat{Z}_{\mathcal{C}(T), \mathbb{Z}^2}^n(A_T, B_T) \geq -\tau_{\beta,n}(\theta_T) \ell_T - \frac{1}{2} \log M_T.$$

This implies that

$$\prod_{T \in \mathcal{L}} \widehat{Z}_{\mathcal{C}(T), Q}^n(A_T, B_T) \geq \exp\left(-\sum_{T \in \mathcal{L}} \tau_{\beta,n}(\theta_T) \ell_T - \sum_{T \in \mathcal{L}} \frac{1}{2} \log M_T - C(\log L)^2\right). \quad (4.28)$$

The sum of the $\log M_T$ terms will be at most $O((\log L)^2)$, as there are $\log L$ many terms contributing at most $O(\log L)$. Finally, as in [14] we can exchange the sum $\sum_{T \in \mathcal{L}} \tau_{\beta,n}(\theta_T) \ell_T$ by an integral at

⁹The proof of [36, Lemma A.6] refers to [17, Thm. 4.16]. The upper bound stated in that theorem was later found to be erroneous, but only the lower bound is required for that proof.

a cost of $O(\log L)$ (it will be a constant cost for each T), resulting in the bound

$$\exp\left(-\sum_{T \in \mathcal{L}} \tau_{\beta,n}(\theta_T) \ell_T - \sum_{T \in \mathcal{L}} C \log M_T\right) \geq \exp\left(-\int_{\gamma_{\text{opt}}} \tau_{\beta,n}(\theta_s) ds - O((\log L)^2)\right). \quad (4.29)$$

To summarize, by combining Eqs. (4.27) to (4.29), we get

$$\sum_{\gamma \in \mathcal{P}_Q(A,B) \cap \mathcal{C}} \exp(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_Q(\gamma)) \geq \exp\left(-\int_{\gamma_{\text{opt}}} \tau_{\beta,n}(\theta_s) ds - O((\log L)^2)\right).$$

Now to handle the area term $\frac{1}{N_n} \mathcal{A}_F(\gamma)$ which we removed from the above analysis, note that for every $\gamma \subset \mathcal{C}$, we have

$$\mathcal{A}_F(\gamma) = \mathcal{A}_F(\gamma_{\text{opt}}) + O(N_n^{2/3}(\log L)^a (N_n^{2/3}(\log L)^a)^{1/2} (\log L)^2) = \mathcal{A}_F(\gamma_{\text{opt}}) + O(N_n(\log L)^{3a/2+2}).$$

Thus, we get

$$\sum_{\gamma \in \mathcal{P}_Q(A,B) \cap \mathcal{C}} \exp\left(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_Q(\gamma) + \frac{\mathcal{A}_F(\gamma)}{N_n}\right) \geq e^{(\log L)^{3a/2+2}} e^{\frac{\mathcal{A}_F(\gamma_{\text{opt}})}{N_n}} \exp\left(-\int_{\gamma_{\text{opt}}} \tau(\theta_s) ds - O((\log L)^2)\right). \quad (4.30)$$

We now use the fact that γ_{opt} deviates at most $O(N_n^{1/3}(\log L)^{2a})$ from \overline{AB} , so since $b > 2a$ it never leaves the region F . Thus, it is the same as if $\mathcal{A}_F(\gamma_{\text{opt}})$ were defined with no area restriction to F , and as remarked in [14, Eq (5.17)], we also have

$$-\int_{\gamma_{\text{opt}}} \tau(\theta_s) ds + \frac{1}{N_n} \mathcal{A}_F(\gamma_{\text{opt}}) = \mathcal{G}_1^n(\ell, \theta) + o(1).$$

We can now plug this into Eq. (4.30) to obtain

$$\sum_{\gamma \in \mathcal{P}_Q(A,B) \cap \mathcal{C}} \exp(-\mathcal{E}_\beta^*(\gamma) + \mathfrak{I}_Q(\gamma) + \frac{1}{N_n} \mathcal{A}_F(\gamma)) \geq \exp(\mathcal{G}_1^n(\ell, \theta) + O((\log L)^{3a/2+2} + (\log L)^2)), \quad (4.31)$$

which together with Eqs. (4.24) and (4.26) provides the final lower bound on the denominator.

By combining the lower bound on the denominator in Eq. (4.31) with the upper bound on the numerator in Eq. (4.23), we obtain that for the choice of $a > 4$ (so that the $(\log L)^{2a}$ term in Eq. (4.23) dominates), the fraction in Eq. (4.15) is upper bounded by

$$\frac{Z_{Q,Q;F,1}^n(A, B \mid \mathfrak{U}, \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\sqcup)}{Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}_1, \mathcal{G}_2, \mathcal{G}^\sqcup)} \leq \exp(-c(\log L)^{2a}) \quad (4.32)$$

for some $c > 0$. ■

4.2. Level-lines contain translated Wulff shapes. In this subsection, we will use Theorem 4.8 to prove Theorem 4.4. The strategy will be to begin with a small Wulff shape, and use Theorem 4.8 to “grow” it to a large Wulff shape (one argues that if a small Wulff shape is present, then w.h.p. a larger one is also present). The smallest starting Wulff shape we need has diameter $\leq \frac{1}{2}$ (see Remark 4.3). In order to jump-start this growth process, we use the following consequence of the analysis of [35].

Claim 4.17. *For every $0 < \varepsilon < \frac{1}{10}$ there exists β_0 such that, for every $\beta > \beta_0$, if V is a domain such that $\Lambda_{(1-\varepsilon)L} \subset V \subset \Lambda_L$, then w.h.p., $\phi \sim \pi_V^0$ satisfies the event $\mathcal{E}_0(\Lambda_{(1-2\varepsilon)L})$.*

Indeed, in [35, Sec 4.4] the authors argue (as part of the event denoted there \mathcal{R}_{H-1}) that $\phi \sim \pi_{\Lambda_L}^0$ contains a $*$ -adjacent circuit C of sites at height at least $H-1$ encapsulating $\Lambda_{L-(H-1)\ell}$ where $\ell \leq L \exp[-c_0(\beta \log L / \log \log L)^{1/2}]$. This is obtained inductively, establishing such a circuit C_j for heights $j = 0, \dots, H-1$, each time applying [35, Prop 4.5] (the main growth tool in that work) to every box Λ_ℓ at distance at least $\log^2 \ell$ from the previously found C_{j-1} ; that tool shows the

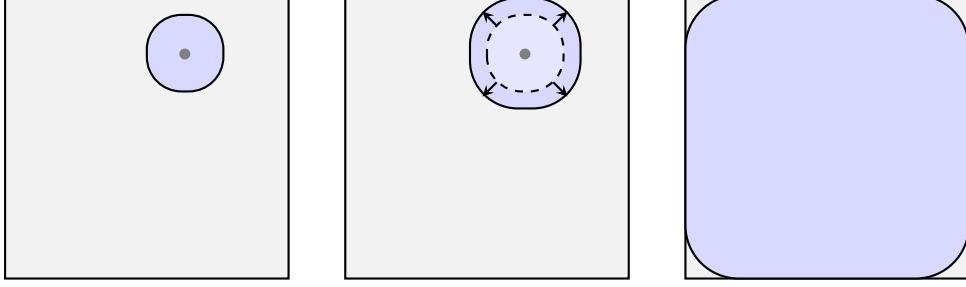


FIGURE 6. The growth procedure used to prove Proposition 4.19. Left: Fix x such that the Wulff shape $\mathcal{W}^n(x, \ell_n)$ (colored blue) is contained in $L\ell_n \mathcal{W}_1(\tau_{\beta, n})$. We start with the event $\mathcal{E}_n(\mathcal{W}^n(x, \ell))$ that \mathfrak{L}_n encapsulates $\mathcal{W}^n(x, \ell)$. Middle: By Lemma 4.20, we can grow this Wulff shape until it reaches a $o(L)$ distance from the boundary. Right: Repeat with all possible starting x from the left picture to obtain \mathcal{D} . We can then start again from the left, now allowing x such that $\mathcal{W}^n(x, \ell_n)$ is contained in \mathcal{D} , and iterate this process.

existence of a circuit at height at least j and distance at most $\varepsilon\ell$ from $\partial\Lambda_\ell$, and ‘‘stitching’’ these circuits together gives C_j . For the final \mathfrak{L}_0 , one chooses ℓ' that may be of order L (still $\ell' \geq L^{1-o(1)}$ yet it can be as large as $\frac{4}{5}L$, hence the upper bound on ε in our assumptions), and proceeds as before, yielding a circuit C_H encapsulating $\Lambda_{(1-\varepsilon)L}$. As this argument is based on applications of the growth tool on small boxes Λ_ℓ at distance at least $\log^2 \ell$ from C_{j-1} (and in the first iteration ∂V), all that is needed is to have $V \supset \Lambda_{(1-\varepsilon)L}$. Thus, in our setting we obtain C_j , for each $j \geq H-1$, encapsulating $\Lambda_{(1-\varepsilon)L-j\ell}$, and finally C_H that encapsulates $\Lambda_{(1-2\varepsilon)L}$.

We will also need the following deterministic fact about the Wulff shape (see e.g., [14, Lemma 3.9] for the proof).

Lemma 4.18. *Fix $\theta \in [0, \pi/4]$ and any τ satisfying the properties of Proposition 3.12. Let \overline{AB} be a line segment of length d and angle θ such that A, B are on $\partial\mathcal{W}_1(\tau)$. Recall that $w_1(\tau)$ is the value of the Wulff functional on $\partial\mathcal{W}_1(\tau)$. Let $\Delta(d, \theta)$ be the vertical distance between the midpoint $X = \frac{A+B}{2}$ and $\partial\mathcal{W}_1$. Then we have that for $d \ll 1$,*

$$\Delta(d, \theta) = \frac{w_1(\tau)d^2(1 + O(d^2))}{16(\tau(\theta) + \tau''(\theta))\cos(\theta)}.$$

The next proposition is the key growth result that shows we can go from the Wulff shape to its translates, to be proven over Lemmas 4.20 to 4.22. (See Fig. 6 for a visualization of the proof.) Fix $1 \leq m \leq \log L$. Recall the definitions of ℓ_n and $\kappa_{n,b}$ stated above Theorem 4.4. As a result of Remark 4.9, we will take $b = 15$ in what follows, and let $\hat{\ell}_{n,m} := \ell_n/(1 - m\kappa_{n,15})$.

Proposition 4.19. *Let n be any fixed integer. Let $1 \leq m \leq \log L$. Let Λ' be any region containing $L\mathcal{L}^n(\hat{\ell}_{n,m}, -(m-1)\kappa_{n,15})$. Under $\pi_{\Lambda'}^{H-n-1}$, given that w.h.p. $\mathcal{E}_n(L\ell_n \mathcal{W}_1(\tau_{\beta, n}))$ holds, then w.h.p. we have $\mathcal{E}_n(L\mathcal{L}^n(\hat{\ell}_{n,m+1}, -m\kappa_{n,15}))$.*

For shorthand, define $\mathfrak{R} := L\mathcal{L}^n(\hat{\ell}_{n,m}, -(m-1)\kappa_{n,15})$. Let $\mathcal{W}^n(x, \ell)$ denote the scaled Wulff shape $L\ell \mathcal{W}_1(\tau_{\beta, n})$ centered at x . Let ℓ_x be the largest value before $\mathcal{W}^n(x, \ell)$ gets within distance $N_n^{1/3}(\log L)^2$ from \mathfrak{R}^c .

Lemma 4.20. *Fix $x \in \mathfrak{R}$ and ℓ such that $\ell_n \leq \ell < \ell_x$. Under $\pi_{\Lambda'}^{H-n-1}$, given that w.h.p. $\mathcal{E}_n(\mathcal{W}^n(x, \ell))$ holds, then w.h.p. we have $\mathcal{E}_n(\mathcal{W}^n(x, \ell_x))$.*

Proof. We show that if $\mathcal{E}_n(\mathcal{W}^n(x, \ell))$ holds, then w.h.p. so does $\mathcal{E}_n(\mathcal{W}^n(x, \ell(1 + L^{-3/4})))$. Pick an angle $\theta \in [0, \pi/4]$. Let \overline{AB} be the line segment such that its horizontal distance is $N_n^{2/3}(\log L)^5$, \overline{AB} has angle θ , and A, B lie on the bottom right quarter of $\mathcal{W}^n(x, \ell)$. Assume that B is to the

right of A . Let R be the rectangle with width $N_n^{2/3}(\log L)^5$ and height $2N_n^{2/3}(\log L)^5$ such that A, B are on the sides of R , and the distance from A to the bottom of R is $N_n^{1/3}(\log L)^{15}$. Let $Z = (Z_1, Z_2)$ be the midpoint of A and B . By Lemma 4.7 and the fact that Theorem 4.8 holds for any Q satisfying the conditions in Lemma 4.7, we have that w.h.p. the interior of \mathfrak{L}_n contains the point $W = (Z_1, W_2)$ where

$$W_2 = Z_2 - \frac{N_n^{1/3}(\log L)^{10}}{8(\tau(\theta) + \tau''(\theta))\cos(\theta)^3} + c(\beta, \theta)N_n^{1/3}(\log L)^{15/2}.$$

By rescaling the result of Lemma 4.18, we obtain that Z lies a vertical distance of Δ' above $\partial\mathcal{W}^n(x, \ell)$, where

$$\Delta' = \ell L \Delta \left(\frac{N_n^{2/3}(\log L)^5}{\ell L \cos \theta}, \theta \right) = \frac{w_1(\tau_{\beta, n}) N_n^{4/3}(\log L)^{10}}{16(\tau(\theta) + \tau''(\theta))\ell L \cos(\theta)^3} + o(1).$$

Hence, as long as, say,

$$1 - \frac{w_1(\tau_{\beta, n}) N_n}{2\ell L} \geq (\log L)^{-1/4}, \quad (4.33)$$

then W is at a height at least $O(N_n^{1/3}(\log L)^{10-1/4})$ below $\partial\mathcal{W}^n(x, \ell)$, whereas the enlarged Wulff shape $\mathcal{W}^n(x, \ell(1 + L^{-3/4}))$ is only distance $O(L^{1/4})$ below $\partial\mathcal{W}^n(x, \ell)$. It is easy to check that Eq. (4.33) is satisfied by the assumption $\ell \geq \ell_n$. By repeating this $O(L)$ times to cover all angles $\theta \in [0, \pi/4]$ and (using symmetry to conclude for other angles), we obtain $\mathcal{E}_n(\mathcal{W}^n(x, \ell(1 + L^{-3/4})))$. We can then repeat this at most $o(L^{3/4} \log L)$ times to obtain $\mathcal{E}_n(\mathcal{W}^n(x, \ell_x))$. ■

Now we can define (deterministic) sets \mathcal{D}_k^n as part of the enlargement/translation procedure. Fix ℓ such that $\ell_n \leq \ell < \ell_x$. We start with \mathcal{D}_0^n defined as the largest rescaling of $\mathcal{W}_1(\tau_{\beta, n})$ that is contained in the unit square S . Now given $\mathcal{D}_k^n(\ell)$, define $\mathcal{D}_{k+1}^n(\ell)$ as follows. For any $\zeta \in \mathcal{D}_k^n(\ell)$, define

$$t_\zeta = \max\{t \geq 1 : t\ell\mathcal{W}_1(\tau_{\beta, n}) + \zeta \subset S\},$$

with $t_\zeta = 0$ if there is no such t . Then, define

$$\mathcal{D}_{k+1}^n(\ell) = \bigcup_{\zeta \in \mathcal{D}_k^n} \{t_\zeta \ell \mathcal{W}_1(\tau_{\beta, n}) + \zeta\}.$$

Note that by construction, $\mathcal{D}_k^n(\ell) \subset \mathcal{D}_{k+1}^n(\ell) \subset S$ for each $k \geq 0$, and there is no dependence of \mathcal{D}_0^n on ℓ . The following lemma showing convergence of the $\mathcal{D}_k^n(\ell)$ to $\mathcal{L}^n(\ell, 0)$ was proved in [14].

Lemma 4.21 ([14, Lemma 6.8]). *The sequence $\mathcal{D}_k^n(\ell)$ converges to $\mathcal{D}_\infty^n := \mathcal{L}^n(\ell, 0)$, and the Hausdorff distance between \mathcal{D}_k^n and \mathcal{D}_∞^n is bounded by c^k for some constant $c \in (0, 1)$.*

We can now apply Lemma 4.20 in the iterative procedure of the \mathcal{D}_k^n sets.

Lemma 4.22. *Let x_0 be the center of \mathfrak{R} . Under $\pi_{\Lambda'}^{H-n-1}$, given that w.h.p. $\mathcal{E}_n(\mathcal{W}^n(x_0, \ell_{x_0}))$ holds, then w.h.p. we have $\mathcal{E}_n(L\mathcal{L}^n(\hat{\ell}_{n,m}, -m\kappa_{n,15}))$.*

Proof. By definition, we have $L(1 - (m-1)\kappa_{n,15})\mathcal{D}_0^n \subset \mathfrak{R}$. Hence, by adding another factor of $\kappa_{n,15}$, we get $L(1 - m\kappa_{n,15})\mathcal{D}_0^n \subset \mathcal{W}^n(x_0, \ell_{x_0})$. We show that with high probability,

$$\mathcal{E}_n(L(1 - m\kappa_{n,15})\mathcal{D}_k^n(\frac{\ell_n}{1-m\kappa_{n,15}})) \implies \mathcal{E}_n(L(1 - m\kappa_{n,15})\mathcal{D}_{k+1}^n(\frac{\ell_n}{1-m\kappa_{n,15}})). \quad (4.34)$$

Indeed, for any $\mathcal{W}^n(x, \ell_n) \subset L(1 - m\kappa_{n,15})\mathcal{D}_k^n(\frac{\ell_n}{1-m\kappa_{n,15}})$, we can apply Lemma 4.20 to conclude that

$$\mathcal{E}_n(\mathcal{W}^n(x, \ell_n)) \implies \mathcal{E}_n(\mathcal{W}^n(x, \ell_x)).$$

Suppose that we could grow $\mathcal{W}^n(x, \ell)$ to $\mathcal{W}^n(x, \tilde{\ell}_x)$, defined as the largest scaling of $\mathcal{W}^n(x, \ell)$ contained in $\Lambda_{L(1-m\kappa_{n,15})}$ as long as $\tilde{\ell}_x \geq \ell$. Then we could apply the above display to all x such

that $\mathcal{W}^n(x, \ell_n) \subset L(1 - m\kappa_{n,15})\mathcal{D}_k^n(\frac{\ell_n}{1-m\kappa_{n,15}})$ to obtain Eq. (4.34). However, it is a-priori possible that $\tilde{\ell}_x > \ell_x$. That is, the additional restriction that we need to consider is that by definition of ℓ_x , we can only grow each Wulff shape until we reach distance $N_n^{1/3}(\log L)^2$ from \mathfrak{R}^c . But, since $\mathcal{D}_k^n \subset \mathcal{D}_\infty^n$, the ending shape $L(1 - m\kappa_{n,15})\mathcal{D}_{k+1}^n(\frac{\ell_n}{1-m\kappa_{n,15}})$ is already at a distance $N_n^{1/3}(\log L)^2$ away from \mathfrak{R}^c , so this additional restriction does not change anything. Note also that we only need to repeat this up to $k = O(\log L)$ times to get \mathcal{D}_k^n within distance $1/L$ of \mathcal{D}_∞^n by the bound given in Lemma 4.21, which is all that is needed due to the discretization of the lattice. As noted in Remark 4.9, we can afford this many iterations. ■

The observation that $L\mathcal{L}^n(\hat{\ell}_{n,m+1}, -m\kappa_{n,15}) \subset L\mathcal{L}^n(\hat{\ell}_{n,m}, -m\kappa_{n,15})$ together with Lemmas 4.20 and 4.22 concludes the proof of Proposition 4.19. Now we prove the main theorem of this section.

