

ENTROPIC REPULSION OF 3D ISING INTERFACES CONDITIONED TO STAY ABOVE A FLOOR

REZA GHEISSARI AND EYAL LUBETZKY

ABSTRACT. We study the interface of the Ising model in a box of side-length n in \mathbb{Z}^3 at low temperature $1/\beta$ under Dobrushin's boundary conditions, conditioned to stay in a half-space above height $-\mathfrak{h}$ (a hard floor). Without this conditioning, Dobrushin showed in 1972 that typically most of the interface is flat at height 0. With the floor, for small \mathfrak{h} , the model is expected to exhibit *entropic repulsion*, where the typical height of the interface lifts off of 0. Detailed understanding of the SOS model—a more tractable height function approximation of 3D Ising—due to Caputo et al., suggests that there is a single integer value $-h_n^* \sim -c \log n$ of the floor height, delineating the transition between rigidity at height 0 and entropic repulsion.

We identify an explicit $h_n^* = (c_* + o(1)) \log n$ such that, for the typical Ising interface above a hard floor at $-\mathfrak{h}$, all but an $\varepsilon(\beta)$ -fraction of the sites are propelled to be above height 0 if $\mathfrak{h} < h_n^* - 1$, whereas all but an $\varepsilon(\beta)$ -fraction of the sites remain at height 0 if $\mathfrak{h} \geq h_n^*$. Further, c_* is such that the typical height of the unconditional maximum is $(2c_* + o(1)) \log n$; this confirms scaling predictions from the SOS approximation.

1. INTRODUCTION

Let μ_n^\mp be the distribution of the Ising model with Dobrushin's boundary conditions on the 3D cylinder,

$$\Lambda_n = \llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times \mathbb{Z} = \{-\lfloor \frac{n}{2} \rfloor, \dots, \lceil \frac{n}{2} \rceil\}^2 \times \mathbb{Z}.$$

More precisely, let μ_n^\mp be the distribution over assignments of ± 1 spins to the cells of Λ_n , denoted $\mathcal{C}(\mathbb{Z}^3)$. The probability of a configuration σ is proportional to $\exp(-\beta \mathcal{H}(\sigma))$, where \mathcal{H} counts the number of disagreeing neighboring cells under \mp -boundary conditions that are minus in the upper-half space $\{(x_1, x_2, x_3) : x_3 > 0\}$ and plus in the lower-half space. The parameter β is an inverse temperature parameter which in our context will be a large enough constant. The weak limit of μ_n^\mp as $n \rightarrow \infty$ is denoted $\mu_{\mathbb{Z}^3}^\mp$, which was the famous example of Dobrushin [16] for a Gibbs measure on \mathbb{Z}^3 which is not translational invariant in every direction.

For an Ising configuration $\sigma \sim \mu_n^\mp$ with Dobrushin boundary conditions, its interface is defined as follows: consider every f in $\mathcal{F}(\mathbb{Z}^3)$ (the faces of \mathbb{Z}^3) that separates two disagreeing spins of σ , and let the *interface* \mathcal{I} be the $*$ -connected component of such faces incident to the $\{x_3 = 0\}$ plane outside Λ_n (two faces are $*$ -connected if they have a common bounding vertex). Dobrushin [16] showed that for $\beta > 0$ large enough, with high probability (w.h.p.) the interface \mathcal{I} is *rigid* (hence the conclusion on $\mu_{\mathbb{Z}^3}^\mp$) at height 0; that is, for some fixed $\varepsilon_\beta > 0$ going to 0 as β goes to infinity,

$$|\mathcal{I} \cap (\llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times \{0\})| \geq (1 - \varepsilon_\beta)n^2, \quad \text{w.h.p. over } \mathcal{I} \sim \mu_n^\mp.$$

In this work we consider the *entropic repulsion* effect of a *hard floor* constraint for μ_n^\mp : namely, let

$$\mathbf{I}_\mathfrak{h}^\text{fl} = \{\mathcal{I} \subset \llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times [-\mathfrak{h}, \infty)\} \tag{1.1}$$

be the event that the interface is everywhere above height $-\mathfrak{h}$, and define the measure “with a floor at $-\mathfrak{h}$,”

$$\widehat{\mu}_n^\mathfrak{h} = \mu_n^\mp(\cdot \mid \mathbf{I}_\mathfrak{h}^\text{fl}) \quad \text{for } \mathfrak{h} \geq 0. \tag{1.2}$$

The hard constraint of confining \mathcal{I} to a half-space above height $-\mathfrak{h}$ creates an energy vs. entropy competition: wherever the interface is to remain flat at height 0, its downward oscillations would be capped at depth \mathfrak{h} . Indeed, in several other more tractable models, a conditioning of this sort was shown to propel the interface above height 0 (see §1.1 for related work). Of course, one expects this to hold only for small values of \mathfrak{h} , whereas for large enough \mathfrak{h} , the effect of a hard floor at height $-\mathfrak{h}$ should be unnoticeable, and the interface should remain flat at height 0 (this would certainly be the case if one should take \mathfrak{h} to be larger than the height of the global maximum in a typical interface $\mathcal{I} \sim \mu_n^\mp$).

Our main result, Theorem 1.1, identifies the exact \mathfrak{h} delineating this phase transition: for any β large enough, we identify a critical integer h_n^* such that, for all $\mathfrak{h} < h_n^*$, a typical interface $\mathcal{I} \sim \widehat{\mu}_n^\mathfrak{h}$ would have (say) at most a 0.01 fraction of the faces of \mathcal{I} be at height 0, whereas for $\mathfrak{h} > h_n^*$ at least (say) 0.99 of the points would be such. This height is defined by

$$h_n^* = h_n^*(\beta) := \inf\{\mathfrak{h} \geq 1 : \alpha_n > \log n - 2\beta\}, \tag{1.3}$$




FIGURE 1. An Ising interface in Λ_n for $n = 500$ conditioned to lie above a floor at height $h = 0$.

in which the quantities $\alpha_h = \alpha_h(\beta)$ are given by

$$\alpha_h = -\log \mu_{\mathbb{Z}^3}^{\mp} \left(\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right) \xleftrightarrow[\mathbb{R}^2 \times [0, \infty)]{+} (\mathbb{Z} + \frac{1}{2})^2 \times \{h - \frac{1}{2}\} \right), \quad (1.4)$$

where $v \xleftrightarrow[A]{+} w$ denotes there is a path of adjacent or diagonally adjacent plus spins between v and w in A .

Theorem 1.1. *There exist $\beta_0, c_0 > 0$ so that, if $\varepsilon_\beta = 1/(c_0 e^\beta)$ and h_n^* is as in (1.3), then for every $\beta > \beta_0$, with probability at least $1 - \exp(-\varepsilon_\beta n)$ the interface $\mathcal{I} \sim \hat{\mu}_n^h$ satisfies*

$$|\mathcal{I} \cap (\llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times \{0\})| < \varepsilon_\beta n^2 \quad \text{if } h < h_n^* - 1, \quad (1.5)$$

$$|\mathcal{I} \cap (\llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times \{0\})| > (1 - \varepsilon_\beta) n^2 \quad \text{if } h \geq h_n^*. \quad (1.6)$$

Furthermore, (1.5) also holds if we replace $\llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times \{0\}$ by $\llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times (-\infty, h_n^* - h - 1)$.

While one may easily be convinced that the hard floor constraint in $\hat{\mu}_n^h$ will become unnoticeable once h exceeds the typical value of M_n , the maximum height of the unconstrained interface $\mathcal{I} \sim \mu_n^\mp$ (which, by reflection symmetry, is the same as the typical value of the minimum height, whence \mathcal{I} will stay in $\llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times [-h, \infty)$ w.h.p.), in fact h_n^* is asymptotically 1/2 of that value, as the next remark states.

Remark 1.2. In [22], the authors derived a law of large numbers for the maximum height M_n of $\mathcal{I} \sim \mu_n^\mp$: as $n \rightarrow \infty$ one has $M_n / \log n \rightarrow 2/\alpha$ in μ_n^\mp -probability, for the quantity $\alpha = \alpha(\beta)$ defined there as

$$\alpha := \lim_{h \rightarrow \infty} \frac{\alpha_h}{h} \in [4\beta - C, 4\beta + e^{-4\beta}], \quad (1.7)$$

with the α_h 's as given in (1.4) (the tightness, and moreover Gumbel tails, of $M_n - \mathbb{E}[M_n]$ were proved in [21]). Theorem 1.1 shows that the critical h for repulsion in $\hat{\mu}_n^h$ is $h_n^* = (\frac{1}{\alpha} + o(1)) \log n$, which is $\frac{1+o(1)}{2} \mathbb{E}[M_n]$ for $\mathcal{I} \sim \mu_n^\mp$, in line with known scaling results for the SOS approximation of the 3D Ising model (see §1.1).

Akin to the behavior of the SOS model, the case $h = h_n^* - 1$ can depend on n , as the next remark details.

Remark 1.3. Theorem 1.1 extends to treat $h = h_n^* - 1$ whenever the deterministic quantity $\lambda_n := \log n - \alpha_{h_n^*}$, which is known to belong to $[-2\beta - \varepsilon_\beta, 2\beta]$ for all n (the upper bound is by definition (1.3) whereas the lower bound is by results in [22]) does not fall in a certain ε_β -fraction of this interval. For instance, (1.5) extends to $h = h_n^* - 1$ for $\lambda_n \geq 2 \log \beta$, while (1.6) extends to $h = h_n^* - 1$ for $\lambda_n \leq \log \beta$ (see Remarks 5.8 and 6.6).