Proof of Theorem 4.4. First, additionally assume that we have

$$\mathcal{E}_{n+1}(L\mathcal{L}^n(\hat{\ell}_{n,H-n}, -(H-n-1)\kappa_{n,15})).$$

We reveal the $H-n-1$ level-line in this case. This reveals an interior region that contains $L\mathcal{L}^n(\hat{\ell}_{n,H-n}, -(H-n-1)\kappa_{n,15})$ with a boundary of sites with height $\geq H-n-1$. By monotonicity, we can drop this boundary condition down to exactly $H-n-1$. Then, we can apply Proposition 4.19 to conclude that with high probability we have $\mathcal{E}_n(L\mathcal{L}^n(\hat{\ell}_{n,H-n}, -(H-n)\kappa_{n,15}))$, which implies the theorem as $\hat{\ell}_{n,H-n} \leq \ell_n(1+L^{-1/2})$.

It thus suffices to justify the above assumption. We again use monotonicity, this time dropping the floor to $-(H-n-1)$. This allows us to control \mathfrak{L}_1 , the 1-level-line with a floor at 0 by the 1-level-line with a floor at height $-(H-n-1)$, which is then the same as the $(H-n)$ level-line when the floor is at 0 and boundary conditions are at height $H-n-1$. By Proposition 4.19 (whose application is justified by Claim 4.17 and Remark 4.3) with $m=1$, this level-line (and hence \mathfrak{L}_1) contains $L\mathcal{L}^n(\hat{\ell}_{n,2}, -\kappa_{n,15})$. On this event, we can reveal \mathfrak{L}_1 , and by monotonicity drop the induced ≥ 1 boundary conditions to be equal to 1. Then we can repeat, studying the k level-line, \mathfrak{L}_k , for $2 \leq k \leq H-n$ by dropping the floor to $-(H-n-k)$, and concluding that \mathfrak{L}_k contains $L\mathcal{L}^n(\hat{\ell}_{n,k+1}, -k\kappa_{n,15})$. ■

5. UPPER BOUND

Our goal in this section is the following bound on the distance of \mathfrak{L}_n from the bottom side of Λ_L .

Theorem 5.1. *In the setting of Theorem 1.1, fix $n \geq 0$ and $K > 0$, let $\rho_n(x)$ be the maximum vertical distance of \mathfrak{L}_n above $x + (\frac{L}{2}, 0)$ for $-N_n^{2/3} \leq x \leq N_n^{2/3}$, and set $\sigma_n^2 > 0$ as per Definition 3.20. Then every weak limit point $\mathbf{Y}_n(t)$ of the process $Y_n(t) := N_n^{-1/3} \rho_n(tN_n^{2/3})$ ($t \in [-K, K]$), as $L \rightarrow \infty$, satisfies*

$$\mathbf{Y}_n \preceq \text{FS}_{\sigma_n}.$$

Moreover, for every fixed m , every weak limit point $(\mathbf{Y}_n(t))_{n \leq m}$ of the processes $(Y_n(t))_{n \leq m}$ satisfies

$$(\mathbf{Y}_n)_{n \leq m} \preceq \bigotimes_{n \leq m} \text{FS}_{\sigma_n}.$$

Observe that $\rho_n(x)$ is a decreasing function of ϕ . Hence, en route to proving Theorem 5.1 in this section we may apply monotonicity arguments that are decreasing (e.g., decreasing individual floor constraints or heights of boundary vertices; see Claim A.1).

Let R be the rectangle of size $N_n^{2/3}(\log L)^{25} \times 2N_n^{2/3}(\log L)^{25}$ centered at the $x = 0$ such that the bottom side of R is at distance $2N_n^{1/3} \exp(-c\beta\sqrt{\log L/\log\log L})$ above the bottom side of Λ_L , for c as in Theorem 4.4. Let A, B be on the sides of ∂R such that the distance from A, B to the bottom of R is $N_n^{1/3}(\log L)^{16}$. The dimensions of R and the location of A, B satisfy the

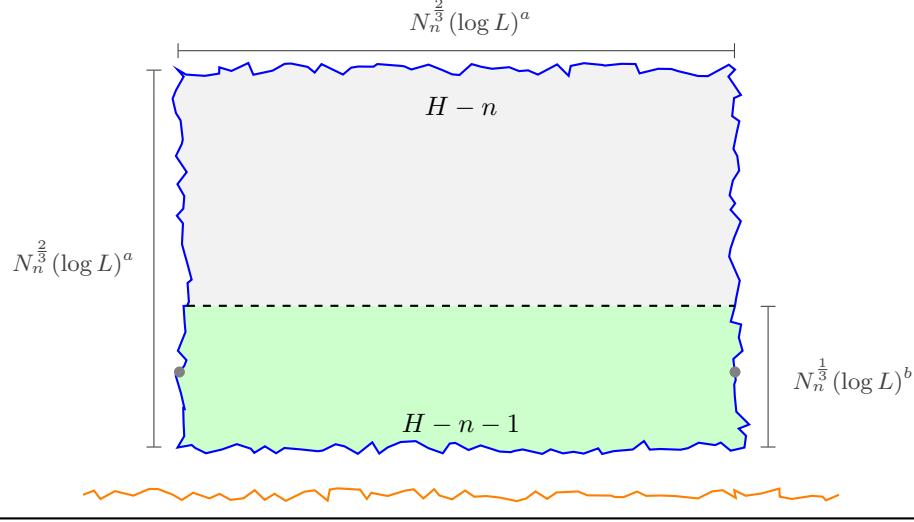


FIGURE 7. An instantiation of the domain Q used in the proof of the upper bound on \mathfrak{L}_n . The orange line is the lower level-line \mathfrak{L}_{n+1} , and the floor (the constraint that $\phi_x \geq 0$) is only present in the green region. The two gray points on the boundary mark where the boundary conditions change from $H - n - 1$ to $H - n$.

requirements of Lemma 4.7. So, let Q, ξ be any domain and boundary condition satisfying the properties of Lemma 4.7 with respect to this R, A, B (see Fig. 7). By Theorem 4.4 and Lemma 4.7, for the stochastic domination of a single \mathbf{Y}_n , it suffices to prove Theorem 5.1 where \mathfrak{L}_n is the $H - n$ level-line in π_Q^ξ .

Next, by FKG, we can first remove the floor in some parts of Q , and condition on decreasing events. Let F be the portion of Q below the line $y = 2N_n^{1/3}(\log L)^{16}$. Let \mathcal{G}^\sqcup be the event that \mathfrak{L}_n stays above the horizontal line \mathcal{H} distance $2(\log L)^2$ above the highest point on the bottom side of Q . Then, $\pi_Q^\xi \succeq \pi_{Q;F}^\xi(\cdot \mid \mathcal{G}^\sqcup)$. Note that the distance between \mathcal{H} and the bottom of Λ_L is $o(N_n^{1/3})$, and goes to 0 under the rescaling of Y_n . For simplicity, in the rest of this section we will refer to the y -coordinate of \mathcal{H} as 0, and use ‘‘height’’ to denote the vertical distance of a point above \mathcal{H} (as opposed to referencing ϕ as a height function).

From here, we can follow the same preliminary steps done in the proof of Theorem 4.8 to move to a polymer model. By Lemma 4.13, we can condition further on the event \mathcal{G}_1 defined there, which says that we have the cluster expansion form of the measure given by Proposition 2.3. That is, (recalling Eq. (3.4) and the notation in Eq. (4.2)) we have

$$\pi_{Q;F}^\xi(\gamma \mid \mathcal{G}^\sqcup, \mathcal{G}_1) \propto \exp \left(-\mathcal{E}_\beta^*(\gamma) + \frac{\mathcal{A}_F(\gamma)}{N_n} + \mathfrak{I}_Q(\gamma) \right) = q_{Q;F,1}^n(\gamma). \quad (5.1)$$

Hence, by Eq. (4.9) we can now move to the polymer model with weights $q_{Q;F,1}^n(\gamma)$ and partition function

$$Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}^\sqcup) = \sum_{\gamma \in \mathcal{P}_Q(A, B) \cap \mathcal{G}^\sqcup} q_{Q;F,1}^n(\gamma).$$

Recalling the animal decomposition from Definition 3.3 and the notational comment following it, we extend the above definitions to weights and sums over animals Γ . Now define W as ‘‘ Q with no top or bottom’’, i.e. the extension of Q to an infinite vertical strip containing it at the top and bottom sides of Q . (Both the choice of the strip and exact point of extension are irrelevant.) By the event \mathcal{G}^\sqcup and the construction of Q , we first argue that except for a bad set of γ making up a

$o(1)$ fraction of $Z_{Q,Q;F,1}^n(\cdot \mid \mathcal{G}^\perp)$ and of $Z_{Q,W;F,1}^n(\cdot \mid \mathcal{G}^\perp)$, we have

$$q_{Q;F,1}^n(\gamma) = (1 + o(1))q_{W;F,1}^n(\gamma),$$

and we can instead study the polymer model with weights $q_{W;F,1}^n$ and partition function $Z_{Q,W;F,1}^n(\cdot \mid \mathcal{G}^\perp)$. By the decay properties of \mathfrak{I} , it suffices to show that γ stays distance at least $(\log L)^2$ away from the top and bottom of Q , besides an exceptional set of γ . If γ approaches the top of Q then it must have length greater than $1.1N_n^{1/3}(\log L)^{16}$, which can be ruled out via a Peierls argument mapping to a minimal length path from A to B . In particular, the distance from γ to the top of Q is at least $(\log L)^2$. Regarding interactions with the bottom of Q , in the proof of Lemma 4.14 (see Eq. (4.10)) it is shown that

$$Z_{Q,Q;F,1}^n(A, B \mid \max_{b \in \gamma} |\partial D_b| \leq (\log L)^2, \mathcal{G}^\perp) \geq (1 - o(1))Z_{Q,Q;F,1}^n(A, B \mid \mathcal{G}^\perp),$$

and the same proof also holds when replacing the interaction domain with W . For γ in the left side of the inequality, the distance between γ and the bottom side of Q is at least $(\log L)^2$.

Random walk preliminaries and notation. Let $\tilde{\mathbf{P}}$ denote the polymer measure on Γ with weight $q_{W;F,1}^n$ and partition function $Z_{Q,W;F,1}^n(A, B)$, so our starting point is $\tilde{\mathbf{P}}(\cdot \mid \mathcal{G}^\perp)$. First observe that in Eq. (5.1), we could have replaced $\mathcal{A}_F(\gamma)$ with $-|D_1 \cap F|$, the area below γ intersected with F , by just a renormalization which does not change the measure $\tilde{\mathbf{P}}$ (in fact we only started with $\mathcal{A}_F(\gamma)$ to keep the same notation as Section 4). Now define $\mathcal{A}_F^\downarrow(\Gamma)$ as the area of the region below the linear interpolation of the cone-points of Γ intersected with F . Similarly for an irreducible component $\Gamma^{(i)}$ between cone-points u, v , we define $\mathcal{A}_F^\downarrow(\Gamma^{(i)})$ as the area below \overline{uv} intersected with F . In Lemma 5.2 and Remark 5.3, we show that except for a bad set of Γ occurring with $\tilde{\mathbf{P}}(\cdot \mid \mathcal{G}^\perp)$ -probability $o(1)$, we can replace $|D_1 \cap F|$ with $|\mathcal{A}_F^\downarrow(\Gamma)|$ at the multiplicative cost of $(1 + o(1))$ to the weight of Γ . We will make this change and call the new measures \mathbf{P} and $\mathbf{P}(\cdot \mid \mathcal{G}^\perp)$.

Now observe that every diamond in Γ is fully contained in W , and hence $\mathfrak{I}_W(\Gamma^{(i)}) = \mathfrak{I}_{\mathbb{Z}^2}(\Gamma^{(i)})$ for irreducible components $\Gamma^{(i)}$. For \mathbf{P} , recall from Proposition 3.19 that, after renormalizing the weights by $e^{\mathfrak{h}_{(1,0)}^n \cdot (B-A)}$, every irreducible component $\Gamma^{(i)}$ in Γ has weight given by the probability $\mathbb{P}^{\mathfrak{h}_{(1,0)}^n}(\Gamma^{(i)})e^{-\frac{1}{N_n}\mathcal{A}_F^\downarrow(\Gamma^{(i)})}$. By the product structure of the weights (Eq. (3.7)), this implies that the law of Γ in between any two cone-points u, v is an area tilted random walk bridge with increments taking values in the set of irreducible components, independent of the portion of Γ before u and after v . This random walk measure will be denoted by $\mathbf{P}_{u,v}(\Gamma) := \mathbf{P}(\Gamma \mid u, v \in \text{Cpts}(\Gamma))$, where it is implicit that the measure is on the restriction of Γ in between u and v . It will be useful also to define the one ended random walk measure \mathbf{P}_u that starts at u , so that $\mathbf{P}_{u,v} = \mathbf{P}_u(\cdot \mid v \in \text{Cpts}(\Gamma))$. We can also define the above notation for the case when there is no area term, replacing \mathbf{P} by $\hat{\mathbf{P}}$. Finally, the notation for the expectation with respect to the above measures will be to replace \mathbf{P} with \mathbf{E} .

Proof strategy. The strategy for proving Theorem 5.1 is as follows:

Step 1: We show in Lemma 5.2 that w.h.p. we have the event Bdd , informally stating that the first and last cone-points of Γ are not far away from A and B , and the maximum distance between any adjacent cone-points of Γ is suitably small. This moreover allows us to modify the area tilt so that it only depends on the cone-points of Γ , and not the full path. Calling the first and last cone-points A', B' , we move to the measure $\mathbf{P}_{A',B'}(\cdot \mid \mathcal{G}^\perp, \text{Bdd})$ for the next step.

Step 2: We next show in Lemma 5.4 that w.h.p. in the first and last quadrant of the interval $[A'_1, B'_1]$, we can find points with height at most $1000N_n^{1/3}$. Denoting these points \bar{u}, \bar{v} , we proceed to the next step with the measure $\mathbf{P}_{u,v}(\cdot \mid \mathcal{G}^\perp, \text{Bdd})$.

Step 3: Moving inwards, we show in Lemma 5.8 that for any $\varepsilon > 0$, with probability $1 - \varepsilon$, the cone-points with x -coordinates $-(\log \log L)N_n^{2/3}$ and $(\log \log L)N_n^{2/3}$, call them \bar{u}', \bar{v}' , are at height at most $K_\varepsilon N_n^{1/3}$. Calling these points \bar{u}', \bar{v}' , we move to the next step with the measure $\mathbf{P}_{u',v'}(\cdot \mid \mathcal{G}^\perp, \text{Bdd})$.

Step 4: In Lemma 5.12 we show that with high probability, we have the event **Repel** that $\text{Cpts}(\Gamma)$ stays above N^δ in the interval J that is $N^{8\delta}$ away from the endpoints \bar{u}' and \bar{v}' . Call the first and last cone-points in J by \bar{u}'', \bar{v}'' and consider the measure $\mathbf{P}_{\bar{u}'',\bar{v}''}(\cdot \mid \mathcal{G}^\perp, \text{Bdd}, \text{Repel})$. Since $\text{Bdd} \cap \text{Repel}$ implies \mathcal{G}^\perp , the conditioning on \mathcal{G}^\perp can be removed. The resulting measure $\mathbf{P}_{\bar{u}'',\bar{v}''}(\cdot \mid \text{Bdd}, \text{Repel})$ is such that the marginal on $\text{Cpts}(\Gamma)$ is a 2D area-tilted random walk, as (recalling also Remark 5.3) we have removed everything not measurable with respect to $\text{Cpts}(\Gamma)$. The convergence of such a random walk to the Ferrari–Spohn diffusion was proven in [25].

Reducing to a 2D RW with an area-tilt. We begin with establishing the boundedness of irreducible components.

Lemma 5.2. *Let Bdd be the set of Γ such that $|\text{Cpts}(\Gamma)| \geq 2$, and the maximum displacement $\max_i(|X(\Gamma^{(i)})|, |X(\Gamma^{(L)})|, |X(\Gamma^{(R)})|)$ is no more than $(\log L)^{50}$. Then,*

$$\tilde{\mathbf{P}}(\text{Bdd}^c \mid \mathcal{G}^\perp) \leq e^{-c(\log L)^{50}},$$

or equivalently,

$$Z_{Q,W;F,1}^n(\text{Bdd}^c \mid \mathcal{G}^\perp) \leq e^{-c(\log L)^{50}} Z_{Q,W;F,1}^n(\mathcal{G}^\perp).$$

Proof. Since the area tilt can change the weight of events by at most a factor of $e^{(\log L)^{41}}$, it suffices to prove the bounds with respect to $\tilde{Z}_{Q,W}^n$ and weights \tilde{q}_W^n (i.e. having removed the area term). By Lemma 3.11 we can assume we have cone-points. Now by construction of Q , all diamonds of Γ are entirely contained in W , whence $\tilde{q}_W^n(\Gamma^{(i)}) = \tilde{q}_{\mathbb{Z}^2}^n(\Gamma^{(i)})$. Ignoring the event \mathcal{G}^\perp by an upper bound, we can thus write

$$\begin{aligned} \tilde{Z}_{Q,W}^n(A, B \mid |X(\Gamma^{(L)})| > (\log L)^{50}, |\text{Cpts}(\Gamma)| \geq 2, \mathcal{G}^\perp) \\ &\leq \sum_{\substack{u \in \mathcal{Y}^\blacktriangle(A) \\ v \in \mathcal{Y}^\blacktriangleright(B) \\ |A-u| > (\log L)^{50}}} \sum_{\Gamma^{(L)} \in \mathcal{A}_L \cap \mathcal{P}_Q(A, u)} \tilde{q}_W^n(\Gamma^{(L)}) \sum_{\Gamma^{(R)} \in \mathcal{A}_R \cap \mathcal{P}_Q(v, B)} \tilde{q}_W^n(\Gamma^{(R)}) \sum_{\substack{k \geq 1 \\ \Gamma^{(1)} \circ \dots \circ \Gamma^{(k)} \in \mathcal{P}_{\mathbb{Z}^2}(u, v)}} \prod_{i=1}^k \tilde{q}_{\mathbb{Z}^2}^n(\Gamma^{(i)}) \\ &= \sum_{\substack{u \in \mathcal{Y}^\blacktriangle(A) \\ v \in \mathcal{Y}^\blacktriangleright(B) \\ |A-u| > (\log L)^{50}}} \sum_{\Gamma^{(L)} \in \mathcal{A}_L \cap \mathcal{P}_Q(A, u)} \tilde{q}_W^n(\Gamma^{(L)}) \sum_{\Gamma^{(R)} \in \mathcal{A}_R \cap \mathcal{P}_Q(v, B)} \tilde{q}_W^n(\Gamma^{(R)}) \tilde{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(u, v). \end{aligned}$$

Fixing u , every $\Gamma^{(L)} \in \mathcal{P}_Q(A, u)$ has length at least $|A - u|$, so by Proposition 3.19 we have

$$\sum_{\Gamma^{(L)} \in \mathcal{A}_L \cap \mathcal{P}_Q(A, u)} \tilde{q}_W^n(\Gamma^{(L)}) \leq C e^{-\nu_g \beta |A - u|} e^{-\tau_{\beta, n}(A - u)},$$

with an analogous bound for $\Gamma^{(R)}$. Thus, using that $\tilde{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(u, v) \leq e^{-\tau_{\beta, n}(v - u)}$ and the convexity of the surface tension (Item (iii) in Proposition 3.12), we obtain an upper bound of

$$\begin{aligned} \tilde{Z}_{Q,W}^n(A, B \mid |X(\Gamma^{(L)})| > (\log L)^{50}, |\text{Cpts}(\Gamma)| \geq 2, \mathcal{G}^\perp) \\ &\leq C e^{-\tau_{\beta, n}(B - A)} \sum_{\substack{u \in \mathcal{Y}^\blacktriangle(A) \\ |A-u| > (\log L)^{50}}} e^{-\nu_g \beta |A - u|} \sum_{v \in \mathcal{Y}^\blacktriangleright(B)} e^{-\nu_g \beta |B - v|}. \end{aligned}$$

The number of points u such that $|A - u| = r$ is $O(r)$, so that the sum over u is bounded by $e^{-\frac{\nu_g}{2}\beta(\log L)^{50}}$. By the same logic, the sum over v is bounded by a constant. Hence, we have

$$\begin{aligned}\widehat{Z}_{Q,W}^n(A, B \mid |X(\Gamma^{(L)})| > (\log L)^{50}, |\mathsf{Cpts}(\Gamma)| \geq 2, \mathcal{G}^\perp) &\leq Ce^{-\frac{\nu_g}{2}\beta(\log L)^{50}}e^{-\tau_{\beta,n}(B-A)} \\ &\leq Ce^{-c\beta(\log L)^{50}}\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A, B) \\ &\leq Ce^{-c\beta(\log L)^{50}}\widehat{Z}_{Q,W}^n(A, B),\end{aligned}$$

where we used Item (ii) of Proposition 3.15 in the second line and Lemmas 3.16 and 3.17 in the third.