In the special case $h = 0$, Theorem 1.1, together with a short and self-contained result (Claim 6.3), translates to the following: with high probability, $\mathcal{I} \sim \hat{\mu}_n^0$ has

$$|\mathcal{I} \cap (\llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times \llbracket h_n^* - 1, (1 + \varepsilon_\beta) h_n^* \rrbracket)| > (1 - \varepsilon_\beta) n^2. \quad (1.8)$$

In view of Theorem 1.1, we conjecture the lower bound here is exact (once the interface has reached a height h_n^* above the floor, it has no further incentive to rise), so that the upper bound $(1 + \varepsilon_\beta) h_n^*$ in (1.8) can be replaced by h_n^* . More generally, in Theorem 1.1, we conjecture that for every h , the interface $\mathcal{I} \sim \hat{\mu}_n^h$ has $(1 - \varepsilon_\beta) n^2$ of its faces be at height either $h_n^* - h - 1$ or $h_n^* - h$.

Remark 1.4. While our results are stated for the cylinder $\Lambda_n = \llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \times \mathbb{Z}$, all proofs go through mutatis mutandis on the box $\Lambda_n = \llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^3$ as the interface never feels the effect of the boundary conditions on the top and bottom, except with e^{-cn} probability.

Remark 1.5. Our results extend to any dimension $d \geq 3$ and are presented for $d = 3$ to simplify the exposition. Theorem 1.1 would be given analogously, with n^2 replaced by n^{d-1} , and with h_n^* defined as in (1.3) with 2β replaced by $(d-1)\beta$ and with the natural generalization of $\alpha_h(d)$. Pertaining Remark 1.2, when $d > 3$, one has $M_n / \log n \xrightarrow{P} (d-1)/\alpha(d)$, whence the ratio $\mathbb{E}[M_n]/h_n^*$ goes to $d-1$ as $n \rightarrow \infty$.

1.1. Related work. The study of entropic repulsion for interfaces separating stable phases of the Ising model has a long history. We first describe the picture in \mathbb{Z}^2 . Without any floor (i.e., under μ_n^\mp as opposed to $\hat{\mu}_n^h$), the Ising interface in a strip of side-length n with Dobrushin's boundary conditions is known to have \sqrt{n} height fluctuations [14], and converge to a Brownian bridge under a diffusive rescaling at all $\beta > \beta_c(d)$ [24, 26]. In the presence of a floor at height zero, either by means of plus boundary conditions in the entire lower half-space or by means of conditioning on the interface being restricted to the upper half-space, one could easily deduce that the fluctuations remain $O(\sqrt{n})$; convergence to a Brownian excursion was recently shown in [28]. In particular, no matter the conditioning on a floor (at whatever negative height) the entropic repulsion effect does not change the order of the typical height of the interface. We also mention that in this two-dimensional setting, [29] established an equivalence between the floor effects induced by conditioning on the interface to be non-negative, and by plus boundary conditions in the lower half-space.

In dimensions $d \geq 3$, the phenomenology is completely different. Recall that without any floor conditioning, the interface is rigid about its ground state, i.e., the interface is mostly flat at height zero, with $O(1)$ height oscillations about that with exponential tails [16, 38]. In [21, 22] the authors studied more refined features regarding the typical shape of this interface near its high points, and used that to deduce that the maximum height is tight and has uniform Gumbel tails about its median $m_n^* = (\frac{2}{\alpha} + o(1)) \log n$.

When introducing a hard floor, either via conditioning or via a set of plus spins in the lower half-space, far less is known, though this problem has been discussed in the physics literature at least since [6]. That work proposed using the solid-on-solid (SOS) model [1], a distribution over $\varphi : \llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2 \rightarrow \mathbb{Z}$ with Hamiltonian, $\sum_{x \sim y} |\varphi_x - \varphi_y|$ as an approximation to the Ising interface on which such questions can be studied. The work [6] identified the order of the typical height the SOS surface rises to due to the entropic repulsion effect at sufficiently low temperatures, and in [8, 9] the exact asymptotics of this height were determined. In particular, it was found in [8, 9] that in the SOS model, the typical height of the surface above a floor at height zero, is exactly half of the maximum height in the absence of a floor (note this matches the scaling relation between h_n^* and m_n^* we found for the Ising model: see Remark 1.2). Beyond that, [10] established that the macroscopic level lines converge to the Wulff shape, and have $n^{1/3+o(1)}$ fluctuations along the sides of the box. The work [32] generalized the results of [8, 9] of this entropic repulsion phenomenon to integer valued $|\nabla\varphi|^p$ interfaces (the SOS model being the $p = 1$ case). Let us also mention that the question of entropic repulsion has been extensively studied in the context of the discrete Gaussian free field [4, 5, 12, 13]. In all such height function approximations, identification of the height to which the interface rises given a floor at height zero is effectively equivalent to identification of the floor height $-h$ at which the phase transition in the form of Theorem 1.1 occurs.

Despite this progress for understanding the entropic repulsion phenomenon on the SOS approximation, for the actual Ising interface, much less was known. A non-quantitative delocalization result that the typical height of the interface goes to infinity when the floor is at height zero was established by means of correlation inequalities in [17, 18]. The works [25, 33] sketched an argument that this interface rises to a height that is between $(c/\beta) \log n$ and $(C/\beta) \log n$ for two different constants $C > c > 0$. However, we are not aware of any previous work on the 3D Ising model pinpointing the location of an entropic repulsion phenomenon to within $o(\log n)$, let alone identifying its exact height.

Understanding the entropic repulsion phenomenon also opens the door to studying *wetting phenomena* for interfaces between phases: see [6] as well as the surveys [30, 40] for details on the below discussion. Here, one can for instance change the interaction strength for the edges along the floor from 1 to some $J \geq 0$; the interface then undergoes a transition at some $J_w(\beta) \leq 1$ between a *partial wetting* regime ($J < J_w$) where the typical interface height is $O(1)$, and a *complete wetting* regime ($J > J_w$) in which the typical interface height diverges. In $d = 2$, this transition can be deduced for the Ising model by means of the machinery of [7, 35, 36]. In $d \geq 3$, wetting for the SOS model was first studied in [11], and $J_w(\beta)$ was precisely identified in [31]; for

the Ising model, some partial results can be obtained, as in [2, 3, 18]. A related physical phenomenon occurs when instead keeping $J = 1$, but introducing an external field $-\lambda < 0$ which pushes the interface down, competing with the repulsion away from the floor. This competition induces a physical phenomenon known as *pre-wetting*, whereby as $\lambda = \lambda_n \downarrow 0$, the interface interpolates between the delocalized one at $\lambda = 0$ and one having uniformly bounded height oscillations when λ is uniformly bounded away from zero. In $d = 2$, this interpolation is smooth, and the interface behavior is now well-understood, exhibiting KPZ fluctuation exponents and local convergence to the Ferrari–Spohn diffusion [19, 27, 39]. In the $d \geq 3$ case, less is known rigorously, aside from the first few jumps of the interface for λ small but independent of the system size [3].

1.2. Proof ideas. The proof of Theorem 1.1 proceeds in two parts corresponding to each of (1.5) and (1.6). A more detailed formulation of (1.5) can be found in Theorem 5.1, and a more detailed formulation of (1.6) can be found in Theorem 6.1. Unlike [16, 21–23] the proofs in this paper are primarily probabilistic in nature, with all combinatorial objects we use having already appeared in the previous work [23].

1.2.1. A priori regularity estimates. The first step in our analysis consists of establishing that, typically,

- (i) all but an ε_β -fraction of the faces of \mathcal{I} are horizontal faces for which the intersection of \mathcal{I} with the column through them is a singleton—call these *ceiling faces* following Dobrushin [16]; and
- (ii) all but an ε_β -fraction of the ceiling faces belong to connected components—referred to as *ceilings*—that are reasonably large and regular, being of size at least $n^{1.9}$ and with a boundary of size at most $n^{1+o(1)}$.

In light of these facts, we may restrict our attention to the heights of faces of \mathcal{I} belonging to ceilings as in (ii). This is crucial because the estimates we import from [23] on large deviations of the maximal height oscillations interior to a ceiling are only sharp inside ceilings with a uniformly bounded isoperimetric dimension.

Such a priori estimates are established using the basic fact that, if $\Gamma \subset \mathbf{I}_h^\uparrow$ is such that $\mu_n^\mp(\Gamma) \ll \mu_n^\mp(\mathbf{I}_h^\uparrow)$, then Γ is also atypical under the conditional measure $\hat{\mu}_n^h = \mu_n^\mp(\cdot \mid \mathbf{I}_h^\uparrow)$. Concretely, the lower bound on $\mu_n^\mp(\mathbf{I}_0^\uparrow)$ given in the next proposition allows us to rule out any event Γ which, e.g., has $\mu_n^\mp(\Gamma) \leq \exp(-1.1n \log n)$.

Proposition 1.6. *There exists $\beta_0 > 0$ such that, for every fixed $\beta > \beta_0$ and every sufficiently large n , the probability of the event \mathbf{I}_0^\uparrow , i.e. that the 3D Ising interface \mathcal{I} is a subset of $[-\frac{n}{2}, \frac{n}{2}]^2 \times \mathbb{R}_+$, satisfies*

$$e^{-(1+\varepsilon_\beta)n \log n} \leq \mu_n^\mp(\mathbf{I}_0^\uparrow) \leq e^{-(1-\varepsilon_\beta)n \log n},$$

where ε_β is some sequence going to 0 as β increases.

The lower bound in this proposition is relatively straightforward, and we establish it in §4 en route to deriving the mentioned preliminary regularity estimates. The upper bound is more delicate, and while not needed for our arguments, in §5.3 we obtain it as a corollary of our proof of (1.5) (namely, from the more detailed result in Theorem 5.2, which rules out the existence of large ceilings below height $h_n^* - h - 1$ under $\hat{\mu}_n^\mp$).

1.2.2. Idea of proof of (1.5). We begin with a natural approach for ruling out the existence of a ceiling \mathcal{C} (a connected component of (horizontal) ceiling faces of \mathcal{I}), whose height $\text{ht}(\mathcal{C})$ is $h_n^* - h - 1 - k$ for $k \geq 1$ and whose projection on height 0, together with every finite component it bounds (in these components the interface can exhibit local oscillations relative to \mathcal{C}) is a set S with area $|S| \geq \theta^k n^2$ for some $0 < \theta(\beta) < 1$. (The aim would be to show that such a \mathcal{C} is atypical under $\hat{\mu}_n^h$, and then take a union bound over k .) Denote the maximal downward oscillation inside \mathcal{C} by $\bar{M}_S^\downarrow = \text{ht}(\mathcal{C}) - \min\{x_3 : (x_1, x_2, x_3) \in \mathcal{I}, (x_1, x_2) \in S\}$, and suppose for the moment that the following bound holds: $\mu_n^\mp(\bar{M}_S^\downarrow < h) \lesssim \exp(-|S|e^{-\alpha_h})$ for any h , with α_h from (1.4). (The actual bound available to us—see (2.9)—has a few small differences (error terms, a restriction on h) in addition to one major proviso we will soon describe.) This estimate is quite intuitive in light of the fact ([21, 22]) that the μ_n^\mp -probability of a local oscillation of height at least h below a site in the bulk is approximately $e^{-\alpha_h}$; indeed, if the oscillations in every site of S were mutually independent, we would get the bound $(1 - e^{-\alpha_h})^{|S|} \approx \exp(-|S|e^{-\alpha_h})$. In the event that a ceiling \mathcal{C} as described above exists, we are guaranteed to have $\bar{M}_S^\downarrow < h_n^* - k$ by the implicit conditioning on \mathbf{I}_h^\uparrow in $\hat{\mu}_n^h$. To rule this out, we shift the interface up by k , effectively sending \mathcal{C} to height $h_n^* - h - 1$, which reduces the weight of the interface by about $\exp(-4(\beta - C)kn)$ due to the necessary addition of $4kn$ faces to the interface. The benefit of doing so is that, even though we can still only say that $\bar{M}_S^\downarrow < h_n^* - k$ (the local oscillations within \mathcal{C} are unaffected by the shift), the shifted interface is only guaranteed to have $\bar{M}_S^\downarrow < h_n^*$ by the implicit conditioning in $\hat{\mu}_n^h$, at which point one could utilize the fact that $\mu_n^\mp(\bar{M}_S^\downarrow < h_n^* - k \mid \bar{M}_S^\downarrow < h_n^*) \approx \mu_n^\mp(\bar{M}_S^\downarrow < h_n^* - k) \lesssim \exp(-|S|e^{-\alpha_{h_n^* - k}})$

(where one also uses a sharp *lower* bound on the denominator in the ratio $\mu_n^\mp(\bar{M}_S^\downarrow < h_n^* - k)/\mu_n^\mp(\bar{M}_S^\downarrow < h_n^*)$). In summary, the probability of encountering such a ceiling \mathcal{C} can be bounded from above by approximately

$$\exp(4(\beta - C)kn - |S|e^{-\alpha_{h_n^* - k}}).$$

The sequence (α_h) is known [21] to satisfy $\alpha_{h_n^* - k} \leq \alpha_{h_n^*} - (4\beta - C)k$, whereas $e^{-\alpha_{h_n^*}} \leq e^{2\beta}/n$ by the definition of h_n^* in (1.3). Plugging these in the bound above, together with the fact that $|S| \geq \theta^k n^2$, yields

$$\exp(4(\beta - C)kn - e^{-2\beta}(\theta e^{4\beta - C})^k n),$$

in which the second term in the exponent dominates the first term for $\theta := e^{-\beta}$ when β is large enough.

This approach highlights the competition between energy and entropy (the terms $4\beta kn$ and $|S|e^{-\alpha_{h_n^* - k}}$) propelling the interface from being rigid at height 0. It is worthwhile comparing this argument to the simpler analysis of entropic repulsion in the SOS model: there, to rule out a large subset \mathcal{C} at height less than h , one can lift the configuration by k , and then plant downward spikes in an ε_β -subset of \mathcal{C} . In the 3D Ising model, the large deviation rate of the height above/below a site in the bulk is not governed by a deterministic shape (such as the single column spike in the SOS model), but rather by a distribution over complex oscillations; hence, we instead consider interfaces with $\bar{M}_S^\downarrow < h_n^* - k$, where no oscillation dipped below a certain height.

One obvious difficulty in carrying out the above approach is due to having, in lieu of the simplified case of one given ceiling \mathcal{C} , an abundance of randomly located ceilings. One needs to reveal (here and in what follows, we use the term *reveal* or *expose* to refer to conditioning on the realizations of the objects under consideration) their boundaries, but not their interior, in order to appeal to the bounds of [23], and do so with care, as the boundary of one ceiling may overlap with that of another. Moreover, a single *wall*—a connected component of non-ceiling faces—may give rise to multiple ceilings at different heights. However, these details can be handled by exposing the outermost walls and reevaluating the ceiling landscape given the exposed walls (and the ceilings they nest): we may identify a slab at height $h_n^* - h - 1 - k$ where the total area in the ceilings \mathcal{C}_i exceeds $\theta^k n^2$, and attempt to apply the above argument to each of the S_i 's separately, for a bound of $\exp(-\sum |S_i|e^{-\alpha_{h_n^* - k}})$.

The main obstacle, as hinted above, is actually the upper bound we have for $\mu_n^\mp(\bar{M}_S^\downarrow < h)$, which is only valid in conjunction with an indicator that every wall in the bulk of S must be relatively small—namely, the diameter of every such wall should be at most $\exp(ch) \lesssim n^{\varepsilon_\beta}$ (see Proposition 2.28 and the event $\mathcal{G}_{S^\circ}^{\mathfrak{D}_\circ}$ in the key estimate (2.9) in it). The source of this constraint is that, in order to prove said bound in [23], the region S was partitioned into boxes B_i of side length L , whose interiors B_i° were inspected carefully (while ignoring $B_i \setminus B_i^\circ$); the constraint on the maximum diameter of any wall meant that no wall in B_i° could reach ∂B_i , leading to the desired independence of local oscillations between B_i° 's. We stress that this obstacle is real, rather than a limitation of the proof technique: a single wall of size $(c/\beta)n \log n$ can raise/lower an area of order n^2 by height $(c/\beta) \log n$, forming correlations between the heights in its interior as it nests sites in ceilings. The cost of such a wall is $\exp(-c'n \log n)$, which can still be larger than $\mu_n^\mp(\mathbb{I}_h^{\text{fl}})$, and indeed, this is precisely the mechanism of entropic repulsion, where most of the interface is raised by a single large wall...

To overcome this obstacle, we *reveal every mesoscopic wall*—in the sense of having diameter larger than a threshold of $e^{c(h_n^* - k)} \vee \log n$; see (5.1)—in \mathcal{I} and carry the above approach in this conditional space. This is achieved in Lemma 5.6, the key to showing (1.5). The following are worthwhile mentioning from its proof:

- (1) Conditioning on all mesoscopic walls can reveal a complicated set of walls nesting one another, supporting a collection of ceilings \mathcal{C}_i (the full interface \mathcal{I} will modify these \mathcal{C}_i 's via nested microscopic walls). Whereas we would like to appeal to Proposition 2.28 simultaneously for every \mathcal{C}_i which is at height $h_n^* - h - 1 - k$, there are two prerequisites on its interior S_i to qualify an application of that proposition: (i) S_i should be simply-connected; (ii) the isoperimetric dimension of S_i should be bounded by an absolute constant. Item (ii) is handled by our a priori regularity estimates, as we only treat S_i 's with area at least $n^{1.9}$ and boundary at most $n^{1+o(1)}$. However, Item (i) is simply false when one mesoscopic wall nests another... Our remedy for this is to further expose the walls along a minimal collection of faces that connect the boundaries of all S_i 's, which will make the unrevealed portion of the S_i 's simply-connected by definition.
- (2) Controlling the set of faces that connect the S_i 's is imperative: too large of a set may potentially deform the regularity of the S_i 's, increasing their isoperimetric dimension beyond the scope of Proposition 2.28. Here we rely on the fact that there is a total of $n^{1+o(1)}$ faces in the union of the boundaries of the S_i 's. A classical bound on the *Euclidean Traveling Salesman Problem* (TSP) then shows that we can connect an arbitrary face from each one via a minimum set of size at most $n^{3/2+o(1)}$ faces. Thus, revealing the

(microscopic, by the conditioning) walls along said set modifies the area and boundary of each S_i by at most $n^{3/2+o(1)}$, keeping its isoperimetric dimension (we had $|S_i| \geq n^{1.9}$ and $|\partial S_i| \leq n^{1+o(1)}$) in check.

- (3) Recall that in the simplified outline one needed a sharp upper bound for $\mu_n^\mp(\bar{M}_S^\downarrow < h_n^* - k \mid \bar{M}_S^\downarrow < h_n^*)$. In the actual proof, we must instead estimate $\mu_n^\mp(\bar{M}_S^\downarrow < h_n^* - k \mid \bar{M}_S^\downarrow < h_n^*, \mathcal{G}_{S^\circ}^{\mathcal{D}}$), where $\mathcal{G}_{S^\circ}^{\mathcal{D}}$ is the event that no walls in S° , the bulk of S , are mesoscopic. We infer this bound from a sharp upper bound on $\mu_n^\mp(\bar{M}_S^\downarrow < h_n^* - k, \mathcal{G}_{S^\circ}^{\mathcal{D}})$ and a sharp lower bound on $\mu_n^\mp(\bar{M}_S^\downarrow < h_n^*, \mathcal{G}_{S^\circ}^{\mathcal{D}})$. The former is provided by (2.9) in Proposition 2.28, as discussed above. The latter is available from (2.8), which is given in terms of a slightly different event $\mathcal{G}_{S^\circ}^{\mathcal{m}}$, yet in our application this event implies the required event $\mathcal{G}_{S^\circ}^{\mathcal{D}}$.

In Section 5 we give a short proof (Claim 5.4) for entropic repulsion at $\mathbf{h} < (1 - \varepsilon_\beta)h_n^*$ (away from the critical threshold); see the discussion following it for more on the obstacles that arise at $\mathbf{h} = (1 + o(1))h_n^*$.