The proof that the irreducible components $|X(\Gamma^{(i)})|$ have the same bound follows similarly, so the details are omitted. The lemma now concludes by a union bound. \blacksquare

Remark 5.3. Given Γ , let S be the linear interpolation of $\mathsf{Cpts}(\Gamma)$. On the event Bdd , the area in between Γ and S between two cone-points of Γ is at most $O(\log L)^{50}$. There are (deterministically) at most $N_n^{2/3}(\log L)^{25}$ cone-points, so the total difference in area between Γ and S is at most $O(N_n^{2/3}(\log L)^{75})$. After dividing this area by N_n , this is $o(1)$. Hence, as commented in the preliminaries of this section, we can change the area tilt term in Eq. (5.1) to $-\mathcal{A}_F^\downarrow(\Gamma)/N_n$ at the cost of a multiplicative $(1 + o(1))$.

Next, we show that Γ will drop to height $O(N_n^{1/3})$. For an easier read of the proof, we leave the dimensions in terms of a, b but our application is for $a = 25$ and $b = 16$.

Lemma 5.4. *Fix $a, b > 0$ such that $a > 3b/2$, let $I := [-\frac{1}{2}N_n^{2/3}(\log L)^a, -\frac{1}{2}N_n^{2/3}(\log L)^a]$, and let \mathbf{u}, \mathbf{v} be such that $\mathbf{u}_1 < \mathbf{v}_1$ are within an additive $(\log L)^{50}$ -term of the two endpoints of I , and $\mathbf{u}_2, \mathbf{v}_2 \in N_n^{1/3}(\log L)^b \pm (\log L)^{50}$. Let Drop be the event that there exist cone-points of Γ in first and last quarters of I which lie below height $KN_n^{1/3}$. Then for some constant $K > 0$,*

$$\mathbf{P}_{\mathbf{u}, \mathbf{v}}(\mathsf{Drop} \mid \mathcal{G}^\perp, \mathsf{Bdd}) \geq 1 - e^{-c(\log L)^a}.$$

Proof. Roughly, the sought bound is achieved by forcing the 2D random walks to proceed as follows:

- In the two intervals of length $N_n^{2/3}/(\log L)^b$ closest to \mathbf{u}, \mathbf{v} , drop by height $N_n^{1/3}/(\log L)^{b/2}$. This is repeated up to $(\log L)^{3b/2}$ times from each side so the height drops to at most $\frac{K}{2}N_n^{1/3}$.
- Not exceed height $KN_n^{1/3}$ throughout the remaining middle interval.

We will argue that the probability cost associated to these two items in a random walk with no area-tilt is $\exp[-c(\log L)^{3b/2}]$ and $\exp[-c(\log L)^a]$ respectively, whereas the associated area of such Γ has $\mathcal{A}_F^\downarrow(\Gamma) = O((\log L)^{3b/2} + (\log L)^a)N_n$, giving a term of the same order as above in the probability.

By a union bound, we may as well assume that Drop only refers to the existence of such a cone-point in the first quarter of I . Isolating the area term, we have

$$\mathbf{P}_{\mathbf{u}, \mathbf{v}}(\mathsf{Drop}^c \mid \mathcal{G}^\perp, \mathsf{Bdd}) = \frac{\hat{\mathbf{E}}_{\mathbf{u}, \mathbf{v}}[e^{-\frac{1}{N_n}\mathcal{A}_F^\downarrow(\Gamma)}\mathbb{1}_{\mathsf{Drop}^c} \mid \mathcal{G}^\perp, \mathsf{Bdd}]}{\hat{\mathbf{E}}_{\mathbf{u}, \mathbf{v}}[e^{-\frac{1}{N_n}\mathcal{A}_F^\downarrow(\Gamma)} \mid \mathcal{G}^\perp, \mathsf{Bdd}]}.$$

On the event $\mathsf{Drop}^c \cap \mathsf{Bdd}$, we know that Γ lies above height $KN_n^{1/3} - (\log L)^{50} =: K'N_n^{1/3}$ over an interval of length $N_n^{2/3}(\log L)^a/4$. Thus, we have $\mathcal{A}_F^\downarrow(\Gamma) \geq K'(\log L)^a N_n/4$, whence

$$\hat{\mathbf{E}}_{\mathbf{u}, \mathbf{v}}[e^{-\frac{1}{N_n}\mathcal{A}_F^\downarrow(\Gamma)} \mid \mathcal{G}^\perp, \mathsf{Bdd}] \leq e^{-K'(\log L)^a/4}.$$

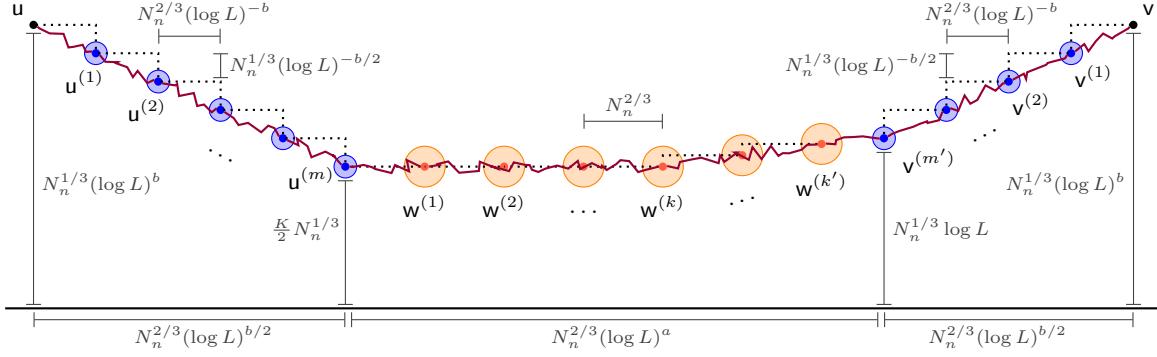


FIGURE 8. The dropping points $z^{(j)}$ and target balls B_j around them in proving Lemma 5.4. We lower bound the probability that $\Gamma : u \rightarrow v$ hits each target ball along the way. Between $w^{(k')}$ and $v^{(m')}$, the size of the balls decreases in order. It is now rare for a random walk started from the orange ball around $w^{(k')}$ to hit the much smaller blue ball around $v^{(m')}$. Controlling this probability is easier if this transition occurs away from the bottom, hence we first climb back up to height $N_n^{1/3}(\log L)$.

It suffices to now show the key lower bound, which captures the dropping behavior of Γ to the equilibrium height of $O(N_n^{1/3})$:

$$\hat{\mathbf{E}}_{u,v}[e^{-\frac{1}{N_n}\mathcal{A}_F^\downarrow(\Gamma)} \mid \mathcal{G}^\square, \text{Bdd}] \geq e^{-c^*(\log L)^a}, \quad (5.2)$$

whence we can take $K' > 8c^*$ to conclude.

To prove this lower bound, we break I into three intervals $I = I_1 \cup I_2 \cup I_3$ such that $|I_1| = |I_3| = N_n^{2/3}(\log L)^{3b/2}$. Let Low be the event that the maximum height of the cone-points of Γ in the interval I_2 is at most $KN_n^{1/3}$. Then, on $\text{Bdd} \cap \text{Low}$, the maximum height of Γ in I_2 is at most $KN_n^{1/3} + (\log L)^{50}$, in which case

$$\frac{\mathcal{A}_F^\downarrow(\Gamma)}{N_n} \leq \frac{1}{N_n} N_n^{1/3}(\log L)^b (|I_1| + |I_3|) + \frac{C}{N_n} KN_n^{1/3} |I_2| \leq (\log L)^{3b/2} + C' K (\log L)^a.$$

Hence,

$$\begin{aligned} \hat{\mathbf{E}}_{u,v}[e^{-\frac{1}{N_n}\mathcal{A}_F^\downarrow(\Gamma)} \mid \mathcal{G}^\square, \text{Bdd}] &\geq e^{-(\log L)^{3b/2} - C' K (\log L)^a} \hat{\mathbf{P}}_{u,v}(\text{Low} \mid \text{Bdd}, \mathcal{G}^\square) \\ &= e^{-(\log L)^{3b/2} - C' K (\log L)^a} \hat{\mathbf{P}}_u(\text{Low} \mid \text{Bdd}, \mathcal{G}^\square, v \in \text{Cpts}(\Gamma)) \end{aligned}$$

As long as $a > 3b/2$, the $(\log L)^a$ term dominates the prefactor, which is in the form of the desired lower bound in Eq. (5.2). It now remains to prove a lower bound of the same form for the probability. We trivially have

$$\hat{\mathbf{P}}_u(\text{Low} \mid \text{Bdd}, \mathcal{G}^\square, v \in \text{Cpts}(\Gamma)) \geq \hat{\mathbf{P}}_u(\text{Low}, \mathcal{G}^\square, v \in \text{Cpts}(\Gamma) \mid \text{Bdd}).$$

First, let \mathcal{G}_*^\square be the event that the minimum height of any cone-point of Γ is at least $(\log L)^{50}$. Then, $\text{Bdd} \cap \mathcal{G}_*^\square \subset \mathcal{G}^\square$, so that

$$\hat{\mathbf{P}}_u(\text{Low}, \mathcal{G}^\square, v \in \text{Cpts}(\Gamma) \mid \text{Bdd}) \geq \hat{\mathbf{P}}_u(\text{Low}, \mathcal{G}_*^\square, v \in \text{Cpts}(\Gamma) \mid \text{Bdd}).$$

By this move, we have reduced to a situation where all the events in question are measurable with respect to $\text{Cpts}(\Gamma)$. Hence, we can turn to the marginal on $\text{Cpts}(\Gamma)$, which has the law of a 2D random walk with increment law $\mathbb{P}(x, y) = \mathbb{P}^{h_{(1,0)}}(X(\Gamma) = (x, y) \mid |(x, y)| \leq (\log L)^{50})$. Let S denote this random walk started at u when the measure referred to is $\hat{\mathbf{P}}_u(\cdot \mid \text{Bdd})$.

Let

$$m := \lfloor (\log L)^{3b/2} - \frac{K}{2}(\log L)^{b/2} \rfloor \quad , \quad m' := \lfloor (\log L)^{3b/2} - (\log L)^{b/2+1} \rfloor \quad ,$$

and define the points

$$\begin{aligned} \mathbf{u}^{(j)} &= \left(\mathbf{u}_1 + jN_n^{2/3}(\log L)^{-b}, \mathbf{u}_2 - jN_n^{1/3}(\log L)^{-b/2} \right) \quad (j = 0, \dots, m) , \\ \mathbf{v}^{(j)} &= \left(\mathbf{v}_1 - jN_n^{2/3}(\log L)^{-b}, \mathbf{v}_2 - jN_n^{1/3}(\log L)^{-b/2} \right) \quad (j = 0, \dots, m') , \end{aligned}$$

so that, recalling $|\mathbf{u}_2 - N_n^{1/3}(\log L)^b| < (\log L)^2$, we have $|\mathbf{u}_2^{(m)} - \frac{K}{2}N_n^{1/3}| = O(N_n^{1/3}(\log L)^{-b/2})$ (accounting for integer rounding in m), and similarly $|\mathbf{v}_2^{(m')} - N_n^{1/3}\log L| = O(N_n^{1/3}(\log L)^{-b/2})$. Further set

$$k := \lfloor (\log L)^a - (m + m')(\log L)^{-b} \rfloor - \lceil \log L - \frac{K}{2} \rceil \quad , \quad k' = \lfloor (\log L)^a - (m + m')(\log L)^{-b} \rfloor - 1 \quad ,$$

and define the points

$$\mathbf{w}^{(j)} = \begin{cases} \left(\mathbf{u}_1^{(m)} + j\lfloor N_n^{2/3} \rfloor, \mathbf{u}_2^{(m)} \right) & j = 1, \dots, k , \\ \left(\mathbf{u}_1^{(m)} + j\lfloor N_n^{2/3} \rfloor, \mathbf{u}_2^{(m)} + (j - k)N_n^{1/3} \right) & j = k + 1, \dots, k' . \end{cases}$$

so that the jump from $\mathbf{w}^{(k')}$ to $\mathbf{v}^{(m')}$ is such that

$$\lfloor N_n^{2/3} \rfloor \leq \mathbf{v}_1^{(m')} - \mathbf{w}_1^{(k')} \leq 2\lfloor N_n^{2/3} \rfloor .$$

We will control the probability that in time $m + m' + k' + 1$, the random walk \mathbf{S} passes through appropriately sized balls around the points $\mathbf{u}^{(j)}, \mathbf{w}^{(j)}, \mathbf{v}^{(j)}$, as illustrated in Fig. 8. That is, \mathbf{S} will first drop down along the points $\mathbf{u}^{(j)}$ during the interval I_1 , then remain low at the points $\mathbf{w}^{(j)}$ along the interval I_2 , and finally climb back up to \mathbf{v} along the points $\mathbf{v}^{(j)}$ along the interval I_3 . We defined separately the above points as $\mathbf{u}, \mathbf{w}, \mathbf{v}$ for ease of definitions and an emphasis on the above geometric description, but it will be convenient now to just define one sequence of points $\mathbf{z}^{(j)}$ such that

$$\mathbf{z}^{(j)} = \begin{cases} \mathbf{u}^{(j)} & \text{for } 0 \leq j \leq m \\ \mathbf{w}^{(j-m)} & \text{for } m < j \leq m + k' \\ \mathbf{v}^{(m+m'+k'+1-j)} & \text{for } m + k' < j \leq m + m' + k' + 1 . \end{cases}$$

Now along this sequence, for $1 \leq j \leq m + m' + k' + 1$, define $\alpha = \mathbb{E}^{\mathbf{h}_{(1,0)}^n}[X(\Gamma)_1 \mid |X(\Gamma)| \leq (\log L)^{50}]$, the expectation in the x -coordinate of an increment. (The expectation in the y -coordinate is 0 by the symmetry of the model.) Then define $\ell_j := \frac{1}{\alpha}(\mathbf{z}_1^{(j)} - \mathbf{z}_1^{(j-1)})$, noting that with the exception of $j = m + k'$, we have $\ell_j \leq \ell_{j+1}$. Define \mathcal{B}_j as the ball of radius $\sqrt{\ell_j}$ centered about $\mathbf{z}^{(j)}$ in \mathbb{Z}^2 , except when $j = m + k'$ (the index where $\ell_j > \ell_{j+1}$), where we set \mathcal{B}_j to have radius $\sqrt{\ell_{j+1}}$. For convenience, define $\mathcal{B}_0 = \mathbf{z}^{(0)}$. Observe now that for all j we have

$$|\tilde{\mathbf{z}}^{(j)} - \mathbf{z}^{(j)}| \leq \sqrt{\ell_{j+1}} . \quad (5.3)$$

Indeed for $j \neq m + k'$, we have by construction that $|\tilde{\mathbf{z}}^{(j)} - \mathbf{z}^{(j)}| \leq \sqrt{\ell_j} \leq \sqrt{\ell_{j+1}}$. For $j = m + k'$, since we adjusted the radius of \mathcal{B}_j to be $\sqrt{\ell_{j+1}}$, we directly have $|\tilde{\mathbf{z}}^{(j)} - \mathbf{z}^{(j)}| \leq \sqrt{\ell_{j+1}}$.

We now show that there exists $C > 0$ such that for any j and any $\tilde{\mathbf{z}}^{(j)} \in \mathcal{B}_j$, we have

$$\left| \mathbf{z}^{(j+1)} - \hat{\mathbb{E}}_{\tilde{\mathbf{z}}^{(j)}}[\mathbf{S}(\ell_{j+1})] \right| \leq C\sqrt{\ell_{j+1}} . \quad (5.4)$$

Indeed, we have

$$\begin{aligned} \hat{\mathbb{E}}_{\tilde{\mathbf{z}}^{(j)}}[\mathbf{S}(\ell_{j+1})] &= (\tilde{\mathbf{z}}_1^{(j)} + \ell_{j+1}\alpha_1, \tilde{\mathbf{z}}_2^{(j)}) \\ &= (\tilde{\mathbf{z}}_1^{(j)} + \mathbf{z}_1^{(j+1)} - \mathbf{z}_1^{(j)}, \tilde{\mathbf{z}}_2^{(j)}) . \end{aligned}$$

The bound on the x -coordinates in Eq. (5.4) now follows from Eq. (5.3). For the y -coordinates,

$$|\tilde{z}_2^{(j)} - z_2^{(j+1)}| \leq |\tilde{z}_2^{(j)} - z_2^{(j)}| + |z_2^{(j+1)} - z_2^{(j)}| \leq 2\sqrt{\ell_{j+1}},$$

using Eq. (5.3) and the fact that $|z_2^{(j+1)} - z_2^{(j)}| \leq \sqrt{\ell_{j+1}}$ by the construction of the sequence $z^{(j)}$.

With Eq. (5.4) in hand, we aim to prove the following claim.