1.2.3. *Idea of proof of (1.6).* We now describe our approach for showing that in the regime of $\mathbf{h} \geq h_n^*$, all but $\varepsilon_\beta n^2$ faces of \mathcal{I} are *at most* at height zero (the fact that all but $\varepsilon_\beta n^2$ faces of \mathcal{I} are *at least* at height zero follows straightforwardly from the FKG inequality and rigidity under the unconditional measure μ_n^\mp).

Consider a ceiling \mathcal{C} at height $\text{ht}(\mathcal{C}) \geq 1$ whose interior S has area at least θn^2 . To demonstrate the basic idea that lets rigidity prevail over entropic repulsion in this range of \mathbf{h} , suppose that the interface gives rise to this ceiling by means of a cylindrical wall W of the form $\partial S \times [0, \text{ht}(\mathcal{C})]$. In this situation, one can compare the original interface \mathcal{I} with the interface \mathcal{I}' obtained by truncating the height of the wall W to be $\text{ht}(\mathcal{C}) - 1$: the weight of \mathcal{I}' will increase by about $\exp((\beta - C)|\partial S|)$, so this simple Peierls argument can rule out the existence of such a ceiling \mathcal{C} as long as it does not violate the hard floor constraint. By definition, $\mathcal{I}' \in \mathbf{I}_\mathbf{h}^{\text{fl}}$ and only if \bar{M}_S^\downarrow , the maximal downward oscillation inside \mathcal{C} , satisfies $\bar{M}_S^\downarrow < \text{ht}(\mathcal{C}) + \mathbf{h}$. The probability of the latter event is exponentially small in n , yet it can still be outweighed by the energy gain due to truncating W .

Concretely, let $\mathbf{I}_W \subset \mathbf{I}_\mathbf{h}^{\text{fl}}$ be the set interfaces containing the above wall W , and further let $\tilde{\mathbf{I}}_W \subset \mathbf{I}_W$ be the subset of interfaces which in addition satisfy $\bar{M}_S^\downarrow \leq h_n^*$ (so that, in particular, $\bar{M}_S^\downarrow < \text{ht}(\mathcal{C}) + \mathbf{h}$). We may reduce $\hat{\mu}_n^\mathbf{h}(\mathcal{I} \in \tilde{\mathbf{I}}_W \mid \mathcal{I} \in \mathbf{I}_W)$, akin to the discussion in the previous section, to $\mu_n^\mp(\bar{M}_S^\downarrow \leq h_n^* \mid \bar{M}_S^\downarrow \leq \text{ht}(\mathcal{C}) + \mathbf{h})$, and appeal to Proposition 2.28—however this time we are interested in a sharp lower bound on this quantity. To that end, we can use the sharp lower bound on $\mu_n^\mp(\bar{M}_S^\downarrow \leq h_n^*) \geq \exp(-(1 + \varepsilon_\beta)|S|e^{\alpha h_n^* + 1})$ from (2.8) in that proposition (only the upper bound (2.9) had the extra event $\mathcal{G}_{S^\circ}^{\mathcal{D}}$ that greatly complicated matters), and simply drop the conditioning (only decreasing the probability in doing so). By the definition of h_n^* in (1.3) and the fact that the sequence is known ([22]) to have $\alpha_{h+1} \geq \alpha_h + 4\beta - C$, we have $\alpha_{h_n^*+1} \geq e^{-2\beta+C}/n$, and combining this with the Peierls bound on $\hat{\mu}_n^\mathbf{h}(\tilde{\mathbf{I}}_W)$ shows that $\hat{\mu}_n^\mathbf{h}(\mathbf{I}_W)$ is bounded from above by approximately

$$\exp(-(\beta - C)|\partial S| + |S|e^{-2\beta+C}/n) \leq \exp(-(\beta - C + e^{-2\beta+C})|\partial S|),$$

using here the isoperimetric inequality $|S| \leq |\partial S|n/4$. Notice that this argument did not actually need S to have area θn^2 ; e.g., if the boundary of S were at least $\log n$, one could already rule out \mathcal{C} via a union bound; our actual proof (applicable to a general wall W as opposed to the cylinder from the toy example above), along the same vein, will rule out any ceiling whose supporting wall W contains at least $n^{9/10}$ faces.

The main obstacle in using this approach to rule out general ceilings \mathcal{C} at positive height is that, in the 3D Ising model, the only (tractable) approach one has to shrinking the height of a subset of the interface is the deletion of a collection of walls: even a single wall may have a complicated landscape of overhangs and nested walls near its boundary, so devising a notion of truncating its height to arrive at a valid interface, while gaining the energy of $|\partial S|$ without an entropic cost, is highly nontrivial. (E.g., how do we reconcile the effect of having walls nested in W shift due to the truncation and collide with W ? Deleting those would impact the entropy, whereas modifying them in any other way can, in addition, lead to similar ripple effects.)

Furthermore, even if we decide to delete the entire wall W , our Peierls map cannot consist of that alone: due to the long-range interactions that the deletion of W and the resulting shifts induce between interface faces, this operation must also be followed by the deletion of certain other walls that interact too strongly with W —either the *groups of walls* in Dobrushin's original work [16], or the *one-sided wall clusters* in [23]. Roughly put, the wall cluster criterion says that if W nests a wall W' whose number of faces is larger than its distance to W , then we must also tag W' for deletion, and process its inner walls recursively (see Def. 2.20). Hence, if we delete W and its wall cluster \mathbf{W} , and then aim to analyze \bar{M}_S^\downarrow , we have the following problem: whereas revealing only W and its exterior gave no information on the interface in S (so we were free to apply

Proposition 2.28), the act of revealing its wall cluster \mathbf{W} robs us of that feature, by imposing the constraint that the size of every wall W' that remained unexposed in S should not exceed its distance to \mathbf{W} ...

Our remedy for this problem is establishing a certain monotonicity principle, which may find other uses. Intuitively, conditioning on wall sizes being below various thresholds, albeit complicated (here, the threshold for a given wall depends on its distance to the wall cluster), should only yield better control over oscillations. Making this intuition precise is the observation that results relying on Peierls maps that only *deleted walls*—such were, e.g., key results in [23] that we use here—can also be established in the above conditional setting (if an interface satisfied the wall size constraints, so will the interface corresponding to a subset of its walls). As such, we can extend said results from [23], and notably the lower bounds of Proposition 2.28, to this setting (see Proposition 3.3). This observation and the extensions it gives rise to appear in Section 3.

In Section 6 we give a short proof (Claim 6.3) for rigidity at height zero for $h \geq (1 + \varepsilon_\beta)h_n^*$ (away from the critical threshold); see the discussion following it for more on the obstacles that arise at $h = (1 + o(1))h_n^*$.

1.3. Outline of paper. In Section 2, we overview the notation we use in the paper and recall many of the key definitions and tools from [16] and [23] used in our analysis. In Section 3, we extend several of the preliminary estimates of [23] to also hold under a monotone conditioning event, which will later be needed for showing (1.6). In Section 4, we establish an a priori bound on the probability that the interface lies above the floor at height 0, and use it to prove preliminary regularity estimates for the interface above a floor at $-h$. In Section 5, we treat the case of $h < h_n^* - 1$, establishing (1.5) of Theorem 1.1. Then in Section 6, we handle $h \geq h_n^*$, establishing (1.6) of Theorem 1.1.

2. PRELIMINARIES

In this section, we recall the decomposition of the interface into walls and ceilings and their groupings into wall clusters. We then recall the results of [23] that we will require in this paper regarding the law of the maximum oscillation inside a ceiling, conditionally on all walls outside that ceiling. Throughout this, we will also introduce any notation that will be used throughout the paper.

2.1. Notation. We begin by describing the underlying geometry on which we will be working, and any related graph notation we will use throughout the paper.

Lattice notation. The underlying graphs we consider throughout this paper are rectangular subsets of \mathbb{Z}^3 . To be precise, \mathbb{Z}^3 is the integer lattice graph with vertices at $(x_1, x_2, x_3) \in \mathbb{Z}^3$ and edges between nearest neighbor vertices (at Euclidean distance one). A *face* of \mathbb{Z}^3 is the open set of points bounded by four edges (or four vertices) forming a square of side length one, lying normal to one of the coordinate directions. A *cell* of \mathbb{Z}^3 is the set of points bounded by six faces (or eight vertices) forming a cube of side length one.

Because our interest is the behavior of the *interface* separating plus and minus spins, it will be convenient for us to consider Ising configurations as assignments of ± 1 spins to the vertices of the *dual* graph $(\mathbb{Z}^3)^* = (\mathbb{Z} + \frac{1}{2})^3$; these are naturally identified with the cells of \mathbb{Z}^3 for which they are the midpoint.

More generally, we will frequently identify edges, faces, and cells with their midpoints. A subset $\Lambda \subset \mathbb{Z}^3$ specifies an edge, face, and cell collection via the edges, faces, and cells of \mathbb{Z}^3 all of whose bounding vertices are in Λ . We will denote the resulting edge set by $\mathcal{E}(\Lambda)$, face set by $\mathcal{F}(\Lambda)$, and cell set by $\mathcal{C}(\Lambda)$.

Two edges are adjacent if they share a vertex, two faces adjacent if they share a bounding edge, and two cells adjacent if they share a bounding face. We denote adjacency by the notation \sim . It will also be useful to have a notion of connectivity in \mathbb{R}^3 (as opposed to \mathbb{Z}^3); we say that an edge/face/cell is **-adjacent*, denoted \sim^* , to another edge/face/cell if they share a bounding vertex. A connected component of faces (respectively **-connected* component of faces) is a maximal set of faces such that for any pair of faces in that set, there is a sequence of adjacent faces (resp., **-adjacent*) faces starting at one face and ending at the other.