Claim 5.5. *If we denote, for $j = 0, \dots, m + m' + k'$,*

$$p_j(z, \mathcal{B}) := \hat{\mathbf{P}}_z \left(S(\ell_{j+1}) \in \mathcal{B}, S(t)_2 \in [\frac{1}{2}z_2, \frac{3}{2}z_2] \forall t \leq \ell_{j+1} \mid \text{Bdd} \right),$$

then there exists a constant $c = c(\beta) > 0$ such that for every $0 \leq j \leq m + m' + k'$ except $j = m + k'$, we have

$$\min_{z \in \mathcal{B}_j} p_j(z, \mathcal{B}_{j+1}) \geq c \quad \text{for } j \in \{0, \dots, m + m' + k'\} \setminus \{m + k' - 1\}, \quad (5.5)$$

$$\min_{z \in \mathcal{B}_j} p_j(z, \mathcal{B}_{j+1}) \geq c(\log L)^{-b} \quad \text{for } j = m + k' - 1. \quad (5.6)$$

Furthermore, we have

$$\min_{z \in \mathcal{B}_j} p_j(z, \{v\}) \geq cN_n^{-2/3}(\log L)^b \quad \text{for } j = m + m' + k'. \quad (5.7)$$

Proof of Claim 5.5. By Eq. (5.4) and the local limit theorem for the \mathbb{Z}^2 random walk in [38], there exists some constant $c_1 = c_1(\beta) > 0$ such that for every j and every $z \in \mathcal{B}_j$,

$$\hat{\mathbf{P}}_z(S(\ell_{j+1}) \in \mathcal{B}_{j+1} \mid \text{Bdd}) \geq c_1 \frac{|\mathcal{B}_{j+1}|}{\ell_{j+1}}.$$

Recall that for all $j \neq m + k' - 1$ we defined \mathcal{B}_{j+1} to be the ball of radius $\sqrt{\ell_{j+1}}$ about $z^{(j+1)}$. For that exceptional $j = m + k' - 1$, the radius is $\sqrt{\ell_{j+2}} = N_n^{1/3}(\log L)^{-b/2} = \sqrt{\ell_{j+1}}(\log L)^{-b/2}$. Thus, we can infer from the above display that there exists some constant $c_2 = c_2(\beta) > 0$ such that

$$\min_{z \in \mathcal{B}_j} \hat{\mathbf{P}}_z(S(\ell_{j+1}) \in \mathcal{B}_{j+1} \mid \text{Bdd}) \geq c_2 \quad \text{for all } j \neq m + k' - 1,$$

and

$$\min_{z \in \mathcal{B}_{m+k'-1}} \hat{\mathbf{P}}_z(S(\ell_{m+k'}) \in \mathcal{B}_{m+k'} \mid \text{Bdd}) \geq c_2(\log L)^{-b}.$$

Similarly, at $j = m + m' + k'$, we may target the singleton $\{v\}$ and obtain the lower bound

$$\min_{z \in \mathcal{B}_{m+m'+k'}} \hat{\mathbf{P}}_z(S(\ell_{m+m'+k'+1}) = v \mid \text{Bdd}) \geq \frac{c_2}{\ell_{m+m'+k'+1}} = \frac{c_2}{N_n^{2/3}(\log L)^{-b}}.$$

Turning to establish Eq. (5.5) for $j \neq m + k' - 1$, it suffices via a union bound to show that, for every such j and all $z \in \mathcal{B}_j$, we have

$$\hat{\mathbf{P}}_z \left(\max_{t \leq \ell_{j+1}} |S(t)_2 - z_2| > \frac{1}{2}z_2 \mid \text{Bdd} \right) \leq c_2/2.$$

Indeed, by construction, for all j and $z \in \mathcal{B}_j$ we have $z_2 \geq (K/4)N_n^{1/3} > (K/4)\sqrt{\ell_{j+1}}$ for all j , and said bound readily holds for our random walk $S(\cdot)$ by Hoeffding's inequality (for large enough K). To establish Eq. (5.6), we use that the location of the transition point $j = m + k' - 1$ was chosen precisely so that z_2 would then be suitably large. Namely, as $\mathcal{B}_{m+k'-1}$ is centered at $z^{(m+k'-1)}$, we have $z_2 = (1 + o(1))(k' - k)N_n^{1/3} = (1 + o(1))N_n^{1/3}\log L$ for all $z \in \mathcal{B}_{m+k'-1}$. In particular, for the exceptional $j = m + k' - 1$ we have by Hoeffding's inequality (recall $\ell_{j+1} = N_n^{2/3}$ for this j) that

$$\hat{\mathbf{P}}_z \left(\max_{t \leq \ell_{j+1}} |S(t)_2 - z_2| > \frac{1}{2}z_2 \mid \text{Bdd} \right) < L^{-100}.$$

Finally, the last inequality holds true also for the final $j = m + m' + k'$ (in which case for all $z \in \mathcal{B}_j$ we have $z_2 \geq (1 - o(1))N_n^{1/3}(\log L)^b$ vs. an interval length $\ell_{j+1} = N_n^{2/3}(\log L)^{-b}$), thereby establishing Eq. (5.7) and concluding the proof of Claim 5.5. ■

Equipped with the above claim, we can now apply the Domain Markov property after every ℓ_j steps of the random walk, and use the fact that the intersection over j of the events of the form $S(t)_2 \in [\frac{1}{2}z_2, \frac{3}{2}z_2]$ imply $\text{Low} \cap \mathcal{G}_*^\square$, so that Claim 5.5 obtains that

$$\hat{\mathbf{P}}_u(\text{Low}, \mathcal{G}_*^\square, v \in \text{Cpts}(\Gamma) \mid \text{Bdd}) \geq e^{-c(m+m'+k')} N_n^{-2/3} \geq e^{-c'(\log L)^a},$$

for some constant $c' = c'(\beta) > 0$ when $a > 3b/2$. This completes the proof of Lemma 5.4. ■

Remark 5.6. At this point, we have established that (a) w.h.p., when looking at the rectangle delimiting an interval of length $N_n^{2/3}(\log L)^a$ centered on the bottom boundary, the intersection of Γ with its left and right sides, denoted \bar{u}, \bar{v} respectively, is at distance at most $KN_n^{1/3}$ from the bottom boundary; (b) the law of Γ in said rectangle is that of a 2D random walk $S(\cdot)$ (cone-points), with i.i.d. animal decorations between every two cone-points, started at \bar{u} and conditioned (i) to hit \bar{v} , and (ii) that \mathcal{L}_n , a subset of Γ stays at nonnegative heights, thereafter tilted by $\exp(-\mathcal{A}_F^\downarrow(\Gamma)/N_n)$. The subtle point is that the marginal on cone-points is not simply a 2D random walk bridge conditioned to be nonnegative (and tilted by the area): the event that \mathcal{L}_n is nonnegative is not measurable w.r.t. the marginal on cone-points, and could potentially have a delicate pinning effects. Further disruptive is the area term: without random walk estimates, we cannot control the area and rule out the case where it pushes Γ towards the boundary. We wish to rule out these scenarios: indeed, if this had been a standard random walk, entropic repulsion would repel it to height N_n^δ , where conditioning on \mathcal{L}_n or $\text{Cpts}(\Gamma)$ being nonnegative is basically the same. The strategy is as follows:

1. Show (in Lemma 5.8) that at $\pm(\log \log L)N_n^{2/3}$ the height of Γ is at most $K_\varepsilon N_n^{2/3}$ with probability at least $1 - \varepsilon$. (This would have followed from [25, §6.6 and Prop. 6.2] if we had a random walk on cone-points, but the conditioning on nonnegative \mathcal{L}_n vs. cone-points precludes that. Also note that the effect of \mathcal{L}_n vs. cone-point conditioning could have alternatively been addressed by [27, Proposition 13], but the area tilt precludes that approach as well.)
2. Establish (Lemma 5.11) that, without an area term, in the above interval, Γ stays above height N_n^δ away from its two endpoints, except with probability $O(N_n^{-\delta})$.
3. Deduce the same (Lemma 5.12) for the area-tilted Γ by showing that the area term $\frac{\mathcal{A}_F^\downarrow(\Gamma)}{N_n}$ in the above interval is $O((\log \log L)^{3/2})$.

At the end of these steps, we arrive at a 2D random walk on the interval J (we may restrict attention to the cone-points), with endpoints at height at most $K_\varepsilon N_n^{1/3}$, conditioned to be above N_n^δ and tilted by the area delimited by the random walk (rescaled by $1/N_n$).

Thanks to Lemma 5.4, we move to $\mathbf{P}_{\bar{u}, \bar{v}}(\cdot \mid \mathcal{G}^\square, \text{Bdd})$ with $\bar{u}_1 \in [-\frac{1}{2}N_n^{2/3}(\log L)^a, -\frac{1}{4}N_n^{2/3}(\log L)^a]$, $\bar{v}_1 \in [\frac{1}{4}N_n^{2/3}(\log L)^a, \frac{1}{2}N_n^{2/3}(\log L)^a]$, and $\bar{u}_2, \bar{v}_2 \leq KN_n^{1/3}$. We next show that we can find two cone-points \bar{u}', \bar{v}' at the $N_n^{2/3}$ scale (moving inwards from $N_n^{2/3}(\log L)^a$) at height at most $O(N_n^{1/3})$. Define first the event

$$\mathcal{G}_{\text{Cpts}}^\square := \{\text{All cone-points of } \Gamma \text{ have height } \geq 0\}.$$

We start with a tightness result from [25].

Proposition 5.7 ([25, Section 6.6]). *For every $\varepsilon > 0$, there exists a constant $K_\varepsilon > 0$ such that the following holds. For any u, v such that $u_2, v_2 \leq CN_n^{1/3}$, and for every fixed vertical line $x = m$*

in between u, v , the height of (the linear interpolation between points of) S at $x = m$ is at most $K_\varepsilon N_n^{1/3}$ except with probability ε under $\mathbf{P}_{u,v}(\cdot \mid \text{Bdd}, \mathcal{G}_{\text{Cpts}}^\perp)$.

Proof. The measure $\mathbf{P}_{u,v}(\cdot \mid \text{Bdd}, \mathcal{G}_{\text{Cpts}}^\perp)$ is exactly an area tilted 2D random walk conditioned to stay above height 0. This tightness result was shown in [25, Section 6.6], noting Proposition 6.2 there in order to extend the result to scales larger than $N_n^{2/3}$ (e.g., applying to the case when $v_1 - u_1 = O(N_n^{2/3}(\log L)^{25})$). \blacksquare

Lemma 5.8. *Let $\mathcal{E}(K)$ be the event that at the first cone point after $-(\log \log L)N_n^{2/3}$ and the last cone point before $(\log \log L)N_n^{2/3}$, the height of γ is at most $KN_n^{1/3}$. Then, for any $\varepsilon > 0$, there exists a K_ε such that*

$$\mathbf{P}_{\bar{u},\bar{v}}(\mathcal{E}(K_\varepsilon) \mid \text{Bdd}, \mathcal{G}^\perp) \geq 1 - \varepsilon.$$

Proof. It suffices to prove that under $\mathbf{P}_{\bar{u},\bar{v}}$, the first cone-point after $-(\log \log L)N_n^{2/3}$ has height $\leq K_\varepsilon N_n^{2/3}$ except with probability $\varepsilon/2$. Fix any $\delta < 1/3$, and let $u^{(L)}$ be the last cone-point of Γ such that $u_1^{(L)} \leq -(\log \log L)N_n^{2/3}$ (“to the left” of $-(\log \log L)N_n^{2/3}$) and $u_2^{(L)} \leq N_n^\delta$. If there is no such point, define $u^{(L)} = \bar{u}$. Similarly, let $v^{(R)}$ be the last cone-point of Γ with $v_1^{(R)} \geq -(\log \log L)N_n^{2/3}$ (“to the right” of $-(\log \log L)N_n^{2/3}$) such that $v_2^{(R)} \leq N_n^\delta$. If there is no such point, define $v^{(R)} = \bar{v}$. Then, the law of the cone-points of Γ in between $u^{(L)}$ and $v^{(R)}$ under the measure \mathbf{P} is exactly the area tilted 2D random walk from $u^{(L)}$ to $v^{(R)}$ conditioned to stay above N_n^δ (since the event \mathcal{G}^\perp is implied by the cone-points staying above N_n^δ and Bdd , and hence can be ignored in the conditioning). We can now conclude by Proposition 5.7 after translating vertically by N_n^δ . \blacksquare

By Lemma 5.8, we can now move to $\mathbf{P}_{\bar{u}',\bar{v}'}(\cdot \mid \mathcal{G}^\perp, \text{Bdd})$, where $\bar{u}'_1 = -(\log \log L)N_n^{2/3} \pm (\log L)^{50}$, $\bar{v}'_1 = (\log \log L)N_n^{2/3} \pm (\log L)^{50}$, and $\bar{u}'_2, \bar{v}'_2 \leq K_\varepsilon N_n^{1/3}$. Before we can appeal to the convergence to Ferrari–Spohn shown in [25], we still need to handle the issue of having a floor with respect to \mathfrak{L}_n vs. a floor with respect to $\text{Cpts}(\Gamma)$. Furthermore, we want to show that even though \bar{u}'_2, \bar{v}'_2 are potentially lower than $O(N^{1/3})$, the entropic repulsion will bring the height back to $O(N^{1/3})$ on an interval $[-TN^{2/3}, TN^{2/3}]$.

Up until Lemma 5.12, we now forget the area tilt. The next proposition from [27, Proposition 13] deals with the fact that our floor condition is with respect to \mathfrak{L}_n and not $\text{Cpts}(\Gamma)$ when there is no area tilt, capturing that fact that Γ quickly repels away from the floor after which point the two events are essentially the same. The statement in [27] assumed further that u and v are on the floor at height 0, but as commented after [27, Theorem 3], such repulsion results are only easier when the end-points are away from the floor. See also the comment before Proposition 6.3 regarding Ising polymers vs. disagreement polymers.

Proposition 5.9 ([27, Proposition 13]). *There exists a constant $c(\beta) > 0$ such that for all u, v with $u_2, v_2 \geq 0$,*

$$\hat{\mathbf{P}}_{u,v}(\gamma \subset \mathbb{H}_+ \mid \mathcal{G}_{\text{Cpts}}^\perp) \geq c(\beta).$$

In particular, since $\gamma \subset \mathbb{H}_+$ implies $\mathfrak{L}_n \subset \mathbb{H}_+$, we have

$$\hat{\mathbf{P}}_{u,v}(\mathcal{G}^\perp \mid \mathcal{G}_{\text{Cpts}}^\perp) \geq c(\beta).$$

In order to show entropic repulsion in the measure $\hat{\mathbf{P}}_{\bar{u}',\bar{v}'}(\cdot \mid \mathcal{G}^\perp, \text{Bdd})$, we will use the following random walk estimates from [26], recorded here for convenience.

Theorem 5.10 ([26, Thm 5.1]). *Consider the 2D effective random walk $S(\cdot)$. Let τ_G be the hitting time of a set $G \subset \mathbb{Z}^2$. There exists a constant $C > 0$ such that, for every sequence δ_ℓ that goes to 0*

arbitrarily slowly as $\ell \rightarrow \infty$, and for all $u, v \in \mathbb{Z}$ such that $u/\sqrt{\ell}, v/\sqrt{\ell} \in (0, \delta_\ell)$,

$$\hat{\mathbf{P}}_{(0,u)}(\tau_{(\ell,v)} < \tau_{\mathbb{H}_-} < \infty \mid \text{Bdd}) \sim C \frac{h^+(u)h^-(v)}{\ell^{3/2}},$$

where h^\pm are positive harmonic (in particular, asymptotically linear) functions.

Applying the above for $\ell = N_n^{2/3}(\log \log L)$, we show entropic repulsion to N_n^δ .

Lemma 5.11. *Fix a constant $0 < \delta < 1/12$, and consider the event **Repel** that in the interval $J := [-(\log \log L)N_n^{2/3} + N_n^{8\delta}, (\log \log L)N_n^{2/3} - N_n^{8\delta}]$, all cone-points of Γ lie above height N_n^δ . Then, there exists $C > 0$ such that*

$$\hat{\mathbf{P}}_{\bar{u}',\bar{v}'}(\text{Repel} \mid \mathcal{G}^\square, \text{Bdd}) \geq 1 - CN_n^\delta.$$

Proof of Lemma 5.11. By Proposition 5.9 (and Lemma 5.2) it suffices to prove that

$$\hat{\mathbf{P}}_{\bar{u}',\bar{v}'}(\text{Repel}^c \mid \mathcal{G}_{\text{Cpts}}^\square, \text{Bdd}) \leq C' N_n^\delta.$$

Set $\ell = N_n^{2/3}(\log \log L)$. Since $\bar{u}'_2, \bar{v}'_2 \leq K_\varepsilon N_n^{1/3} = o(\sqrt{\bar{v}'_1 - \bar{u}'_1})$, we are in the regime of Theorem 5.10. We follow the computation of [26, Lemma 3.6] (except simpler in our case, as we have already moved to the setting of the effective random walk). We can sum over the first cone-point z below height ℓ^δ in the interval J and apply random walk estimates, obtaining

$$\begin{aligned} \hat{\mathbf{P}}_{\bar{u}',\bar{v}'}(\text{Repel}^c \mid \mathcal{G}_{\text{Cpts}}^\square, \text{Bdd}) &\leq \sum_{z_1=-\ell+\ell^{8\delta}}^{\ell-\ell^{8\delta}} \sum_{z_2=0}^{\ell^\delta} \hat{\mathbf{P}}_{\bar{u}',\bar{v}'}(\tau_z < \infty \mid \mathcal{G}_{\text{Cpts}}^\square, \text{Bdd}) \\ &= \sum_{z_1=-\ell+\ell^{8\delta}}^{\ell-\ell^{8\delta}} \sum_{z_2=0}^{\ell^\delta} \frac{\hat{\mathbf{P}}_{\bar{u}'}(\tau_z < \tau_{\bar{v}'} < \tau_{\mathbb{H}_-} < \infty \mid \text{Bdd})}{\hat{\mathbf{P}}_{\bar{u}'}(\tau_{\bar{v}'} < \tau_{\mathbb{H}_-} < \infty \mid \text{Bdd})} \\ &= \sum_{z_1=-\ell+\ell^{8\delta}}^{\ell-\ell^{8\delta}} \sum_{z_2=0}^{\ell^\delta} \frac{\hat{\mathbf{P}}_{\bar{u}'}(\tau_z < \tau_{\mathbb{H}_-} < \infty \mid \text{Bdd}) \mathbf{P}_z(\tau_{\bar{v}'} < \tau_{\mathbb{H}_-} < \infty \mid \text{Bdd})}{\mathbf{P}_{\bar{u}'}(\tau_{\bar{v}'} < \tau_{\mathbb{H}_-} < \infty \mid \text{Bdd})} \\ &\leq C\ell^{3/2+\delta} \sum_{z_1=-\ell+\ell^{8\delta}}^{\ell-\ell^{8\delta}} \frac{\ell^\delta}{(z_1 + \ell)^{3/2}} \cdot \frac{\ell^\delta}{(\ell - z_1)^{3/2}} \\ &\leq 2C\ell^{3\delta} \sum_{z_1=-\ell+\ell^{8\delta}}^0 \frac{1}{(z_1 + \ell)^{3/2}} \cdot \frac{1}{(1 - \frac{z_1}{\ell})^{3/2}} \\ &\leq 2C\ell^{3\delta} \sum_{z_1=-\ell+\ell^{8\delta}}^0 \frac{1}{(z_1 + \ell)^{3/2}} \\ &\leq 4C\ell^{3\delta}\ell^{-4\delta} = 4C\ell^{-\delta}. \end{aligned}$$

The first two lines follow by a union bound and conditional probability, noting that the event $\tau_{\mathbb{H}_-} < \infty$ occurs with probability 1 as the random walk has no drift in the y -coordinate. The third line follows by the markov property, and the fourth line uses Theorem 5.10. The rest of the computations are purely algebraic, requiring only that $\delta < 1/8$. \blacksquare

We can now add back the area tilt to obtain entropic repulsion in the measure $\mathbf{P}_{\bar{u}',\bar{v}'}(\cdot \mid \mathcal{G}^\square, \text{Bdd})$.