We use the notation $d(A, B) = \inf_{x \in A, y \in B} d(x, y)$ to denote the Euclidean distance in \mathbb{R}^3 between two sets A, B . We then let $B_r(x) = \{y : d(y, x) \leq r\}$. When these balls are viewed as subsets of edges/faces/cells, we include all those edges/faces/cells whose midpoint falls in $B_r(x)$.

Subsets of \mathbb{Z}^3 . As mentioned, the primary subsets of \mathbb{Z}^3 with which we will be concerned are of the form of cubes and cylinders. In view of that, define the centered $n \times m \times h$ box,

$$\Lambda_{n,m,h} := \left[[-\frac{n}{2}, \frac{n}{2}] \times \left[-\frac{m}{2}, \frac{m}{2}\right] \times \left[-\frac{h}{2}, \frac{h}{2}\right] \subset \mathbb{Z}^3,$$

where

$$[[a, b]] := \{[a], [a] + 1, \dots, [b] - 1, [b]\}.$$

We can then let Λ_n denote the special case of the cylinder $\Lambda_{n,n,\infty}$. The (outer) boundary $\partial\Lambda_n$ of the cell set $\mathcal{C}(\Lambda_n)$ is the set of cells in $\mathcal{C}(\mathbb{Z}^3) \setminus \mathcal{C}(\Lambda_n)$ adjacent to a cell in $\mathcal{C}(\Lambda_n)$.

We are also going to use a dedicated notation for horizontal slices and half-spaces of \mathbb{Z}^3 . For any $h \in \mathbb{Z}$ let \mathcal{L}_h be the subgraph of \mathbb{Z}^3 having vertex set $\mathbb{Z}^2 \times \{h\}$ and correspondingly define edge and face sets $\mathcal{E}(\mathcal{L}_h)$ and $\mathcal{F}(\mathcal{L}_h)$. For a half-integer $h \in \mathbb{Z} + \frac{1}{2}$, let \mathcal{L}_h collect the faces and cells in $\mathcal{F}(\mathbb{Z}^3) \cup \mathcal{C}(\mathbb{Z}^3)$ whose midpoints have half-integer e_3 coordinate h . A certain such set which will recur is \mathcal{L}_0 and its restriction to $\mathcal{F}(\Lambda_n)$ which we denote by $\mathcal{L}_{0,n}$.

Finally we use $\mathcal{L}_{>h} = \bigcup_{h' > h} \mathcal{L}_{h'}$ and $\mathcal{L}_{<h} = \bigcup_{h' < h} \mathcal{L}_{h'}$, and similarly $\mathcal{L}_{\geq h}$ and $\mathcal{L}_{\leq h}$, for half-spaces.

Projections onto \mathcal{L}_0 . Throughout the paper, we will refer to a face as *horizontal* if its normal vector is $\pm e_3$, and as *vertical* if its normal vector is one of $\pm e_1$ or $\pm e_2$.

For a face $f \in \mathcal{F}(\mathbb{Z}^3)$, its *projection* is the edge or face given by

$$\rho(f) = \{(x_1, x_2, 0) : (x_1, x_2, s) \in f \text{ for some } s \in \mathbb{R}\} \subset \mathcal{L}_0.$$

Specifically, the projection of a *horizontal* face is a face in $\mathcal{F}(\mathcal{L}_0)$, while the projection of a *vertical* face is an edge in $\mathcal{E}(\mathcal{L}_0)$. The projection of a collection of faces F is given by $\rho(F) := \bigcup_{f \in F} \rho(f)$, which may consist both of edges and faces of \mathcal{L}_0 .

2.2. The Ising model. An Ising configuration σ on a subset $\Lambda \subset \mathbb{Z}^3$ is an assignment of ± 1 -valued spins to the cells of Λ , i.e., $\sigma \in \{\pm 1\}^{\mathcal{C}(\Lambda)}$. For a finite connected subset $\Lambda \subset \mathbb{Z}^3$, the Ising model on Λ with boundary conditions $\eta \in \{\pm 1\}^{\mathcal{C}(\mathbb{Z}^3)}$ is the probability distribution over $\sigma \in \{\pm 1\}^{\mathcal{C}(\Lambda)}$ given by

$$\begin{aligned} \mu_\Lambda^\eta(\sigma) &\propto \exp[-\beta \mathcal{H}(\sigma)], \quad \text{where} \\ \mathcal{H}(\sigma) &= \sum_{\substack{v, w \in \mathcal{C}(\Lambda) \\ v \sim w}} \mathbf{1}\{\sigma_v \neq \sigma_w\} + \sum_{\substack{v \in \mathcal{C}(\Lambda), w \in \mathcal{C}(\mathbb{Z}^3) \setminus \mathcal{C}(\Lambda) \\ v \sim w}} \mathbf{1}\{\sigma_v \neq \eta_w\}. \end{aligned} \quad (2.1)$$

Throughout this paper, we will be considering the *Dobrushin boundary conditions*, denoted $\eta = \mp$, where $\eta_w = -1$ if w is in the upper half-space ($w_3 > 0$) and $\eta_w = +1$ if w is in the lower half-space ($w_3 < 0$).

Infinite-volume measures. Though (2.1) is defined only on finite graphs, the definition can be extended to infinite graphs via a consistency criterion known as the *DLR conditions* (see e.g., the book [20] for a definition and discussion). On \mathbb{Z}^d , such *infinite-volume measures* arise as weak limits of finite-volume measures, say $n \rightarrow \infty$ limits of the Ising model on boxes of side length n with prescribed sequences of boundary conditions. At low temperatures $\beta > \beta_c(d)$, the Ising model on \mathbb{Z}^d admits multiple infinite-volume Gibbs measures $\mu_{\mathbb{Z}^3}^+$ and $\mu_{\mathbb{Z}^3}^-$ obtained by taking plus and minus boundary conditions on boxes of side-length n and sending $n \rightarrow \infty$. An important consequence of the work of [16] was that when $d \geq 3$, the weak limit $\mu_{\mathbb{Z}^3}^\mp := \lim_{n \rightarrow \infty} \mu_{n,n,n}^\mp$ is a non-translation-invariant DLR measure, and is thus distinct from any mixtures of $\mu_{\mathbb{Z}^3 2}^+$ and $\mu_{\mathbb{Z}^3 2}^-$.

2.3. Interfaces under Dobrushin boundary conditions. Having defined Ising configurations and the Ising measure with Dobrushin boundary conditions, let us now formally define the *interface* separating the plus and minus phases.

Definition 2.1 (Interfaces). For a domain $\Lambda_{n,m,h}$ with Dobrushin boundary conditions, and an Ising configuration σ on $\mathcal{C}(\Lambda_{n,m,h})$, the *interface* $\mathcal{I} = \mathcal{I}(\sigma)$ is defined as follows:

- (1) Extend σ to a configuration on $\mathcal{C}(\mathbb{Z}^3)$ by taking $\sigma_v = +1$ if $v \in \mathcal{L}_{<0} \setminus \mathcal{C}(\Lambda_{n,m,h})$ and $\sigma_v = -1$ if $v \in \mathcal{L}_{>0} \setminus \mathcal{C}(\Lambda_{n,m,h})$.
- (2) Let $F(\sigma)$ be the set of faces in $\mathcal{F}(\mathbb{Z}^3)$ separating cells with differing spins under σ .
- (3) Let $\mathcal{I}^\infty(\sigma)$ be the (maximal) $*$ -connected component of faces in $F(\sigma)$ containing $\mathcal{L}_0 \setminus \mathcal{F}(\Lambda_{n,m,h})$ in $F(\sigma)$. (This is also the unique infinite $*$ -connected component in $F(\sigma)$.)
- (4) The interface $\mathcal{I}(\sigma)$ is the restriction of $\mathcal{I}^\infty(\sigma)$ to $\mathcal{F}(\Lambda_{n,m,h})$.

Taking the $h \rightarrow \infty$ limit $\mu_{n,m,h}^\mp$ to obtain the infinite-volume measure $\mu_{n,m,\infty}^\mp$, the interface defined is easily seen to stay finite almost surely. Thus, $\mu_{n,m,\infty}^\mp$ -almost surely, the above process also defines the interface for configurations on all of $\mathcal{C}(\Lambda_{n,m,\infty})$.

Remark 2.2. Every (finite) interface uniquely defines a configuration with exactly one $*$ -connected plus component and exactly one $*$ -connected minus component. For every \mathcal{I} , we can obtain this configuration, which we call $\sigma(\mathcal{I})$, by iteratively assigning spins to $\mathcal{C}(\Lambda_{n,m,h})$, starting from $\partial\Lambda_{n,m,h}$ and proceeding inwards, in such a way that adjacent sites have differing spins if and only if they are separated by a face in \mathcal{I} . Informally, $\sigma(\mathcal{I})$ is indicating the sites that are in the “plus phase” and “minus phase” given the interface \mathcal{I} .

2.4. The Ising interface conditioned on a floor. In this paper, we study the behavior of the interface $\mathcal{I}(\sigma)$ when conditioned on a *floor event*, i.e., conditioned on the interface lying above a certain horizontal plane (the floor). For us, the floor will either reside at height zero, or into the lower half-space, so that we are examining the entropic repulsion of the interface into the upper half-space. More precisely, we use the following notation to denote the Ising measure with a floor at height $-h$ for $h \geq 0$:

$$\hat{\mu}_n^h(\sigma) = \mu_n^\mp(\sigma \mid \mathcal{I} \subset \mathcal{L}_{\geq -h}) = \mu_n^\mp(\sigma \mid \mathbf{I}_h^{\text{fl}}).$$

2.5. Decomposition of the interface: walls and ceilings. Having defined the key objects of interest in this paper, we now begin to collect the main tools for its analysis. We begin by recalling the classical decomposition of [16] of the Ising interface into *walls* and *ceilings*.

Definition 2.3 (Walls and ceilings). A face $f \in \mathcal{I}$ is a *ceiling face* if it is horizontal and there is no $f' \in \mathcal{I}$, $f' \neq f$ such that $\rho(f) = \rho(f')$. A face $f \in \mathcal{I}$ is a *wall face* if it is not a ceiling face.

A *wall* is a $*$ -connected component of wall faces and a *ceiling* is a $*$ -connected component of ceiling faces.

Intuitively speaking, the walls capture all the oscillations of the interface off of the lowest energy (horizontally flat) interface, while the ceilings capture the flat stretches of the interface. In particular, every ceiling has a unique *height*, denoted $\text{ht}(\mathcal{C})$, since all faces in the ceiling have the same x_3 coordinate.

Definition 2.4. Throughout the paper, we let $\mathfrak{C}(\mathcal{I})$ be the collection of all ceilings of \mathcal{I} , and for every height $h \in \mathbb{Z}$, we let $\mathfrak{C}_h(\mathcal{I}) = \{\mathcal{C} \in \mathfrak{C}(\mathcal{I}) : \text{ht}(\mathcal{C}) = h\}$.

The projections of walls, like non-crossing loop collections in \mathbb{Z}^2 , satisfy certain important nesting relations.

Definition 2.5 (Nesting of walls). For a wall W , the complement (in \mathcal{L}_0) of its projection, denoted

$$\rho(W)^c := (\mathcal{E}(\mathcal{L}_0) \cup \mathcal{F}(\mathcal{L}_0)) \setminus \rho(W),$$

splits into one infinite component, and some finite ones. An edge or face $u \in \mathcal{E}(\mathcal{L}_0) \cup \mathcal{F}(\mathcal{L}_0)$ is said to be *interior* to (or *nested in*) a wall W , denoted by $u \Subset W$, if u is not in the infinite component of $\rho(W)^c$. A wall W' is nested in a wall W , denoted $W' \Subset W$, if every element of $\rho(W')$ is interior to W . Similarly, a ceiling \mathcal{C} is nested in a wall W if every element of $\rho(\mathcal{C})$ is interior to W .

We can then identify the connected components of $\rho(W)^c$ with the ceilings incident to W .

Lemma 2.6 ([16]). *For a projection of the walls of an interface, each connected component of that projection (as a subset of edges and faces) corresponds to a single wall. Moreover, there is a 1-1 correspondence between the ceilings adjacent to a standard wall W and the connected components of $\rho(W)^c$. Similarly, for a wall W , all other walls $W' \neq W$ can be identified to the connected component of $\rho(W)^c$ they project into, and in that manner they can be identified to the ceiling of W to which they are interior.*

The above correspondence can be made more transparent by introducing the following notion.

Definition 2.7. For a wall W , the ceilings incident to W can be decomposed into *interior ceilings* of W (those ceilings identified with the finite connected components of $\rho(W)^c$), and a single exterior ceiling, called the *supporting ceiling* of W , identified with the infinite connected component of $\rho(W)^c$.

Definition 2.8. The hull of a ceiling \mathcal{C} , denoted $\dot{\mathcal{C}}$ is the minimal simply-connected set of horizontal faces containing \mathcal{C} . The hull of a wall, \dot{W} is the union of W with the hulls of its interior ceilings.

Observation 2.9. *For every interior ceiling \mathcal{C} of W , the projection of its hull, $\rho(\dot{\mathcal{C}})$, is exactly the finite component of $\rho(W)^c$ it projects into. On the other hand, the projection of the hull of the floor of W is all of $\mathcal{L}_{0,n}$. Finally, the set $\rho(\dot{W})$ is the union of $\rho(W)$ with all the finite components of $\rho(W)^c$.*

Finally, we can assign the index points of $\mathcal{L}_{0,n}$ the walls of an interface \mathcal{I} as follows.

Remark 2.10. Given an interface \mathcal{I} , for every face $x \in \mathcal{F}(\mathcal{L}_0)$, assign x the wall W of \mathcal{I} if $x \in W$ and x shares an edge with $\rho(W)$. If there is no W for which this is the case, let $W_x = \emptyset$. Importantly, this labeling scheme is such that x is only assigned one wall, but the same wall may be assigned to many index faces.

We also introduce a notion of a *nested sequence of walls* which will recur.

Definition 2.11. To any edge/face/cell x , we can assign a *nested sequence of walls* $\mathfrak{W}_x = \bigcup_s W_{u_s}$ consisting of all walls nesting $\rho(x)$ (by Definition 2.5, this forms a nested sequence of walls).

2.6. The standard wall representation. A key property of the wall and ceilings decomposition of [16] is that neither the ceilings, nor crucially the vertical positions of the walls, are needed to reconstruct the interface. This is formalized by the following notion of standard walls, also defined in [15].

Definition 2.12 (Standard walls). A wall W is a *standard wall* if there exists an interface \mathcal{I}_W such that \mathcal{I}_W has exactly one wall given by W . A collection of standard walls is *admissible* if any two standard walls in the collection have pairwise vertex disjoint projections.

Definition 2.13 (Standardization of walls). For every wall W of \mathcal{I} , we can define its *standardization* $\Theta_{\text{ST}}W$ which is the translate of the wall by $(0, 0, -s)$ where s is the height of its supporting ceiling. (We leave the dependence of s , and therefore the dependence of $\Theta_{\text{ST}}W$ on the rest of \mathcal{I} to be contextually understood.)

We then have the following important bijection between interfaces and their standard wall representation, defined as the collection of standard walls given by standardizing all walls of \mathcal{I} .

Lemma 2.14 ([16]). *The map sending an interface to its standard wall representation is a bijection between the set of all valid interfaces and the set of all admissible collections of standard walls.*

We note the following important observation based on the bijection given by Lemma 2.14.

Observation 2.15. *Consider interfaces \mathcal{I} and \mathcal{J} , such that the standard wall representation of \mathcal{I} contains that of \mathcal{J} (and additionally has the standard walls $\Theta_{\text{ST}}\mathbf{W} = (\Theta_{\text{ST}}W_1, \dots, \Theta_{\text{ST}}W_r)$). There is a 1-1 map between the faces of $\mathcal{I} \setminus \mathbf{W}$ and the faces of $\mathcal{J} \setminus \mathbf{H}$ where \mathbf{H} is the set of faces in \mathcal{J} projecting into $\rho(\mathbf{W})$. Moreover, this bijection can be encoded into a map $f \mapsto \theta_{\downarrow}f$ that only consists of vertical shifts, and such that all faces projecting into the same connected component of $\rho(\mathbf{W})^c$ undergo the same vertical shift.*

2.7. The induced distribution on interfaces. Aside from the decomposition of the Ising interface into walls and ceilings, the other key tool used by [16] to establish rigidity of the interface was an expression for the Ising distribution over interfaces as a perturbation, by means of the cluster expansion [34], of a Gibbs measure on interfaces with weight $\exp(-\beta|\mathcal{I}|)$. The perturbative term takes into account the interaction of bubbles of the low-temperature Ising configurations in the minus phase above and plus phase below the interface, with the interface itself.

Here and throughout the paper, let $\mu_n^{\mp} = \mu_{\Lambda_n}^{\mp} = \lim_{h \rightarrow \infty} \mu_{n,n,h}^{\mp}$.

Theorem 2.16 ([16, Lemma 1]). *There exist $\beta_0 > 0$, $\bar{c} > 0$, $\bar{K} > 0$, and a function $\mathbf{g} = \mathbf{g}_{\beta}$ such that the following holds for every $\beta > \beta_0$. For every interface \mathcal{I} ,*

$$\mu_n^{\mp}(\mathcal{I}) \propto \exp\left(-\beta|\mathcal{I}| + \sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I})\right),$$

and for every $\mathcal{I}, \mathcal{I}'$ and $f \in \mathcal{I}$ and $f' \in \mathcal{I}'$,

$$|\mathbf{g}(f, \mathcal{I})| \leq \bar{K} \tag{2.2}$$

$$|\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(f', \mathcal{I}')| \leq \bar{K} e^{-\bar{c}\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}')} \tag{2.3}$$

where $\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}')$ is the largest radius around the origin on which $\mathcal{I} - f$ (\mathcal{I} shifted by the midpoint of the face f) is congruent to $\mathcal{I}' - f'$. That is to say,

$$\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}') := \sup\{r : (\mathcal{I} - f) \cap B_r(0) \equiv (\mathcal{I}' - f') \cap B_r(0)\},$$

where the congruence relation \equiv is equality as subsets of \mathbb{R}^3 .

We will say that the radius $\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}')$ is *attained* by a face $g \in \mathcal{I}$ (resp., $g' \in \mathcal{I}'$) of minimal distance to f (resp., f') if the presence of that face prevents $\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}')$ from being any larger.

2.8. Excess energy of interfaces and walls. Given the representation of Theorem 2.16, we see that the ratio between the weight of \mathcal{I} and the weight of \mathcal{I}' will have a term of the form $\exp(-\beta(|\mathcal{I}| - |\mathcal{I}'|))$. This difference in sizes of the interfaces can be thought of as the *excess energy* of one interface over the other, and will recur throughout the paper.

Definition 2.17 (Excess energy). For two interfaces \mathcal{I}, \mathcal{J} , the *excess energy* of \mathcal{I} with respect to \mathcal{J} , denoted $\mathbf{m}(\mathcal{I}; \mathcal{J})$, is given by

$$\mathbf{m}(\mathcal{I}; \mathcal{J}) := |\mathcal{I}| - |\mathcal{J}|.$$

Evidently, for any valid interface \mathcal{I} , we have that $\mathbf{m}(\mathcal{I}; \mathcal{L}_{0,n}) \geq 0$.

We relate this notion back to the walls and ceilings, and assign excess energies to walls as follows.

Definition 2.18. For a wall W , recall that $\mathcal{I}_{\Theta_{\text{ST}}W}$ is the interface whose only wall is the standardization of W , and define its excess energy by

$$\mathbf{m}(W) = \mathbf{m}(\mathcal{I}_{\Theta_{\text{ST}}W}; \mathcal{L}_{0,n}).$$

This excess energy can alternatively be expressed as $\mathbf{m}(W) = |W| - |\mathcal{F}(\rho(W))|$; one then finds that

$$\mathbf{m}(W) \geq \frac{1}{2}|W|, \quad \text{and} \quad \mathbf{m}(W) \geq |\rho(W)| = |\mathcal{E}(\rho(W))| + |\mathcal{F}(\rho(W))|. \quad (2.4)$$

Evidently, any two faces $x, y \in \mathcal{L}_{0,n}$ both nested in W must satisfy $d(x, y) \leq \mathbf{m}(W)$.