Lemma 5.12. *For the event **Repel** defined in Lemma 5.11,*

$$\mathbf{P}_{\bar{u}',\bar{v}'}(\text{Repel} \mid \mathcal{G}^\square, \text{Bdd}) \geq 1 - o(1).$$

Proof. We have

$$\mathbf{P}_{\bar{u}',\bar{v}'}(\text{Repel}^c \mid \mathcal{G}^\square, \text{Bdd}) = \frac{\hat{\mathbf{E}}_{\bar{u}',\bar{v}'}[e^{-\frac{\mathcal{A}_F^\downarrow(\Gamma)}{N_n}} \mathbb{1}_{\text{Repel}^c} \mid \mathcal{G}^\square, \text{Bdd}]}{\hat{\mathbf{E}}_{\bar{u}',\bar{v}'}[e^{-\frac{\mathcal{A}_F^\downarrow(\Gamma)}{N_n}} \mid \mathcal{G}^\square, \text{Bdd}]} \leq \frac{\hat{\mathbf{E}}_{\bar{u}',\bar{v}'}[\mathbb{1}_{\text{Repel}^c} \mid \mathcal{G}^\square, \text{Bdd}]}{\hat{\mathbf{E}}_{\bar{u}',\bar{v}'}[e^{-\frac{\mathcal{A}_F^\downarrow(\Gamma)}{N_n}} \mid \mathcal{G}^\square, \text{Bdd}]} \quad (5.8)$$

The numerator is $\leq CN_n^\delta$ by Lemma 5.11. To lower bound the denominator, we wish to show that with at least constant probability in $\hat{\mathbf{P}}_{\bar{u}',\bar{v}'}(\cdot \mid \mathcal{G}^\square, \text{Bdd})$, the maximum of Γ is at most $O(N_n^{1/3} \sqrt{\log \log L})$, for then we will have a lower bound on the denominator by $e^{-c(\log \log L)^{3/2}}$, which is negligible as $e^{c(\log \log L)^{3/2}} \ll N_n^\delta$. By Proposition 5.9 again, we can lower bound this probability under $\hat{\mathbf{P}}_{\bar{u}',\bar{v}'}(\cdot \mid \mathcal{G}_{\text{Cpts}}^\square, \text{Bdd})$ instead. Yet this is now a 2D random walk whose start and end-points are of a lower order than the square root of the horizontal distance between them. By [26, Theorem 5.3], this converges to a Brownian excursion, where the bound on the maximum is a well known fact. \blacksquare

Proof of Theorem 5.1. As per Steps 1 to 4 in the strategy outlined at the beginning of this section, by Lemmas 5.2, 5.4, 5.8 and 5.12, we can move to the measure $\mathbf{P}_{\bar{u}',\bar{v}'}(\cdot \mid \text{Repel}, \mathcal{G}^\square, \text{Bdd})$. As noted in the proof of Lemma 5.12, we can reveal the first and last cone-points \bar{u}'', \bar{v}'' in J , and by the linearity of the cones we have $|\bar{u}''_2 - \bar{u}'_2|, |\bar{v}''_2 - \bar{v}'_2| \leq N^{8\delta}$. On the interval J in between \bar{u}'' and \bar{v}'' , the event $\text{Repel} \cap \text{Bdd}$ already implies \mathcal{G}^\square , so we have reduced to the measure $\mathbf{P}_{\bar{u}'',\bar{v}''}(\cdot \mid \text{Repel}, \text{Bdd})$. At this point, we have an area-tilted 2D random walk bridge conditioned to stay above height N_n^δ with increment law $\mathbb{P}(x, y) = \mathbb{P}^{h_{(1,0)}^n}(X(\Gamma) = (x, y) \mid |(x, y)| \leq (\log L)^{50})$, as we have removed every constraint on Γ that is not measurable with respect to $\text{Cpts}(\Gamma)$. By the exponential tails on the increments in Proposition 3.19, changing the increment law to $\mathbb{P}^{h_{(1,0)}^n}(X(\Gamma) = (x, y))$ only changes the law of the random walk in total variation by $e^{-c(\log L)^{50}}$. Moreover, the endpoints \bar{u}'', \bar{v}'' satisfy $\bar{v}''_1 - \bar{u}''_1 \geq TN^{2/3}$ for a constant T sufficiently large depending on K_ε , and $\bar{u}''_2, \bar{v}''_2 \leq K_\varepsilon N_n^{1/3}$. In particular, we have reduced to the starting point of [25, Sec. 6], where it was shown that the area-tilted 2D random walk bridge above a floor with increment law $\mathbb{P}^{h_{(1,0)}^n}(X(\Gamma) = (x, y))$ converges to FS_{σ_n} (see [25, Eq. (6.55)]) in the limit $L \rightarrow \infty$ followed by $T \rightarrow \infty$.

Thus far we have established that, w.h.p., for every fixed $n \geq 0$, the process $Y_n(t)$ is stochastically dominated by a process $Z_n(t)$ (the 2D random walk with an area-tilt) that weakly converges to FS_{σ_n} as $L \rightarrow \infty$. Moreover, as the lower level-lines are w.h.p. below the randomly constructed Q , the proof in fact showed that, w.h.p., $Y_n(t) \preceq Z_n(t)$ conditionally on \mathcal{L}_k for $k > n$ (and in particular, conditionally on the rescaled portions of these level-lines, $Y_k(t)$ for $k > n$). Hence, the stronger statement concerning the joint law follows from the next elementary fact.

Fact 5.13. *Suppose $Y_1, \dots, Y_m, Z_1, \dots, Z_m$ are \mathbb{R}^d -valued random variables such that $Y_n \preceq Z_n$ for each n . Suppose furthermore that for each n , conditional on $\{Y_k, k > n\}$, we have $Y_n \preceq Z_n$. Then, the joint law satisfies $(Y_1, \dots, Y_m) \preceq (Z'_1, \dots, Z'_m)$ where $\{Z'_k\}_{k=1}^m$ are independent with $Z'_k \stackrel{d}{=} Z_k$.*

Note that here we conditioned on the lower level-lines \mathcal{L}_k for $k > n$; we cannot condition on the upper level-lines \mathcal{L}_k for $k < n$ (unlike the proof of Theorem 6.1, the matching lower bound for Theorem 5.1, where we condition on the upper level-lines but cannot condition on the lower ones). The reason for this is that, in our proof, we construct a rectangle Q of height $N_n^{2/3}$ surrounding the typical location of \mathcal{L}_n , and while it is positioned so as not to intersect the lower level-line \mathcal{L}_{n+1} , it will impede on the support of \mathcal{L}_{n-1} . (The necessity for such a rectangle, as opposed to one of height (say) $N^{1/3}(\log L)^c$, is that the latter entails pinning issues with the top side of Q , and the present tools cannot resolve those when said boundary is random and wiggly.) \blacksquare

Remark 5.14. In the proof of the convergence of the area-tilted 2D random walk to FS_σ in [25, Sec. 6], the initial hypothesis that the height of the endpoints (our \bar{u}''_2, \bar{v}''_2) is replaced by

[25, Eq. (6.55)], requiring that they would belong to $[c_\varepsilon N_n^{1/3}, K_\varepsilon N_n^{1/3}]$. Said assumption qualifies for an application of a coupling tool ([25, Prop. 6.2]). At this stage, both here and in the setting of [25], the heights of the endpoints could potentially be $o(N_n^{1/3})$. This is not an issue for their argument, as the proof of [25, Prop. 6.2] remains valid even without the assumption of the $c_\varepsilon N_n^{1/3}$ lower bound.

Alternatively, it is not difficult to show that one can move from the assumption $\bar{u}_2'', \bar{v}_2'' \leq K_\varepsilon N_n^{1/3}$ to $\bar{u}_2'', \bar{v}_2'' \in [c_\varepsilon N_n^{1/3}, K_\varepsilon N_n^{1/3}]$ with probability at least $1 - \varepsilon$. Indeed, it suffices to consider u'' (then applying the same argument to v'' by a union bound). Let g_L^ε be the maximum integer such that $u_2'' \geq g_L^\varepsilon$ with probability at least $1 - \varepsilon$. If $\liminf_{L \rightarrow \infty} g_L^\varepsilon N_n^{-1/3} > 0$ then u'' satisfies the sought condition (with c_ε as this \liminf value). Otherwise, we may apply the same logic to the “last” cone-point u''' in the interval of length $N_n^{2/3}$ starting at u'' (i.e., the cone-point x with maximizing $x_1 - u''_1$ out of those with $x_1 - u''_1 < N_n^{2/3}$), looking at the maximum integer h_L^ε such that $u'''_2 \geq h_L^\varepsilon$ with probability at least $1 - \varepsilon$. If $\liminf_{L \rightarrow \infty} h_L^\varepsilon N_n^{-1/3} > 0$, we may use u''' as our desired endpoint. It remains to handle the case where $u_2'', u'''_2 < \delta_L N_n^{1/3}$ for $\delta_L = o(1)$ (for $\delta_L = (g_L^\varepsilon \vee h_L^\varepsilon) N_n^{-1/3}$). For such u'', u''' , by [26, Thm. 5.3], the 2D random walk with *no area-tilt* converges to a Brownian excursion. Consequently, we claim that there exists some c_ε such that the cone-points at the middle of this interval—concretely, take the first cone-point w such that $w_1 - u''_1 > N_n^{2/3}/2$ —are at height at least $c_\varepsilon N_n^{1/3}$ with probability at least $1 - \varepsilon$. Indeed, we may bound the complement as in Eq. (5.8): the denominator is uniformly bounded away from zero, e.g., we may bound it from below by $\frac{1}{2} \exp(-A)$ where A is the median of the area of a standard Brownian excursion B_t ($0 < t < 1$), whereas the numerator is bounded by $\mathbb{P}(B_{1/2} < \varepsilon)$.

6. LOWER BOUND

Our goal in this section is the following bound on the distance of \mathfrak{L}_n from the bottom side of Λ_L .

Theorem 6.1. *In the setting of Theorem 1.1, fix $n \geq 0$ and $K > 0$, let $\rho_n(x)$ be the maximum vertical distance of \mathfrak{L}_n above $x + (\frac{L}{2}, 0)$ for $-N_n^{2/3} \leq x \leq N_n^{2/3}$, and set $\sigma_n^2 > 0$ as per Definition 3.20. Then every weak limit point $\mathbf{Y}_n(t)$ of the process $Y_n(t) := N_n^{-1/3} \rho_n(t N_n^{2/3})$ ($t \in [-K, K]$), as $L \rightarrow \infty$, satisfies*

$$\mathbf{Y}_n \succeq \mathsf{FS}_{\sigma_n}.$$

Moreover, for every fixed m , every weak limit point $(\mathbf{Y}_n(t))_{n \leq m}$ of the processes $(Y_n(t))_{n \leq m}$ satisfies

$$(\mathbf{Y}_n)_{n \leq m} \succeq \bigotimes_{n \leq m} \mathsf{FS}_{\sigma_n}.$$

Since $\rho_n(x)$ is a decreasing function of ϕ , en route to proving Theorem 6.1 in this section we may apply monotonicity arguments that are increasing (e.g., raising the heights of boundary vertices).

Fix $T > K$ (eventually we will take $T \rightarrow \infty$). Let R be the $3TN_n^{2/3} \times N_n^{1/3}(\log L)$ rectangle centered at $x = 0$ such that the bottom side of ∂R coincides with the bottom of $\partial \Lambda_L$. Let $A = (-TN_n^{2/3}, 0)$ and $B = (TN_n^{2/3}, 0)$. We start with a simpler analogue of Lemma 4.7, the proof of which we also postpone to Appendix B.

Lemma 6.2. *For R, A, B , as defined above, assume we know that w.h.p. under $\pi_{\Lambda_L}^0$, all of R lies in the exterior of \mathfrak{L}_{n-1} (for $n = 0$, there is no assumption). Then, there exists a $\pi_{\Lambda_L}^0$ -measurable distribution on connected regions $Q \subset R$ with marked boundary conditions ξ satisfying Items 1 to 3 below, such that the following holds. If \mathcal{A}_1 is the area of the interior of \mathfrak{L}_n intersected with Q , and \mathcal{A}_2 is the area above the $(H - n)$ level-line in Q under π_Q^ξ , then $\mathcal{A}_1 \subset \mathcal{A}_2$ w.h.p. under $\pi_{\Lambda_L}^0$.*

(1) Q is simply connected,

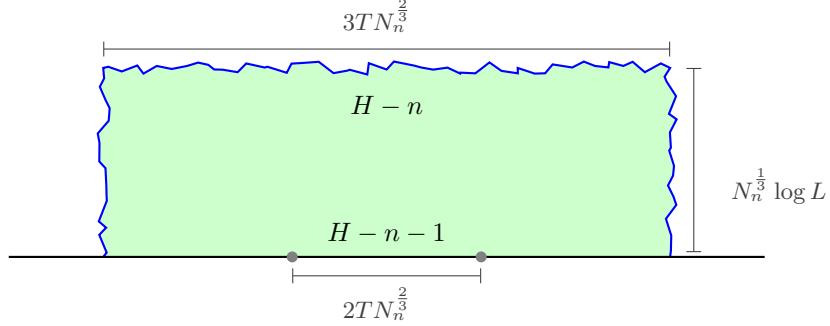


FIGURE 9. An instantiation of the domain Q used in the proof of the lower bound of \mathfrak{L}_n . The bottom boundary of Q coincides with the bottom boundary of Λ_L . The two gray points on the boundary mark the change in boundary conditions from $H - n - 1$ to $H - n$. In contrast with Fig. 7, the conditioning on $\phi_x \geq 0$ is present in all of Q .

- (2) $\text{dist}(\partial Q, \partial R) \leq (\log L)^2$, and the bottom side of ∂Q coincides with the bottom side of ∂R ,
- (3) The boundary conditions ξ assigns height $H - n - 1$ on the straight line between A and B , and $H - n$ on the remainder of ∂Q .

Let Q, ξ be any domain and boundary condition satisfying Lemma 6.2 (see Fig. 9). It now suffices to prove Theorem 6.1 where \mathfrak{L}_n is the $H - n$ level-line in π_Q^ξ . We start by showing the lower bound for \mathfrak{L}_0 so that w.h.p. the y -coordinate of \mathfrak{L}_0 is $O(N_0^{1/3})$ above the interval $[-\frac{3}{2}TN_1^{2/3}, \frac{3}{2}TN_1^{2/3}]$. This justifies the assumption in Lemma 6.2 for $n = 1$, since $N_1^{1/3} \leq N_0^{1/3} e^{-c\beta\sqrt{\log L/\log\log L}}$ by Eq. (2.3). We can then prove the lower bound for \mathfrak{L}_1 , and proceed inductively.

In the rest of this section, we will use “height” to refer to the vertical distance above the bottom of Λ_L . Let \mathcal{G}^\square be the event that \mathfrak{L}_n stays under the horizontal line \mathcal{H} at distance $2(\log L)^2$ below the lowest point on the top side of Q . Then, $\pi_Q^\xi \preceq \pi_Q^\xi(\cdot \mid \mathcal{G}^\square)$. We can now follow the same preliminary steps as in Section 5 to move to a polymer model. By Lemma 4.13, we can condition further on the event \mathcal{G}_1 defined there, so that by Proposition 2.3 (and Eq. (3.4)) we have

$$\pi_Q^\xi(\gamma \mid \mathcal{G}^\square, \mathcal{G}_1) \propto \exp\left(-\mathcal{E}_\beta^*(\gamma) + \frac{\mathcal{A}_Q(\gamma)}{N_n} + \mathfrak{I}_Q(\gamma)\right) = q_{Q;Q,1}^n(\gamma). \quad (6.1)$$

By Eq. (4.9) we can move to the polymer model with weights $q_{Q;Q,1}^n$ and partition function

$$Z_{Q;Q,1}^n(A, B \mid \mathcal{G}^\square) = \sum_{\gamma \in \mathcal{P}_Q(A, B) \cap \mathcal{G}^\square} q_{Q;Q,1}^n(\gamma).$$

Forgetting the area term for now, the first step is to control the effect of the interactions with the bottom of Q , effectively eliminating any potential pinning effects. This was shown in the half-space \mathbb{H}_+ in [27] for Ising polymers. However, as mentioned in the end of the introduction there, the proof is more robust and allows for more complicated geometries in γ (e.g. our connected components of bonds), and more complicated energies (e.g. our $\mathcal{E}_\beta^*(\gamma)$), as long as the polymer model in question features the Ornstein–Zernicke results found in Lemma 3.11 and Proposition 3.19. The particular structure of the Ising polymer model is only used to show that the increment measure has $1 - \varepsilon_\beta$ mass on three basic irreducible components (see [27, Definition 18]), which simplifies the proof of certain random walk estimates. The same is true for the class of polymer models considered in this paper, for the same reason that allowed us to transfer the results of [17] in Section 3: the assumption in Eq. (3.8) that $e^{-\mathcal{E}_\beta^*(\gamma)}$ decays exponentially in $\beta\mathcal{N}(\gamma)$ implies that the lowest energy disagreement polymers coincide with those of the Ising polymer setting, and the increase in entropy of other polymers is negligible compared to their energy.

Proposition 6.3 ([27, Theorem 2]). *There exists a constant $C(\beta) > 0$ such that*

$$\frac{1}{C(\beta)} \widehat{Z}_{\mathbb{H}_+, \mathbb{Z}^2}^n(A, B) \leq \widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B) \leq C(\beta) \widehat{Z}_{\mathbb{H}_+, \mathbb{Z}^2}^n(A, B).$$

In our setting, we will need the following corollary.