For a collection of walls \mathbf{W} , we define $\mathbf{m}(\mathbf{W}) = \sum_{W \in \mathbf{W}} \mathbf{m}(W)$. It then becomes evident that the excess energy of an interface, $\mathbf{m}(\mathcal{I}; \mathcal{L}_{0,n})$ is exactly given by the excess energy of its wall collection.

2.9. Wall clusters and rigidity inside ceilings. Up to this point, all of the above consisted of simple reformulations of concepts from the classical work [16]. The main result there was exponential tails on the wall through a face $x \in \mathcal{L}_{0,n}$ for any x . In [23], these ideas were extended and a concept called *wall clusters* was introduced to establish that these exponential tails hold even conditionally on arbitrary nesting walls.

Towards the definition of wall clusters, we first define a notion of *close nesting* between walls.

Definition 2.19. A wall W' is closely nested in a wall W if W' is nested in W ($W' \Subset W$) and

$$d_\rho(W, W') := d(\rho(W), \rho(W')) \leq \mathbf{m}(W').$$

Definition 2.20. For a wall W , define the *wall cluster* $\text{Clust}(W)$ as follows:

- (1) Initialize $\text{Clust}(W)$ with W .
- (2) Iteratively, add to $\text{Clust}(W)$ all walls W'' that are closely nested in some wall $W' \in \text{Clust}(W)$.

Remark 2.21. The key upshot of wall clusters is that unlike the analogous grouping scheme in [16] (called *groups of walls*), the wall cluster of W is completely nested in W , i.e., $\rho(\text{Clust}(W)) \subset \rho(\dot{W})$.

In [23] we used the following bounds on the effect on the Ising measure of removing a wall cluster at a point $x \in \mathcal{L}_{0,n}$, and on the number of possible wall clusters through x .

Lemma 2.22 ([23, Lemma 3.10]). *There exists $C > 0$ such that for every $\beta > \beta_0$ the following holds. For every fixed \mathcal{I} , and every $x \in \mathcal{L}_{0,n}$, if \mathcal{J} is the interface obtained (per Lemma 2.14) by removing $\Theta_{\text{ST}}\text{Clust}(W_x)$ from the standard wall representation of \mathcal{I} , then*

$$\left| \log \frac{\mu_n^\mp(\mathcal{I})}{\mu_n^\mp(\mathcal{J})} + \beta \mathbf{m}(\text{Clust}(W_x)) \right| \leq C \mathbf{m}(\text{Clust}(W_x)).$$

Lemma 2.23 ([23, Lemma 3.13]). *Fix $x \in \mathcal{L}_{0,n}$, and fix any set of walls $(W_z)_{z \in A}$ for some $A \subset \mathcal{L}_{0,n}$ with $x \notin A$. There exists a universal constant $C > 0$ such that the number of wall clusters $\text{Clust}(W_x)$ compatible with $(W_z)_{z \in A}$, and having $\mathbf{m}(\text{Clust}(W_x)) = M$ is at most C^M .*

Combining the above, [23, Section 3] established rigidity in a set $S \subset \mathcal{L}_{0,n}$ even conditionally on the shape of the interface outside S . In order to phrase this formally, let us introduce some notation. For a collection of walls \mathbf{W} , let $\mathbf{I}_{\mathbf{W}}$ be the set of all interfaces having $\Theta_{\text{ST}}\mathbf{W}$ in its standard wall representation. Notice that if $\mathbf{W} = (W_z)_{z \notin S}$ and $\rho(\mathbf{W}) \cap S = \emptyset$, then for every $\mathcal{I} \in \mathbf{I}_{\mathbf{W}}$, there is a single ceiling $\mathcal{C}_{\mathbf{W}}$ having $\rho(\mathcal{C}_{\mathbf{W}}) \supset S$. In particular, given the event $\{\mathcal{I} \in \mathbf{I}_{\mathbf{W}}\}$, the minimal energy interface is then the one for which $W_x = \emptyset$ for all $x \in S$. The next results give exponential tails of walls of $x \in S$ and height oscillations above $x \in S$ about $\text{ht}(\mathcal{C}_{\mathbf{W}})$, conditionally on $\mathbf{I}_{\mathbf{W}}$.

Theorem 2.24 ([23, Thm. 3.1]). *There exists $C > 0$ so that for all $\beta > \beta_0$ the following holds. Fix any two admissible collections of walls (W_1, \dots, W_r) , and $\mathbf{W} = (W_z)_{z \in A}$ such that $\rho(\mathbf{W}) \cap \rho(\bigcup_{i \leq r} \dot{W}_i) = \emptyset$. Then*

$$\mu_n^\mp(\mathbf{I}_{W_1, \dots, W_r} \mid \mathbf{I}_{\mathbf{W}}) \leq \exp\left(-(\beta - C) \sum_{i \leq r} \mathbf{m}(W_i)\right).$$

As a consequence we obtain the following for nested sequences of walls inside S . In what follows, for a set $S \subset \mathcal{L}_{0,n}$ and a site $x \in S$, let $\mathfrak{W}_{x,S}$ be the collection of walls nesting x , and themselves fully nested in S . Noticing that the number of possible nested wall collections $\mathfrak{W}_{x,S}$ having $\mathbf{m}(\mathfrak{W}_{x,S}) = M$ is at most C^M for some universal constant C , and enumerating over these possible choices, we obtain the following corollary.

Corollary 2.25 ([23, Cor. 3.2]). *There exists $C > 0$ so that for all $\beta > \beta_0$ the following holds. Let $S_n \subset \mathcal{L}_{0,n}$ be such that $\mathcal{L}_0 \setminus S_n$ is connected, and let $\mathbf{W}_n = (W_z)_{z \notin S_n}$ be an admissible collection of walls with $\rho(\mathbf{W}_n) \subset S_n^c$. For every $x \in S_n$ and every $r \geq 1$,*

$$\mu_n^\mp(\mathbf{m}(\mathfrak{W}_{x,S_n}) \geq r \mid \mathbf{I}_{\mathbf{W}_n}) \leq \exp(-(\beta - C)r).$$

We will also use the following extension of Corollary 2.25 to the case of $x_1, \dots, x_N \in S_n$ for some $N \leq r$.

Corollary 2.26. *In the setting of Corollary 2.25, if $x_1, \dots, x_N \in S_n$ for some $N \leq r$ then we have*

$$\mu_n^\mp\left(\mathbf{m}\left(\bigcup_{i \leq N} \mathfrak{W}_{x_i, S_n}\right) \geq r \mid \mathbf{I}_{\mathbf{W}_n}\right) \leq \exp(-(\beta - C)r).$$

Proof. The number of possible collections of nested sequences of walls $\bigcup_{i \leq N} \mathfrak{W}_{x_i, S_n}$ having total excess energy $\mathbf{m}(\bigcup_{i \leq N} \mathfrak{W}_{x_i, S_n}) = M$ can be seen to be at most C^{N+M} for some universal constant C as follows: there are 2^{N+M} many ways to partition the excess energy into M_1, \dots, M_N amongst the N points, and then one can systematically enumerate over the wall collections $\mathfrak{W}_{x_i, S_n} \setminus \bigcup_{j < i} \mathfrak{W}_{x_j, S_n}$ having excess area M_i for each i , costing in total a further C^M . Summing Theorem 2.24 over all such collections and all $M \geq r$, one then obtains the desired. \blacksquare

2.10. The law of the maximum height oscillation in a ceiling. A key ingredient in our proof will be tail bounds on the maximum height oscillations inside a ceiling and conditionally on the interface outside that ceiling. Towards this we introduce some notation. For each face x , let $\text{ht}_{\mathcal{I}}(x) = \{h : (x_1, x_2, h) \in \mathcal{I}\}$. For a subset $S \subset \mathcal{L}_{0,n}$, define

$$M_S^\uparrow(\mathcal{I}) := \max_{f \in S} \max \text{ht}_{\mathcal{I}}(f), \quad \text{and} \quad M_S^\downarrow(\mathcal{I}) := \min_{f \in S} \min \text{ht}_{\mathcal{I}}(f).$$

We will be especially interested in studying the behavior of these quantities when conditioning on $\mathbf{I}_{\mathbf{W}_n}$ for some collection $\mathbf{W} = \{W_z : z \notin S\}$ such that $\rho(\mathbf{W}) \subset S^c$. In this situation, in the interface $\mathcal{I}_{\mathbf{W}}$, there is a single ceiling $\mathcal{C}_{\mathbf{W}}$ for which $S \subset \rho(\mathcal{C}_{\mathbf{W}})$; this will be the reference height with which the oscillations inside h will be compared: as such, define the maximal upwards and downwards oscillations about $\text{ht}(\mathcal{C}_{\mathbf{W}})$ as

$$\bar{M}_S^\uparrow(\mathcal{I}) = M_S^\uparrow(\mathcal{I}) - \text{ht}(\mathcal{C}_{\mathbf{W}}), \quad \text{and} \quad \bar{M}_S^\downarrow(\mathcal{I}) = \text{ht}(\mathcal{C}_{\mathbf{W}}) - M_S^\downarrow(\mathcal{I}).$$

Also, for a subset $S \subset \mathcal{L}_{0,n}$ (understood contextually), recall from [23] the following two events for $A \subset S$:

$$\mathcal{G}_A^{\mathbf{m}}(r) = \bigcap_{x \in A} \{\mathbf{m}(\mathfrak{W}_{x,S}) < r\}, \quad \mathcal{G}_A^{\mathfrak{D}}(r) = \bigcap_{x \in A} \{\text{diam}(\mathfrak{W}_{x,S}) < r\}. \quad (2.5)$$

We begin with providing certain cruder bounds on the behavior of $\bar{M}_S^\uparrow(\mathcal{I}), \bar{M}_S^\downarrow(\mathcal{I})$ quantities for arbitrary sets S , that can be proved using only the conditional exponential tail bounds obtained in Section 2.9.