Corollary 6.4. *For the above defined Q and endpoints A, B , one has*

$$\widehat{Z}_{Q, Q}^n(A, B | \mathcal{G}^\square) = (1 + o(1)) \widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B).$$

Proof. First note that by a simple Peierls argument (mapping the polymer to the straight line between A, B), we can rule out the event that $|\gamma| > 1.1|B - A|$ regardless of what the domain or interaction with the boundary. In particular, this implies that

$$\widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B | \mathcal{G}^\square) = (1 + o(1)) \widehat{Z}_{Q, Q}^n(A, B | \mathcal{G}^\square),$$

and by the decay of the Φ functions (since by the above Peierls argument we can assume on \mathcal{G}^\square that γ stays $(\log L)^2$ away from the top and sides of ∂Q , and the bottom side of ∂Q coincides with ∂H), we also have

$$\widehat{Z}_{Q, \mathbb{H}_+}^n(A, B | \mathcal{G}^\square) = (1 + o(1)) \widehat{Z}_{Q, Q}^n(A, B | \mathcal{G}^\square).$$

Combining the last two displays results in

$$\widehat{Z}_{Q, Q}^n(A, B | \mathcal{G}^\square) = (1 + o(1)) \widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B | \mathcal{G}^\square). \quad (6.2)$$

Next, to eliminate the conditioning on \mathcal{G}^\square , we turn to $\widehat{Z}_{\mathbb{H}_+, \mathbb{Z}^2}^n$ and argue that

$$\widehat{Z}_{\mathbb{H}_+, \mathbb{Z}^2}^n(A, B | \mathcal{G}^\square) \geq (1 - o(1)) \widehat{Z}_{\mathbb{H}_+, \mathbb{Z}^2}^n(A, B). \quad (6.3)$$

Indeed, suppose first that the domain restriction induces only that the cone-points have nonnegative heights (instead of also forcing all of γ to be nonnegative). Then, [26, Theorem 5.3] shows that Γ with weight $\widehat{q}_{\mathbb{Z}^2}^n$ and partition function $\widehat{Z}_{\mathbb{Z}^2, \mathbb{Z}^2}^n(A, B | \text{height}(\text{Cpts}(\Gamma)) \geq 0)$ converges weakly to a Brownian excursion upon rescaling, whence the probability of reaching height $N_n^{1/3}(\log L)$ is $o(1)$ by standard estimates on the Brownian excursion. We can then use Proposition 5.9 to translate the same result to the our setting with a partition function of $\widehat{Z}_{\mathbb{H}_+, \mathbb{Z}^2}^n$, proving Eq. (6.3). We now deduce the analogue of Eq. (6.3) for interactions in \mathbb{H}_+ as opposed to \mathbb{Z}^2 :

$$\widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B | \mathcal{G}^\square) \geq (1 - o(1)) \widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B). \quad (6.4)$$

Indeed, recall that Lemma 3.18 implies that events that have a probability of $o(1)$ in the measure associated to $\widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B)$ also have probability $o(1)$ in the analogue for $\widehat{Z}_{\mathbb{H}_+, \mathbb{Z}^2}^n(A, B)$ as long as the two partition functions are up to a multiplicative constant apart; Proposition 6.3 provides exactly that hypothesis. Since $\widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B | \mathcal{G}^\square)$ is a subset of the sum in $\widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B | \mathcal{G}^\square)$, Eq. (6.4) yields

$$\widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B | \mathcal{G}^\square) = (1 - o(1)) \widehat{Z}_{\mathbb{H}_+, \mathbb{H}_+}^n(A, B),$$

which, after combining with Eq. (6.2), completes the proof. \blacksquare

Proof of Theorem 6.1. First observe that we can replace $\mathcal{A}_Q(\gamma)$ with $-|D_1|$, the area below γ , in Eq. (6.1) by just a renormalization. Now we are in the same setting as in [7], and we recall the proof of [7, Theorems 7.1] to show the necessary adjustments. For convenience, let $\mathbf{P}_{V, U}^n, \hat{\mathbf{P}}_{V, U}^n$ denote the measures on Γ associated to $Z_{V, U}^n, \widehat{Z}_{V, U}^n$ respectively. For the case of the Ising polymer model with no area tilt $\hat{\mathbf{P}}_{\mathbb{H}_+, \mathbb{H}_+}^n$, the authors prove that w.h.p.,

- (1) Γ has a linear length and number of cone-points ([7, Lemma 5.7], also proved here in Lemma 3.11),

- (2) The left, right, and irreducible components of Γ have size no more than $(\log L)^2$ ([7, Lemmas 5.5, 5.8]), and
- (3) The cone-points of Γ stay above height N_n^δ in the interval $[A_1 + N_n^{4\delta}, B_2 - N_n^{4\delta}]$ ([7, Lemma 5.11]).

The proof of the above inputs relied only on results from [27] and the product structure from the Ornstein–Zernicke theory, both of which we have in our disagreement polymer model, whence their results translate immediately to our setting in $\hat{\mathbf{P}}_{\mathbb{H}_+, \mathbb{H}_+}^n$ for general disagreement polymers. We can then use Corollary 6.4 (with Lemma 3.18) to obtain that the same results hold w.h.p. in $\hat{\mathbf{P}}_{Q, Q}^n(\cdot \mid \mathcal{G}^\square)$.

To reintroduce the area tilt, we first replace $\mathcal{A}_Q(\gamma)$ in Eq. (6.1) by $-|D_1|$ by just a renormalization, which does not change the above measures. Then, the same argument leading to Eqs. (6.3) and (6.4) (i.e., using convergence to Brownian excursion and then Proposition 6.3) proves that for some constant $c = c(T) > 0$,

$$\hat{\mathbf{E}}_{\mathbb{H}_+, \mathbb{H}_+}^n[e^{-|D_1|/N_n}] \geq c.$$

(See also [7, Claim 7.4] for more details.¹⁰) Again Corollary 6.4 and Lemma 3.18 implies the same lower bound with respect to $\hat{\mathbf{E}}_{Q, Q}^n(\cdot \mid \mathcal{G}^\square)$. This in turn implies that the events in Items 1 to 3 hold w.h.p. in $\mathbf{P}_{Q, Q}^n(\cdot \mid \mathcal{G}^\square)$, as per a computation analogous to Eq. (5.8). Finally, as in Remark 5.3, this allows us to replace $|D_1|$ with the area below the linear interpolation of the cone-points of Γ at the cost of a multiplicative $1 + o(1)$ to the weight.

The consequence of all of the above is that in $\mathbf{P}_{Q, Q}^n(\cdot \mid \mathcal{G}^\square)$, upon looking at the first and last cone-points u, v in the interval $[A_1 + N_n^{4\delta}, B_1 - N_n^{4\delta}]$, the law of $\text{Cpts}(\Gamma)$ in between u, v can be coupled to an area-tilted 2D random walk bridge conditioned to stay above height N_n^δ , with increment law given by $\mathbb{P}(x, y) = \mathbb{P}^{h_{(1,0)}}(X(\Gamma) = (x, y))$. The convergence to the Ferrari–Spohn diffusion FS_{σ_n} in the limit $L \rightarrow \infty$ followed by $T \rightarrow \infty$ now follows by [25, Sec. 6].

Thus, we have established that w.h.p., for every fixed $n \geq 0$, the process $Y_n(t)$ stochastically dominates a process $Z_n(t)$ (the 2D random walk with an area-tilt) that weakly converges to FS_{σ_n} . As noted in the induction below Lemma 6.2, the upper level-lines are w.h.p. above the randomly constructed Q , so the proof in fact shows that w.h.p., $Y_n(t) \succeq Z_n(t)$ conditionally on \mathcal{L}_k for $k < n$ (and in particular, conditionally on the rescaled portions of these level-lines, $Y_k(t)$ for $k < n$). Hence, the stronger statement concerning the joint law follows from Fact 5.13 (for the other direction of stochastic domination, which is easily deduced by multiplying the random variables by -1).

Note that this conditioning is in the reverse direction as in the proof of Theorem 5.1, and we cannot instead condition on the lower level-lines \mathcal{L}_k for $k > n$. Indeed, revealing \mathcal{L}_{n+1} reveals a wiggly boundary at $o(N_n^{1/3})$ with only size $(\log L)^2$ perturbations, which appears to be no different from the flat boundary we have in the construction of Q at the macroscopic level. However, our proof relies strongly on the depinning proved in [27], and it is unclear how to extend those results to the case of a wiggly boundary. ■

7. EXTENSION TO ALL $|\nabla\phi|^p$ MODELS

In this section we prove Theorem 1.5, extending Theorem 1.1 to $|\nabla\phi|^p$ models for fixed $p > 1$.

¹⁰More work had to be done in [7] to show convergence to a Brownian excursion because the 2D random walk there was not symmetric in the y -coordinate, and hence the authors could not directly apply [26, Theorem 5.3]. Our setting does feature this symmetry, so we cite [26] for the convergence.

7.1. Large deviations. The following analogues of Eqs. (2.3) to (2.5) for $1 < p < 2$ ([35, Thm. 5.1]) and $2 < p < \infty$ ([35, Thm. 5.5]) are known:

$$\frac{\hat{\pi}_\infty^{(p)}(\phi_o = h)}{\hat{\pi}_\infty^{(p)}(\phi_o = h-1)} \leq \exp \left[-c\beta h^{(p-1) \wedge 1} \right], \quad (7.1)$$

$$\exp \left[-c_1 \beta h^{p \wedge 2} \right] \leq \hat{\pi}_\infty^{(p)}(\phi_o = h) \leq \exp \left[-c_2 \beta h^{p \wedge 2} \right], \quad (7.2)$$

$$\hat{\pi}_\infty^{(p)}(\phi_z = h \mid \phi_o = h) \leq \exp \left[-c\beta h^{(p-1) \wedge 1} \right]. \quad (7.3)$$

As was the case for $p = 2$, our proof will require a refined version of Eq. (7.3), namely that one appearing below in Eq. (7.6) below. Let $\mathcal{B}_R(x)$ be the ball of radius R centered at the site x .

Theorem 7.1. *There exists a constant $C = C(p) > 0$ such that the following holds for any domain V that contains $\mathcal{B}_R(o)$, where $R = Ch^{p-1}$ for $1 < p < 2$ and $R = Ch$ for $2 < p < \infty$. Let z be such that $\mathcal{B}_{R+1}(z) \subset V$. Then there exist absolute constants $c_0, c_1, c_2, c_3, c_4 > 0$ such that*

$$\exp \left[-c_0 \beta h^{(p-1) \wedge 1} \right] \leq \frac{\hat{\pi}_V^{(p)}(\phi_o = h)}{\hat{\pi}_V^{(p)}(\phi_o = h-1)} \leq \exp \left[-c_1 \beta h^{(p-1) \wedge 1} \right], \quad (7.4)$$

$$\exp \left[-c_2 \beta h^{p \wedge 2} \right] \leq \hat{\pi}_V^{(p)}(\phi_o = h) \leq \exp \left[-c_3 \beta h^{p \wedge 2} \right], \quad (7.5)$$

$$\hat{\pi}_V^{(p)}(\phi_z = h \mid \phi_o = h) \leq \exp \left[-c_4 \beta h^{p \wedge \frac{p}{p-1}} \right]. \quad (7.6)$$

Remark 7.2. In Remark 1.7 we defined the set \mathcal{B} of exceptional values of L (extending Remark 1.3 for the \mathbb{Z} GFF) as $\bigcup_{h \geq 1} \llbracket \frac{3}{4} L_h, L_h \rrbracket$ where $L_h = \lceil 5\beta/\hat{\pi}_\infty^{(p)}(\phi_o = h) \rceil$. For this choice, for every h we have $\sum_{k \in \mathcal{B} \cap \llbracket 1, L_h \rrbracket} \frac{1}{k} = O(h) = o(\log L_h)$ —i.e., \mathcal{B} has zero logarithmic density—as Eq. (7.5) shows $\lim_{h \rightarrow \infty} -\frac{1}{h} \hat{\pi}_\infty^{(p)}(\phi_o = h) = h = \infty$. (Precisely, $\frac{1}{\log n} \sum_{k \in \mathcal{B} \cap \llbracket 1, n \rrbracket} \frac{1}{k} = O((\log n)^{-(\frac{p-1}{p} \wedge \frac{1}{2})})$ for all $p > 1$.)

The extensions of the large deviation estimates in Eqs. (7.1) to (7.3) to a general domain V containing $\mathcal{B}_R(o)$ follow immediately from their proofs in [35], which we summarize below. Looking at the $1 < p < 2$ case, the proof begins by showing that the outermost 1 level-line loop containing the origin (denoted Γ_1 , not to be confused with the notation for animals used previously) has size at most Ch^{p-1} except with probability $e^{-\beta h^{p-1}}$, via a Peierls map argument which holds in any domain V . Hence we obtain

Claim 7.3 (extending [35, Lem. 5.3]). *Suppose $\phi_o \geq 1$. Let Γ_1 be the outermost 1 level-line loop surrounding the origin o . There exists a constant $C = C(p) > 0$ such that for any domain V containing o ,*

$$\hat{\pi}_V^{(p)}(|\Gamma_1| > Ch^{p-1} \mid \phi_o \geq h) \leq e^{-\beta h^{p-1}}.$$

This is used in [35, Lem. 5.4] to show the large deviation rate by comparing to the real-valued energy minimizer ϕ^* in a similar (but simpler) strategy as in the $p = 2$ case. The lower bound of Eq. (7.2) is established by taking the probability of a single ϕ which is an integer valued approximation to $h\phi^*$, and is supported on $\mathcal{B}_{R-1}(o)$, thus holding also on $\hat{\pi}_V^{(p)}$. The upper bound of Eq. (7.2) was proved by revealing Γ_1 and using monotonicity to set the boundary conditions outside Γ_1 to 0, and hence follows from Claim 7.3. The proof of Eq. (7.1) also uses the same revealing strategy, combined with a map argument which holds in any domain. Thus, we will have Theorem 7.1 for $1 < p < 2$ once we prove the lower bound in Eq. (7.4) (stated as Lemma 7.5) and the conditional probability bound in Eq. (7.6)¹¹.

¹¹Here we mention how the proof of Eqs. (7.1) and (7.2) extend to $\hat{\pi}_V^{(p)}$ for $1 < p < 2$; the same is valid for (7.3), but we will need a stronger version of that inequality (namely, Eq. (7.6)) anyway.

For the $2 < p < \infty$ case, the lower bound of Eq. (7.2) was again proven by taking a single ϕ^* with $\phi_o^* = h$ that has probability at least $e^{-c\beta h}$, which is already supported on $\mathcal{B}_R(o)$. Both the upper bound of Eq. (7.2) and the bound on the ratio of Eq. (7.1) then apply a map argument, which holds in any domain, together with [35, Claim 5.7], which refers only to the energy of a set of level-line loops. Thus, Theorem 7.1 for $2 < p < \infty$ will be established upon proving the lower bound in Eq. (7.4) (stated as Corollary 7.8) and Eq. (7.6). One extra ingredient which will be needed for this is to bound the effective support of solutions to the large deviation problem (see Lemma 7.7) in $\mathcal{B}_R(o)$ where $R = Ch$ for a sufficiently large absolute constant $C > 0$.

Remark 7.4. We emphasize that for extending the results of [35] for $2 < p < \infty$, only the lower bound of $\hat{\pi}_V^{(p)}(\phi_o = h)$ requires a restriction on V , which is that it must contain the support of a particular ϕ^* defined as $\phi^*(x) = (h - \|x\|_1) \vee 0$. In particular, we can take $R = \sqrt{2}h$, with no dependence on β . Then to prove the additional lower bound on Eq. (7.5) and the bound in Eq. (7.6), we increase R to some other constant Ch , as needed to capture the outermost 1 level-line loop. If instead R had the form $C\beta h$, then our proof of the upper bound for points near the boundary of V in Lemma 7.10 would fail.

Lemma 7.5. *Fix $1 < p < 2$, and V as in Theorem 7.1. There exists a constant $c > 0$ such that for all $h \geq 1$,*

$$\frac{\hat{\pi}_V^{(p)}(\phi_o = h)}{\hat{\pi}_V^{(p)}(\phi_o = h - 1)} \geq \exp[-c\beta h^{p-1}].$$

Remark 7.6. The proof of Lemma 7.5 is valid for all $p \geq 1$, though its lower bound of $\exp[-c\beta h^{p-1}]$ is weaker than the matching upper bounds of $\exp(-c\beta h/\log h)$ for $p = 2$ and $\exp(-c\beta h)$ for $p > 2$.

Proof of Lemma 7.5. It is easy to see that we have

$$\hat{\pi}_V^{(p)}(\min_{x \sim o} \phi_x < 0 \mid \phi_o = h) \leq \varepsilon_\beta.$$

Indeed, by the upper bound of Eq. (7.4), we have that $\hat{\pi}_V^{(p)}(\phi_o = h) = (1 + o(1))\hat{\pi}_V^{(p)}(\phi_o \geq h)$, so it suffices to prove the bound when the conditioning is instead on $\phi_o \geq h$. But by FKG, we can then forget about the conditioning entirely, whence the bound is immediate by a standard Peierls argument.

Now we can write

$$\hat{\pi}_V^{(p)}(\phi_o = h - 1) \leq (1 + \varepsilon_\beta)\hat{\pi}_V^{(p)}(\phi_o = h - 1, \min_{x \sim o} \phi_x \geq 0).$$

On this latter event, we can raise the height of the origin by 1, mapping to the set $\{\phi : \phi_o = h\}$, and this changes the weight of ϕ by a factor of at most $e^{4\beta(h^p - (h-1)^p)} \leq e^{4\beta ph^{p-1}}$, so that

$$\hat{\pi}_V^{(p)}(\phi_o = h - 1, \min_{x \sim o} \phi_x \geq 0) \leq e^{4\beta ph^{p-1}}\hat{\pi}_V^{(p)}(\phi_o = h). \quad \blacksquare$$

We now look to prove the lower bound of Eq. (7.4) for $2 < p < \infty$. For each configuration ϕ with $\phi_o = h$, let $\Gamma_i(\phi)$ for $1 \leq i \leq h$ be the outermost i level-line loop surrounding the origin o , dropping ϕ from the notation if it is clear from context. For a collection of such loops $\{\Gamma_i\}$, define the energy of the collection by

$$\mathcal{E}(\{\Gamma_i\}) = \sum_e \Delta_e^p,$$

where Δ_e denotes the number of Γ_i in the collection that contain the dual bond e . Let the weight of Γ_i with respect to $\{\Gamma_i\}$ be defined as

$$\mathcal{W}(\Gamma_i) = \sum_{e \in \Gamma_i} \Delta_e^p / \Delta_e = \sum_{e \in \Gamma_i} \Delta_e^{p-1}.$$

(The weight depends on the full collection $\{\Gamma_i\}$ and not just Γ_i , but we will omit this from the notation for simplicity as it will be obvious from the context.) Observe that $\mathcal{E}(\{\Gamma_i\}) = \sum_i \mathcal{W}(\Gamma_i)$, and that $|\Gamma_i| \leq \mathcal{W}(\Gamma_i)$.

Lemma 7.7. *Fix $p \geq 1$ and V as in Theorem 7.1. Then, there exists constants $C, c > 0$ such that for all $h \geq 1$, $r \geq Ch$,*

$$\hat{\pi}_V^{(p)}(\mathcal{W}(\Gamma_1) \geq r \mid \phi_o = h) < e^{-(\beta-c)(r-Ch)}.$$

Proof. By a standard Peierls argument (enumerate over the h level-line loops $\{\Gamma_i\}$, where a length- k loop intersects (say) a horizontal ray from the origin at distance at most k),

$$\hat{\pi}_V^{(p)}(\mathcal{E}(\{\Gamma_i\}) \geq C_1 h^2, \phi_o = h) = e^{-(C_1/2)\beta h^2},$$

and hence the lower bound on $\hat{\pi}_V^{(p)}(\phi_o = h)$ in Eq. (7.5) implies that

$$\hat{\pi}_V^{(p)}(\mathcal{E}(\{\Gamma_i\}) \geq C_1 h^2 \mid \phi_o = h) = e^{-\beta h^2} \quad (7.7)$$

provided that $C_1/2 > c_2 + 1$ for c_2 from Eq. (7.5). Hence, we can assume that $\mathcal{E}(\{\Gamma_i\}) \leq C_1 h^2$, whence by pigeonhole principle there exists some Γ_{i^*} such that $\mathcal{W}(\Gamma_{i^*}) \leq C_1 h$. Now consider the map T on the set $\{\phi : \phi_o = h, \mathcal{E}(\{\Gamma_i\}) \leq C_1 h^2\}$ that takes the collection $\{\Gamma_i\}$ in ϕ , deletes the outermost contour Γ_1 , and adds a copy of Γ_{i^*} . This map preserves the fact that the origin is at height h . Calling this resulting collection $\{\Gamma_i(T(\phi))\}$, we obtain that

$$\begin{aligned} -\frac{1}{\beta} \log \frac{\hat{\pi}_V^{(p)}(T(\phi))}{\hat{\pi}_V^{(p)}(\phi)} &= \mathcal{E}(\{\Gamma_i(T(\phi))\}) - \mathcal{E}(\{\Gamma_i(\phi)\}) \leq -\mathcal{W}(\Gamma_1(\phi)) + \sum_{e \in \Gamma_{i^*}(\phi)} (\Delta_e + 1)^p - \Delta_e^p \\ &\leq -\mathcal{W}(\Gamma_1(\phi)) + \sum_{e \in \Gamma_{i^*}(\phi)} p(\Delta_e + 1)^{p-1} \\ &\leq -\mathcal{W}(\Gamma_1(\phi)) + p 2^{p-1} C_1 h \\ &=: -\mathcal{W}(\Gamma_1(\phi)) + C_2 h. \end{aligned}$$

Moreover, the multiplicity of this map is such that for any given ψ in the image of T ,

$$|\{\phi : T(\phi) = \psi, \mathcal{W}(\Gamma_1(\phi)) = k\}| \leq hks^k,$$

for some absolute constant $s > 0$, since given the image ψ we can read off $T(\{\Gamma_i(\phi)\})$, and reconstruct $\{\Gamma_i(\phi)\}$ if we know what i^* and Γ_1 are. There are h choices of i^* , k choices for the length $|\Gamma_1| = j \leq k$, and $s^j \leq s^k$ choices of loops of length j which surround the origin. Hence, if $\mathcal{W}(\Gamma_1) \geq r \geq Ch$, we can apply a Peierls argument with this map to obtain that

$$\begin{aligned} \sum_{\substack{\phi: \phi_o = h, \\ \mathcal{E}(\{\Gamma_i\}) \leq C_1 h^2, \\ \mathcal{W}(\Gamma_1) \geq r}} \hat{\pi}_V^{(p)}(\phi) &\leq \sum_{k \geq r} \sum_{\psi \in \text{Im}(T)} \sum_{\substack{\phi: \phi_o = h, \\ \mathcal{E}(\{\Gamma_i\}) \leq C_1 h^2, \\ \mathcal{W}(\Gamma_1) = k, \\ T(\phi) = \psi}} \hat{\pi}_V^{(p)}(\psi) e^{-\beta(k-C_2 h)} \\ &\leq \sum_{k \geq r} \sum_{\psi \in \text{Im}(T)} \hat{\pi}_V^{(p)}(\psi) e^{-\beta(k-C_2 h)} hks^k \\ &\leq e^{-(\beta-C_3)(r-C_2 h)} \hat{\pi}_V^{(p)}(\phi_o = h). \end{aligned} \quad (7.8)$$

Combining Eqs. (7.7) and (7.8) concludes the proof. ■

The constant C used for $R = Ch$ in Theorem 7.1 should now be taken to be the maximum between $2C$ from Lemma 7.7 and the constant needed previously from [35] (which was just to ensure that the the single ϕ used to lower bound Eq. (7.2) is supported in $\mathcal{B}_R(o)$).

Corollary 7.8. *Fix $2 < p < \infty$ and let V be as in Theorem 7.1. There exists a constant $c > 0$ such that for all $h \geq 1$,*

$$\frac{\hat{\pi}_V^{(p)}(\phi_o = h)}{\hat{\pi}_V^{(p)}(\phi_o = h - 1)} \geq \exp[-c\beta h].$$

Proof. By Lemma 7.7, there exists C such that

$$\hat{\pi}_V^{(p)}(\phi_o = h - 1) = (1 + o(1))\hat{\pi}_V^{(p)}(\phi_o = h - 1, \mathcal{W}(\Gamma_1) \leq Ch).$$

On the latter event, consider the map $T(\phi) = \psi$ which raises the interior of $\Gamma_1(\phi)$ by 1. Then, we have

$$\begin{aligned} -\frac{1}{\beta} \log \frac{\hat{\pi}_V^{(p)}(T(\phi))}{\hat{\pi}_V^{(p)}(\phi)} &= \sum_{e \in \Gamma_1(\phi)} (\Delta_e + 1)^p - \Delta_e^p \\ &\leq p2^{p-1}\mathcal{W}(\Gamma_1(\phi)). \end{aligned}$$

As this map is injective and maps into the set $\{\phi : \phi_o = h\}$, we obtain

$$\hat{\pi}_V^{(p)}(\phi_o = h - 1, \mathcal{W}(\Gamma_1) \leq Ch) \leq \hat{\pi}_V^{(p)}(\phi_o = h)e^{-\beta C'h}$$

for $C' = p2^{p-1}C$. ■

Finally, to show the bound on the conditional probability in Eq. (7.6), we begin with a simple lemma comparing $\hat{\pi}_V^{(p)}$ to $\hat{\pi}_{\mathcal{B}_R(o)}^{(p)}$.

Lemma 7.9. *Let $V \supset \mathcal{B}_R(o)$ be as in Theorem 7.1. Then, there exists $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow \infty$ such that*

$$e^{-\varepsilon_\beta R} \leq \frac{\hat{\pi}_V^{(p)}(\phi_o = h)}{\hat{\pi}_{\mathcal{B}_R(o)}^{(p)}(\phi_o = h)} \leq e^{\varepsilon_\beta R}.$$

Proof. For any $V_1 \subset V_2$, it is a straightforward FKG argument (see [35, Eq. 3.14], or the proof of Lemma 2.8) to show that

$$\hat{\pi}_{V_2}^{(p)}(\phi_o \geq h) \geq e^{-\varepsilon_\beta |\partial V_1 \setminus \partial V_2|} \hat{\pi}_{V_1}^{(p)}(\phi_o \geq h). \quad (7.9)$$

Moreover, by a standard Peierls argument, (see, e.g., the proof of [35, Clm. 3.6]) one obtains that for any V , $h \geq 0$,

$$\hat{\pi}_V^{(p)}(\phi_o > h) \leq \varepsilon_\beta \hat{\pi}_V^{(p)}(\phi_o = h). \quad (7.10)$$

Combining the above two displays immediately implies the lower bound.

For the upper bound, by Lemma 7.7 it suffices to upper bound $\hat{\pi}_V^{(p)}(\mathcal{W}(\Gamma_1) \leq R, \phi_o = h)$ (noting the comment about the constants after Lemma 7.7). The event $\mathcal{W}(\Gamma_1) \leq R$ implies $|\Gamma_1| \leq R$, which in turn implies that there is a chain of sites with height $\phi_x \leq 0$ contained inside $\mathcal{B}_R(o)$ which surrounds the origin, whose length is at most R . By monotonicity, the heights along this chain can be raised to 0. That is, calling the interior of this chain Λ , we have argued that

$$\hat{\pi}_V^{(p)}(\mathcal{W}(\Gamma_1) \leq R, \phi_o = h) \leq \hat{\pi}_V^{(p)}(\mathcal{W}(\Gamma_1) \leq R, \phi_o \geq h) \leq \max_{\substack{o \in \Lambda \subset \mathcal{B}_R(o) \\ |\partial \Lambda| \leq R}} \hat{\pi}_\Lambda^{(p)}(\phi_o \geq h).$$

The upper bound now follows by applying Eq. (7.9) for $V_1 = \Lambda$, $V_2 = \mathcal{B}_R$, along with Eq. (7.10). ■

Proof of Theorem 7.1, Eq. (7.6). With the above lemma in hand (replacing Claim 2.7 from the case $p = 2$), the proof follows from the same argument used to establish Eq. (2.8) in Theorem 2.5, with the only difference that one defines the two events E_1 and E_2 for

$$\Delta := \left\lfloor \frac{c_1}{6^{p+1}} h^{1 \wedge \frac{1}{p-1}} \right\rfloor,$$

with $c_1 > 0$ from the upper bound on $\widehat{\pi}_V^{(p)}(\phi_o = h)/\widehat{\pi}_V^{(p)}(\phi_o = h-1)$, now given by (7.4). Note that for this choice of Δ we have

$$\Delta h^{(p-1) \wedge 1} \asymp \Delta^p \asymp h^{p \wedge \frac{p}{p-1}}.$$

As was argued in Eq. (2.9), one can iterate the upper bound of Eq. (7.4) and thereafter apply Lemma 7.9 to arrive at the analogue here:

$$\widehat{\pi}_V^{(p)}(E_1^c \mid \phi_o = h) \leq \exp \left(- (c_1 - o(1)) \beta \Delta h^{(p-1) \wedge 1} \right).$$

(NB. We apply Lemma 7.9 for $R = Ch^{(p-1) \wedge 1}$ whence the term $e^{\varepsilon_\beta R}$ from said lemma is absorbed in the $o(1)$ -term in the above exponent.) Proceeding as we did below Eq. (2.9), one then has

$$\widehat{\pi}_V^{(p)}(\phi_z = h \mid \phi_o = h, E_1, E_2^c) \leq (1 + \varepsilon_\beta) e^{-\beta(3\Delta^p + (5\Delta)^p)} = (1 + \varepsilon_\beta) e^{-(5^p + 3)\beta\Delta^p},$$

where the constants here, apart from the p -power of the interaction, were unchanged as we did not alter the definition of E_1 and E_2 (but only the definition of Δ). For the same reason we further get

$$\widehat{\pi}_V^{(p)}(\phi_z \geq h + \Delta \mid \phi_o = h, E_1, E_2) \geq (1 - \varepsilon_\beta) e^{-4\beta(6\Delta)^p}$$

and

$$\widehat{\pi}_V^{(p)}(\phi_z \geq h + \Delta \mid \phi_o = h, E_1, E_2) \leq \frac{\exp(-(c_1 - o(1))\beta\Delta h^{(p-1) \wedge 1})}{\widehat{\pi}_V^{(p)}(E_2 \mid \phi_o = h, E_1)},$$

combining to yield

$$\begin{aligned} \widehat{\pi}_V^{(p)}(E_2 \mid \phi_o = h, E_1) &\leq (1 + \varepsilon_\beta) \exp \left(\beta \left(4 \cdot 6^p \Delta - (c_1 - o(1)) h^{(p-1) \wedge 1} \right) \Delta \right) \\ &\leq \exp \left(- (c_1/3) \beta \Delta h^{(p-1) \wedge 1} \right), \end{aligned}$$

as required. ■

Analogously to Lemma 2.8 from the $p = 2$ case, as a consequence of extending the large deviation results to an arbitrary domain containing $\mathcal{B}_R(o)$, we obtain an upper bound on the probability that $\phi_x \geq h$ even for points x close to the boundary.

Lemma 7.10. *Fix $1 < p < 2$, or $2 < p < \infty$. There exists $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow \infty$ such that, for every $V \subset \mathbb{Z}^2$, $x \in V$, and h sufficiently large compared to β ,*

$$\widehat{\pi}_V^{(p)}(\phi_x \geq h) \leq \widehat{\pi}_\infty^{(p)}(\phi_o \geq h)^{1+\varepsilon_\beta}. \quad (7.11)$$

Proof. Consider first $1 < p < 2$. Here we take $R = Ch^{p-1}$. Let $V' = V \cup \mathcal{B}_R(x)$. By Eq. (7.9), we have

$$\widehat{\pi}_{V'}^{(p)}(\phi_x \geq h) \geq e^{-\varepsilon_\beta R^2} \widehat{\pi}_V^{(p)}(\phi_x \geq h),$$

using the fact that $\partial V \setminus \partial V' \subset \mathcal{B}_R(x)$ so that $|\partial V \setminus \partial V'| \leq \pi R^2$. Combining with Lemma 7.9 to compare both $\widehat{\pi}_V^{(p)}$ and $\widehat{\pi}_\infty^{(p)}$ with $\widehat{\pi}_{\mathcal{B}_R(o)}^{(p)}$, we obtain that

$$\widehat{\pi}_V^{(p)}(\phi_x \geq h) \leq \widehat{\pi}_\infty^{(p)}(\phi_x \geq h) e^{\varepsilon_\beta Ch^{2p-2}}.$$

By the lower bound of Eq. (7.5) and the fact as $h \rightarrow \infty$, $h^{2p-2} \ll h^p$ in the regime $1 < p < 2$, we have that

$$e^{\varepsilon_\beta Ch^{2p-2}} \leq \widehat{\pi}_\infty^{(p)}(\phi_o \geq h)^{o(1)},$$

where $o(1)$ is as $h \rightarrow \infty$. We conclude by combining the above two displays with the translation invariance of $\widehat{\pi}_\infty^{(p)}$.

For $2 < p < \infty$, we can apply the same logic as above with $R = Ch$ to obtain that

$$\widehat{\pi}_V^{(p)}(\phi_x \geq h) \leq \widehat{\pi}_\infty^{(p)}(\phi_x \geq h) e^{\varepsilon_\beta Ch^2}.$$

In this case, the error term is of the same order as the large deviation with respect to h (both are $\exp[O(h^2)]$). We instead use the ε_β term to ensure that the error is sufficiently small. That is, by the lower bound of Eq. (7.5) for $2 < p < \infty$, we have

$$e^{\varepsilon_\beta Ch^2} \leq \hat{\pi}_\infty^{(p)}(\phi_o \geq h)^{\varepsilon_\beta},$$

and the proof concludes by the translation invariance of $\hat{\pi}_\infty^{(p)}$. ■

We now have all the tools needed to prove the crucial bound on the probability of having a floor (analogous to Lemma 2.9), which results in the area term in the cluster expansion. Recall the definition of $H^{(p)}$ in Eq. (1.8).

Lemma 7.11. *Fix $1 < p < 2$ or $2 < p < \infty$. Fix $0 < \delta < \frac{1}{3}$ and $n \geq 0$, let $V \subset \mathbb{Z}^2$ be a connected region, let $F \subset V$ be a subset satisfying $|F| \leq L(\log L)^{50}$ and $|\partial F| = O(L^{1-\delta})$, and set $h = H^{(p)} - n$. Then, for β sufficiently large,*

$$\hat{\pi}_V^{(p)}(\phi_x \geq -h, \forall x \in F) = (1 + o(1)) \exp \left(-\hat{\pi}_\infty^{(p)}(\phi_o < -h)|F| \right).$$

Proof. The proof follows exactly as in Lemma 2.9, only plugging in the estimates above when needed. We summarize the crucial relationships between the exponents in Theorem 7.1 for the reader's convenience.

To lower bound the left hand side, we firstly need that the exponent in (the lower bound of) Eq. (7.4) is of a smaller order than in Eq. (7.5). This ensures that for each finite $h = H^{(p)}$ to $h = H^{(p)} - n$, $\hat{\pi}_\infty^{(p)}(\phi_o \geq h) = L^{-1+o(1)}$, a fact used throughout the proof. To control points near ∂F , by Lemma 7.10, we have that for all x ,

$$\hat{\pi}_V^{(p)}(\phi_x < -h) \leq L^{-1+\varepsilon_\beta},$$

which replaces Eq. (2.18). The ε_β in place of $o(1)$ is not an issue as we only needed something smaller than δ for the application in Eq. (2.19) (though it does mean that for every choice of δ , we need to take a sufficiently large β). The rest of the proof of the lower bound just involved FKG and coupling to infinite volume, holding exactly in the same manner.

The upper bound was a more involved “grill” estimate. The key equation to prove (analogous to Eq. (2.24)) is

$$\hat{\pi}_V^{(p)} \left(\mathfrak{D}_s \cap \bigcap_{x \in Q'_{i_j}} \{\phi_x \geq -h\} \mid \mathcal{F}_{j-1} \right) \leq \exp \left(- \left(1 - e^{-c\beta h^{p \wedge \frac{p}{p-1}}} \right) \hat{\pi}_\infty^{(p)}(\phi_o < -h) |Q'_{i_j}| \right), \quad (7.12)$$

where the definitions of $\mathfrak{D}_s, Q'_{i_j}, \mathcal{F}_j$ all remain the same. The exponent is motivated by the bound on the conditional probability in Eq. (7.6). Indeed, the justification for the above equation is to use Bonferroni's inequalities, with the main term to estimate being

$$\sum_{\substack{x, y \in Q'_{i_j} \\ 0 < \text{dist}(x, y) \leq \log^2 L}} \hat{\pi}_\infty^{(p)}(\phi_x < -h, \phi_y < -h) \leq (\log L)^4 \exp(-c\beta h^{p \wedge \frac{p}{p-1}}) \hat{\pi}_\infty^{(p)}(\phi_o < -h) |Q'_{i_j}|,$$

and this follows directly from Eq. (7.6), Eq. (7.5), and the lower bound on Eq. (7.4). We highlight that when justifying this point, as was done in the proof of Lemma 2.9, one needs that

- (a) the probability $\exp[-c\beta h^{(p-1) \wedge 1}]$ given in (the lower bound of) Eq. (7.4) is of a larger order than the probability $\exp[-c\beta h^{p \wedge \frac{p}{p-1}}]$ from Eq. (7.6), and is thus absorbed (see Eq. (2.25));
- (b) for $h = H^{(p)} - n$, which satisfies $h \asymp (\beta \log L)^{1/(p \wedge 2)}$ by Eq. (7.5), the same $\exp[-c\beta h^{(p-1) \wedge 1}]$ term from the (lower bound of) Eq. (7.4) outweighs the $(\log L)^c$ pre-factor.

The remaining computations are unchanged. ■

We can now establish the cluster expansion law. We use the same definitions of disagreement polymers as per Definitions 2.1 and 2.2, with the exception that the energy of γ is now defined to be

$$\mathcal{E}_\beta(\gamma) = \beta \sum_{e \in \gamma} |(\nabla \phi)_e|^p.$$

With this definition, we have the following analog of Proposition 2.3.

Proposition 7.12 (Cluster expansion with a floor for $|\nabla \phi|^p$). *In the setting of the $|\nabla \phi|^p$ model for fixed $p > 1$, with $\pi_V^{(p),\xi}$ replacing π_V^ξ and \mathcal{E}_β as above, the statement of Proposition 2.3 holds unchanged.*

Proof. The proof is identical as in Proposition 2.3. The cluster expansion without a floor only requires that $\mathcal{E}_\beta(\gamma) \geq \mathcal{N}(\gamma)$, where $\mathcal{N}(\gamma)$ is the number of bonds in γ , and we have this. Obtaining from this the cluster expansion with a floor requires only the area estimate of Lemma 7.11. ■

Proof of Theorem 1.5. With Proposition 7.12 replacing Proposition 2.3, we can proceed through Sections 3 to 6 just as in the case $p = 2$, since the only property needed from the polymer model was that the energy $\mathcal{E}_\beta^*(\gamma)$ satisfies Eq. (3.8). Otherwise, the only properties we used of the actual measure π_Λ were the FKG inequality and basic Peierls arguments. We simply highlight that in part (b) of Theorem 4.4, the initial upper bound on the displacement of the $(H - n - 1)$ level line is obtained by the estimate of the ratio N_{n+1}/N_n coming directly from the upper bound in Eq. (2.6), now replaced with the upper bound in Eq. (7.4) at $h = H^{(p)} - n$. What is needed is that this ratio is smaller than $(\log L)^\kappa$ for any $\kappa > 0$, which holds for all $1 < p < \infty$. ■

APPENDIX A. MONOTONICITY OF $|\nabla \phi|^p$ MODELS WITH NON-UNIFORM FLOORS AND CEILINGS

It is well-known that the $|\nabla \phi|^p$ measures enjoy the FKG inequality for any $p \geq 1$; we will use this fact in the presence of non-uniform floors and ceilings, and we include the short proof for completeness.

Claim A.1. *Let $\pi_V^{(p),\xi,\mathbf{a},\mathbf{b}}$ be the $|\nabla \phi|^p$ model ($p \in [1, \infty)$ and $\beta > 0$) on a domain V with boundary conditions ξ and arbitrary floors and ceilings $\mathbf{a}, \mathbf{b} \in (\mathbb{Z} \cup \{\pm\infty\})^V$ (i.e., $\mathbf{a}_x \leq \phi_x \leq \mathbf{b}_x$ for all $x \in V$). This measure satisfies the FKG lattice condition. In particular, if $\xi \leq \xi'$, $\mathbf{a} \leq \mathbf{a}'$ and $\mathbf{b} \leq \mathbf{b}'$ pointwise, then $\pi_V^{(p),\xi,\mathbf{a},\mathbf{b}} \preceq \pi_V^{(p),\xi',\mathbf{a}',\mathbf{b}'}$.*

Proof. Write $\pi = \pi_V^{(p),\xi,\mathbf{a},\mathbf{b}}$ for brevity. It suffices (see, e.g., [23, Thm. 2.22]) to verify the condition

$$\pi(\phi_1 \vee \phi_2) \pi(\phi_1 \wedge \phi_2) \geq \pi(\phi_1) \pi(\phi_2) \quad (\text{A.1})$$

for two configurations ϕ_1, ϕ_2 that differ on exactly 2 sites $x, y \in V$. We may further assume that x, y are nearest-neighbors (or there is equality in Eq. (A.1), as all interactions are nearest-neighbor). Writing $a_i := \phi_i(x)$ and $b_i := \phi_i(y)$, we may further assume without loss of generality that $a_1 > a_2$ and $b_1 < b_2$ (or else again there is trivially an equality). The above then translates into showing

$$|a_1 - b_2|^p + |a_2 - b_1|^p \leq |a_1 - b_1|^p + |a_2 - b_2|^p.$$

Writing $k := a_1 - b_1$, $\Delta_a := a_1 - a_2 > 0$ and $\Delta_b := b_2 - b_1 > 0$, this is equivalent to

$$|k - \Delta_b|^p + |k - \Delta_a|^p \leq |k - \Delta_a - \Delta_b|^p + |k|^p.$$

If $\Delta_a > \Delta_b$ then we have $[(|k - \Delta_a|^p - |k - \Delta_a - \Delta_b|^p)/\Delta_b] < [|k|^p - |k - \Delta_b|^p]/\Delta_b$ by the monotonicity of slopes due to the convexity of the function $x \mapsto |x|^p$ for $p \geq 1$. Similarly, when $\Delta_b > \Delta_a$ the monotonicity of slopes implies $[(|k - \Delta_b|^p - |k - \Delta_a - \Delta_b|^p)/\Delta_a] < [|k|^p - |k - \Delta_a|^p]/\Delta_a$, thus Eq. (A.1) holds in both cases. The final conclusion of the claim is a consequence of FKG in $\pi_V^{(p),\xi,\mathbf{a},\mathbf{b}'}$ if $\xi = \xi'$, whereas the statement for $\xi \neq \xi'$ follows as we may implement ξ via $\mathbf{a}_x = \mathbf{b}_x = \xi_x$ for all $x \in \partial V$ (in the measure on $V \cup \partial V$ with, say, 0 boundary conditions). ■

APPENDIX B. RANDOM BOUNDARY CONSTRUCTION

In this section, we prove Lemmas 4.7 and 6.2. We begin with the following fact, which follows easily by a Peierls argument.

Observation B.1 (e.g., [35, Eq. (4.2)]). *Let \mathcal{C}_x^i be the connected component of x of sites y with $\phi_y \leq H - n - i$. W.h.p. under π_V^k , the interior of \mathfrak{L}_{n+i-1} does not contain x with $|\mathcal{C}_x^i| > (\log L)^2$.*

Proof of Lemma 4.7. Start by revealing all of $\phi|_{\Lambda_L \setminus R}$. Now on R , we will continue to reveal ϕ at certain sites and lower some of the revealed heights in order to obtain Q, ξ in the manner described below, and this defines the π_V^k -measurable distribution on Q .

As the procedure will reveal all of $\phi|_{Q^c}$, the boundary conditions ξ leave two cases to consider: first, that \mathfrak{L}_n does not intersect R at all. Then, either R is entirely in the interior of \mathfrak{L}_n or entirely in the exterior, and the requirement that $A, B \in \partial R$ are in the interior of \mathfrak{L}_n rules out the latter possibility. Hence, \mathcal{A}_1 is all of Q , so $\mathcal{A}_2 \subset \mathcal{A}_1$ deterministically. In the second case, \mathfrak{L}_n intersects R , in which case ξ implies that the restriction $\mathfrak{L}_n|_Q$ coincides with the $H - n$ level-line in π_Q^ξ . In this case, the law of \mathfrak{L}_n has also been changed due to the lowering of heights in the revealing procedure. However, since the area inside a level-line is an increasing function of ϕ , we obtain that $\mathcal{A}_2 \subset \mathcal{A}_1$ by monotonicity.

The procedure to construct Q, ξ is as follows:

- (1) Call $\partial^1 R, \partial^2 R$ the top and bottom arcs of ∂R from A to B , respectively. For every x on $\partial^1 R$ with $\phi_x \leq H - n - 1$, reveal the connected component \mathcal{C}_x^1 of sites y with $\phi_y \leq H - n - 1$ containing x . Similarly, for every x on $\partial^2 R$ such that $\phi_x \leq H - n - 2$, reveal the connected component \mathcal{C}_x^2 of sites y with $\phi_y \leq H - n - 2$ containing x .
- (2) By maximality, the sites in the outer boundary of \mathcal{C}_x^1 have height $\geq H - n$, which we can drop by monotonicity to be exactly $H - n$. Similarly, the sites in the outer boundary of the \mathcal{C}_x^2 have height $\geq H - n - 1$, which we can drop to be exactly $H - n - 1$.
- (3) The remaining sites in $\partial^1 R$ all have height $\geq H - n$, which can be lowered to $H - n$ by monotonicity. Similarly, the remaining sites in $\partial^2 R$ have height $\geq H - n - 1$, which can be lowered to $H - n - 1$.
- (4) By construction, every $x \in \partial^1 R$ is now either at height $H - n$, or enclosed in a boundary of sites with height $H - n$. Similarly, every $x \in \partial^2 R$ is at height $H - n - 1$ or enclosed in a boundary of sites with height $H - n - 1$. Hence, the union of these sites at height $H - n$ and $H - n - 1$ enclose a simply connected domain Q satisfying the conditions of Lemma 4.7 except for Item 3. (Here, we use Observation B.1 to ensure Item 2.)
- (5) To obtain Item 3, let $A^{(0)} = (A_1^{(0)}, A_2^{(0)})$ be the rightmost point on the left side of ∂Q such that $A_2^{(0)} = A_2 + (\log L)^5 + (\log L)^2$. We start the following iterative procedure starting with $A^{(0)}$. Look at the $(\log L)^2 \times 2(\log L)^3$ rectangle with $A^{(0)} + (1, 0)$ as the midpoint of the left side of the rectangle. If there is no point in ∂Q which intersects the rectangle, then set $A' = A^{(0)}$. Otherwise, let $A^{(1)}$ be the the rightmost intersection of ∂Q with the rectangle. Then, repeat starting with $A^{(1)}$, continuing until we set $A' = A^{(k)}$ for some k . This procedure must stop after at most $(\log L)^2$ steps because each $A^{(i+1)}$ is to the right of $A^{(i)}$, while ∂Q only deviates at most $(\log L)^2$ distance from the left side of ∂R . Moreover, the change in the y -coordinate from $A^{(i)}$ to $A^{(i+1)}$ is at most $(\log L)^3$ at each step, so that $|A'_2 - A_2^{(0)}| \leq (\log L)^5$. In particular, $A'_2 \geq A_2 + (\log L)^2$, so that by Observation B.1, A_2 lies on the arc of ∂Q with boundary condition $H - n$.
- (6) Repeat the above procedure analogously on the right side of Q to obtain the point B . Finally, lower the boundary conditions from $H - n$ to $H - n - 1$ on the arcs of Q between A and A' , and B and B' .

The fact that $\mathcal{C}(\overline{A'B'}) \subset Q$ follows because of the assumption that A, B are at least distance $\sqrt{\ell_1}(\log L)^2$ away from the top or bottom of R , and the shifts by up to $(\log L)^5$ to A', B' are irrelevant as $\ell_1 \gg (\log L)^6$. It is straightforward to check that the remaining conditions in Item 3 are now satisfied by A', B' . \blacksquare

We turn now to prove Lemma 6.2, which is essentially the same as Lemma 4.7 but with reversed monotonicity. The minor complication is that the C_x components we wish to reveal now will now be of the form $\{y : \phi_y \geq H - n + 1\}$. The Peierls argument that deletes the contour ∂C_x would now want to shift the sites in C_x down, but this may run into problems with the conditioning that $\phi \geq 0$. This was handled in [35], which we recall here.

Observation B.2 ([35, Theorem 2]). *With high probability under $\pi_{\Lambda_L}^0$, for each $h = 0, \dots, H$, there is at most one h level-line of size $> (\log L)^2$, and there is no $H + 1$ level-line of size $> (\log L)^2$. In particular, if C_x is the connected component of x of sites y with $\phi_y \geq H - n + 1$, then for $n \geq 1$, w.h.p. the exterior of \mathfrak{L}_{n-1} does not contain x with $|C_x| > (\log L)^2$. For $n = 0$, w.h.p. there are no points x with $|C_x| > (\log L)^2$.*

Proof of Lemma 6.2. As described in the beginning of the proof of Lemma 4.7, by monotonicity it suffices to construct Q, ξ via a procedure which reveals $\phi|_{Q^c}$ under $\pi_{\Lambda_L}^0$ raises some of the revealed heights. The procedure to construct Q, ξ is as follows:

- (1) Call $\partial^1 R$ the top and sides of ∂R . For every x on $\partial^1 R$ with $\phi_x \geq H - n + 1$, reveal the connected component C_x of sites y with $\phi_y \geq H - n + 1$ containing x .
- (2) By maximality, the sites in the outer boundary of C_x have height $\leq H - n$, which we can raise by monotonicity to be exactly $H - n$.
- (3) Call $\partial^2 R$ the bottom side of ∂R in between A, B . Call $\partial^3 R$ the remainder of the bottom side of ∂R . Since all the boundary heights on $\partial^2 R \cup \partial^3 R$ are 0, by monotonicity we can raise the boundary condition on $\partial^2 R$ to $H - n - 1$ and on $\partial^3 R$ to $H - n$.
- (4) By construction, every $x \in \partial^1 R$ is now either at height $H - n$, or enclosed in a boundary of sites with height $H - n$. Hence, the union of these sites at height $H - n$ with $\partial^2 R$ and $\partial^3 R$ enclose a simply connected domain Q satisfying the conditions of Lemma 6.2, using Observation B.2 to ensure Item 2. \blacksquare

Acknowledgments. This research was supported by the NSF grant DMS-2451083.

REFERENCES

- [1] M. Basu Roy Chowdhury, P. Caputo, and S. Ganguly. Characterizing Gibbs states for area-tilted Brownian lines. *Ann. Probab.* To appear.
- [2] R. Bauerschmidt, J. Park, and P.-F. c. Rodriguez. The Discrete Gaussian model, I. Renormalisation group flow at high temperature. *Ann. Probab.*, 52(4):1253–1359, 2024.
- [3] R. Bauerschmidt, J. Park, and P.-F. c. Rodriguez. The discrete Gaussian model, II. Infinite-volume scaling limit at high temperature. *Ann. Probab.*, 52(4):1360–1398, 2024.
- [4] R. Brandenberger and C. E. Wayne. Decay of correlations in surface models. *J. Statist. Phys.*, 27(3):425–440, 1982.
- [5] J. Bricmont, A. El Mellouki, and J. Fröhlich. Random surfaces in statistical mechanics: roughening, rounding, wetting, *J. Statist. Phys.*, 42(5-6):743–798, 1986.
- [6] W. K. Burton, N. Cabrera, and F. C. Frank. The growth of crystals and the equilibrium structure of their surfaces. *Philos. Trans. Roy. Soc. London Ser. A*, 243:299–358, 1951.
- [7] P. Caddeo, Y. H. Kim, and E. Lubetzky. On level line fluctuations of SOS surfaces above a wall. *Forum Math. Sigma*, 12:Paper No. e91, 59, 2024.
- [8] M. Campanino, D. Ioffe, and Y. Velenik. Ornstein-Zernike theory for finite range Ising models above T_c . *Probab. Theory Related Fields*, 125(3):305–349, 2003.
- [9] P. Caputo and S. Ganguly. Uniqueness, mixing, and optimal tails for Brownian line ensembles with geometric area tilt. *Probab. Math. Phys.*, 6(1):195–239, 2025.
- [10] P. Caputo, D. Ioffe, and V. Wachtel. Confinement of Brownian polymers under geometric area tilts. *Electron. J. Probab.*, 24:Paper No. 37, 21, 2019.

- [11] P. Caputo, D. Ioffe, and V. Wachtel. Tightness and line ensembles for Brownian polymers under geometric area tilts. In *Statistical mechanics of classical and disordered systems*, volume 293 of *Springer Proc. Math. Stat.*, pages 241–266. Springer, Cham, 2019.
- [12] P. Caputo, E. Lubetzky, F. Martinelli, A. Sly, and F. L. Toninelli. The shape of the $(2+1)D$ SOS surface above a wall. *C. R. Math. Acad. Sci. Paris*, 350(13-14):703–706, 2012.
- [13] P. Caputo, E. Lubetzky, F. Martinelli, A. Sly, and F. L. Toninelli. Dynamics of $(2+1)$ -dimensional SOS surfaces above a wall: Slow mixing induced by entropic repulsion. *Ann. Probab.*, 42(4):1516–1589, 2014.
- [14] P. Caputo, E. Lubetzky, F. Martinelli, A. Sly, and F. L. Toninelli. Scaling limit and cube-root fluctuations in SOS surfaces above a wall. *J. Eur. Math. Soc. (JEMS)*, 18(5):931–995, 2016.
- [15] S. T. Chui and J. D. Weeks. Phase transition in the two-dimensional coulomb gas, and the interfacial roughening transition. *Phys. Rev. B*, 14:4978–4982, Dec 1976.
- [16] A. Dembo, E. Lubetzky, and O. Zeitouni. On the limiting law of line ensembles of Brownian polymers with geometric area tilts. *Ann. Inst. Henri Poincaré Probab. Stat.*, 60(1):113–125, 2024.
- [17] R. Dobrushin, R. Kotecký, and S. Shlosman. *Wulff construction*, volume 104 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1992. A global shape from local interaction, Translated from the Russian by the authors.
- [18] P. L. Ferrari and H. Spohn. Constrained Brownian motion: fluctuations away from circular and parabolic barriers. *Ann. Probab.*, 33(4):1302–1325, 2005.
- [19] J. Fröhlich and T. Spencer. Kosterlitz-Thouless transition in the two-dimensional plane rotator and Coulomb gas. *Phys. Rev. Lett.*, 46(15):1006–1009, 1981.
- [20] J. Fröhlich and T. Spencer. The Kosterlitz-Thouless transition in two-dimensional abelian spin systems and the Coulomb gas. *Comm. Math. Phys.*, 81(4):527–602, 1981.
- [21] R. Gheissari and E. Lubetzky. Approximate domain Markov property for rigid Ising interfaces. *J. Stat. Phys.*, 190(5):99, 2023.
- [22] L. Greenberg and D. Ioffe. On an invariance principle for phase separation lines. *Ann. Inst. H. Poincaré Probab. Statist.*, 41(5):871–885, 2005.
- [23] G. Grimmett. *The random-cluster model*, volume 333 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2006.
- [24] M. Hegde, Y. H. Kim, and C. Serio. Scaling limit and tail bounds for a random walk model of SOS level lines, 2025. Preprint, [arXiv:2502.10384](https://arxiv.org/abs/2502.10384).
- [25] D. Ioffe, S. Ott, S. Shlosman, and Y. Velenik. Critical prewetting in the 2D Ising model. *Ann. Probab.*, 50(3):1127–1172, 2022.
- [26] D. Ioffe, S. Ott, Y. Velenik, and V. Wachtel. Invariance principle for a Potts interface along a wall. *J. Stat. Phys.*, 180(1-6):832–861, 2020.
- [27] D. Ioffe, S. Shlosman, and F. L. Toninelli. Interaction versus entropic repulsion for low temperature Ising polymers. *J. Stat. Phys.*, 158(5):1007–1050, 2015.
- [28] D. Ioffe, S. Shlosman, and Y. Velenik. An invariance principle to Ferrari-Spohn diffusions. *Comm. Math. Phys.*, 336(2):905–932, 2015.
- [29] D. Ioffe and Y. Velenik. Ballistic phase of self-interacting random walks. In *Analysis and stochastics of growth processes and interface models*, pages 55–79. Oxford Univ. Press, Oxford, 2008.
- [30] J. M. Kosterlitz. Kosterlitz–thouless physics: a review of key issues. *Reports on Progress in Physics*, 79(2):026001, jan 2016.
- [31] R. Kotecký and D. Preiss. Cluster expansion for abstract polymer models. *Comm. Math. Phys.*, 103(3):491–498, 1986.
- [32] P. Lammers. A dichotomy theory for height functions, 2022. Preprint, [arXiv:2211.14365](https://arxiv.org/abs/2211.14365).
- [33] P. Lammers and S. Ott. Delocalisation and absolute-value-FKG in the Solid-On-Solid model. *Probab. Theory Related Fields*, 188(1-2):63–87, 2024.
- [34] B. Laslier and E. Lubetzky. Tilted Solid-On-Solid is liquid: scaling limit of SOS with a potential on a slope, 2024. Preprint, [arXiv:2409.08745](https://arxiv.org/abs/2409.08745).
- [35] E. Lubetzky, F. Martinelli, and A. Sly. Harmonic pinnacles in the discrete Gaussian model. *Comm. Math. Phys.*, 344(3):673–717, 2016.
- [36] F. Martinelli and F. L. Toninelli. On the mixing time of the 2D stochastic Ising model with “plus” boundary conditions at low temperature. *Comm. Math. Phys.*, 296(1):175–213, 2010.
- [37] S. Ott and F. Schweiger. Quantitative delocalization for Solid-On-Solid models at high temperature and arbitrary tilt, 2025. Preprint, [arxiv:2505.16804](https://arxiv.org/abs/2505.16804).
- [38] V. Richter. Multi-dimensional local limit theorems for large deviations. *Theory of Probability & Its Applications*, 3(1):100–106, 1958.
- [39] C. Serio. Scaling limit for line ensembles of random walks with geometric area tilts. *Electron. J. Probab.*, 28:Paper No. 124, 14, 2023.

- [40] S. Sheffield. Random surfaces. *Astérisque*, (304):vi+175, 2005.
- [41] J. Villain. Theory of one- and two-dimensional magnets with an easy magnetization plane. II. The planar, classical, two-dimensional magnet. *Journal de Physique*, 36(6):581–590, 1975.

J. CHEN

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY 10012, USA.

Email address: `jlc871@courant.nyu.edu`

E. LUBETZKY

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY 10012, USA.

Email address: `eyal@courant.nyu.edu`