The following bounds were primarily formulated in [23] for upwards oscillations $\bar{M}_{S_n}^\uparrow$, but since their application in this paper is for downwards oscillations, we express them for $\bar{M}_{S_n}^\downarrow$: the two formulations are equivalent by reflection symmetry of Λ_n about \mathcal{L}_0 .

Proposition 2.27 ([23, Prop. 6.1]). *There exist $\beta_0, C > 0$ so the following holds for all $\beta > \beta_0$. Let $S_n \subset \mathcal{L}_{0,n}$ be a simply-connected set such that $|S_n| \rightarrow \infty$ with n , and let $\mathbf{W}_n = \{W_z : z \notin S_n\}$ be such that $\rho(\mathbf{W}) \subset S_n^c$. Then for every $h = h_n \geq 1$,*

$$\mu_n^\mp\left(\bar{M}_{S_n}^\downarrow < h \mid \mathbf{I}_{\mathbf{W}_n}\right) \geq \mu_n^\mp\left(\mathcal{G}_{S_n}^{\mathbf{m}}(4h) \mid \mathbf{I}_{\mathbf{W}_n}\right) \geq \exp\left(-|S_n|e^{-(4\beta-C)h}\right), \quad (2.6)$$

and

$$\mu_n^\mp \left(\bar{M}_{S_n}^\downarrow < h \mid \mathbf{W}_n \right) \leq \exp \left(-|S_n| e^{-(4\beta+C)h} \right). \quad (2.7)$$

Moreover, identical bounds hold if $\bar{M}_{S_n}^\downarrow$ is replaced by $\bar{M}_{S_n}^\uparrow$.

For “nice” sets S , the above bounds on the distributions of $\bar{M}_{S_n}^\downarrow$ and $\bar{M}_{S_n}^\uparrow$ can be sharpened significantly so that the lower and upper bounds essentially match; indeed this was the main result of [23]. It relies on an identification of the exact rate of a deviation of height h above a point $x \in S_n$ as α_h from (1.4); without any conditioning that goes back to [21, 22]. To state this refined form, which is crucial to the identification of the height h_n^* in our main theorems, we first formalize what we mean by “nice” sets.

Definition (isoperimetric dimension of face sets). A simply-connected subset of faces $S \subset \mathcal{F}(\mathcal{L}_{0,n})$ is said to have isoperimetric dimension at most d , denoted $\dim_{\text{ip}}(S) \leq d$, if $|\partial S| \leq |S|^{(d-1)/d}$.

For sets S_n with a uniformly bounded $\dim_{\text{ip}}(S_n)$ we have the following much finer estimates.

Proposition 2.28 ([23, Prop. 6.2]). *There exist $\beta_0, \kappa_0 > 0$ such that the following holds for every fixed $\beta > \beta_0$. Let $S_n \subset \mathcal{L}_{0,n}$ be a sequence of simply-connected sets such that $\dim_{\text{ip}}(S_n) \leq \sqrt{\beta}$ and $\lim_{n \rightarrow \infty} |S_n| = \infty$. Let $\mathbf{W}_n = (W_z)_{z \notin S_n}$ be such that $\rho(\mathbf{W}_n) \subset S_n^c$. For every $\sqrt{\log |S_n|} \leq h \leq \frac{1}{\sqrt{\beta}} \log |S_n|$,*

$$\mu_n^\mp \left(\bar{M}_{S_n}^\downarrow < h \mid \mathbf{W}_n \right) \geq \mu_n^\mp \left(\bar{M}_{S_n}^\downarrow < h, \mathcal{G}_{S_n \setminus S_{n,h}^\circ}^{\text{m}}(4h), \mathcal{G}_{S_{n,h}^\circ}^{\text{m}}(5h) \mid \mathbf{W}_n \right) \geq \exp \left(-(1 + \varepsilon_\beta) |S_n| e^{-\alpha_h} \right), \quad (2.8)$$

and

$$\mu_n^\mp \left(\bar{M}_{S_n}^\downarrow < h, \mathcal{G}_{S_{n,h}^\circ}^{\mathfrak{D}}(e^{\kappa_0 h}) \mid \mathbf{W}_n \right) \leq \exp \left(-(1 - \varepsilon_\beta) |S_n| e^{-\alpha_h} \right), \quad (2.9)$$

where $S_{n,h}^\circ := \{x \in S_n : d(x, \partial S_n) \geq e^{2\kappa_0 h}\}$. Moreover, identical bounds hold if $\bar{M}_{S_n}^\downarrow$ is replaced by $\bar{M}_{S_n}^\uparrow$.

3. EXTENSION OF ESTIMATES WITHIN A CEILING TO BE CONDITIONAL ON NO BIG WALLS

In this section we use the observation that the proof of Theorem 2.24 is monotone w.r.t. wall sizes since it hinges on maps that only delete walls (see (3.2) below) in order to elevate Theorem 2.24, and consequently also the lower bounds in Propositions 2.27 and 2.28, to apply conditionally on events of the following form:

$$\mathcal{E}_S^\eta := \bigcap_{z \in S} \{m(W_z) \leq \eta_z\} \quad \text{for } \eta = \{\eta_z\}_{z \in S}. \quad (3.1)$$

This extension will be important in the proof of (1.6) as hinted in Section 1.2.

The proof of Theorem 2.24 in [23] relied on a Peierls-type argument: a map Φ that deletes wall clusters. Extending it to be conditional on \mathcal{E}_S^η rests on the fact that Φ clearly satisfies

$$m(W_z(\Phi(\mathcal{I})) \leq m(W_z(\mathcal{I})) \quad \text{for every } z \text{ and every } \mathcal{I}, \quad (3.2)$$

and so, in particular, if $\mathcal{I} \in \mathcal{E}_S^\eta$ for some η then necessarily $\Phi(\mathcal{I}) \in \mathcal{E}_S^\eta$. We now formalize the extension and include a proof for completeness.

Theorem 3.1. *In the setting of Theorem 2.24, for any $\eta = \{\eta_z\}_{z \in A^c}$ we have*

$$\mu_n^\mp(\mathbf{W}_{1,\dots,W_r} \mid \mathbf{W}, \mathcal{E}_{A^c}^\eta) \leq \exp \left(-(\beta - C) \sum_{i \leq r} m(W_i) \right).$$

Proof. Consider a map $\Phi = \Phi_{(W_i)_i}$ which takes an interface $\mathcal{I} \in \mathbf{W}_{1,\dots,W_r} \cap \mathbf{W} \cap \mathcal{E}_{A^c}^\eta$, deletes $\bigcup_i \Theta_{\text{ST}} \text{Clust}(W_i)$ from its standard wall representation, and from the resulting standard wall collection, constructs the interface $\Phi(\mathcal{I})$. On the one hand, the map satisfies a weight gain of (see [23, Lemma 3.10])

$$\left| \log \frac{\mu_n^\mp(\mathcal{I})}{\mu_n^\mp(\Phi(\mathcal{I}))} + \beta m(\mathcal{I}; \Phi(\mathcal{I})) \right| \leq C m(\mathcal{I}; \Phi(\mathcal{I})) = C m \left(\bigcup_i \text{Clust}(W_i) \right).$$

On the other hand, for fixed $(W_i)_i$, the number of sets $\bigcup_i \text{Clust}(W_i)$ is bounded by $\exp(C m(\bigcup_i \text{Clust}(W_i)))$ (see [23, Lemma 3.13]). Furthermore, the resulting interface $\Phi(\mathcal{I})$ belongs to $\mathbf{W} \cap \mathcal{E}_{A^c}^\eta$ since it only decreases

the wall collection, and only deletes walls confined to A^c . If we let $M = \mathbf{m}(\bigcup_i \text{Clust}(W_i))$ and apply the map Φ pointwise to $\mathcal{I} \in \mathbf{I}_{W_1, \dots, W_r} \cap \mathbf{I}_{\mathbf{W}} \cap \mathcal{E}_{A^c}^\eta$, mapping them to interfaces $\mathcal{J} \in \mathbf{I}_{\mathbf{W}} \cap \mathcal{E}_{A^c}^\eta$, we obtain

$$\mu_n^\mp(\mathbf{I}_{W_1, \dots, W_r}, \mathbf{I}_{\mathbf{W}}, \mathcal{E}_{A^c}^\eta) \leq \sum_{\mathcal{J} \in \mathbf{I}_{\mathbf{W}}, \mathcal{E}_{A^c}^\eta} \mu_n^\mp(\mathcal{J}) \sum_{k \geq M} e^{-(\beta-C)k} \leq e^{-(\beta-C)M} \mu_n^\mp(\mathbf{I}_{\mathbf{W}}, \mathcal{E}_{A^c}^\eta).$$

Dividing both sides by the probability on the right-hand side then yields the desired conditional probability bound. \blacksquare

Corollary 3.2. *In the setting of Corollary 2.25, for any $\eta = \{\eta_z\}_{z \in S_n}$ we have*

$$\mu_n^\mp(\mathbf{m}(\mathfrak{W}_{x, S_n}) \geq r \mid \mathbf{I}_{\mathbf{W}_n}, \mathcal{E}_{S_n}^\eta) \leq \exp(-(\beta-C)r).$$

We stress that the monotonicity in excess energy of Φ does not carry to the more complicated maps in [21–23] where the quantities α_h governing the large deviation rates in Proposition 2.28 arise. Still, we are able to extract the following extension of the lower bounds (2.6) and (2.8) stated in Propositions 2.27 and 2.28, which will be needed in Section 6. (In what follows, we retain the definition of $S_{n,h}^\circ$ from the latter proposition in terms of $2\kappa_0$ rather than κ_0 —placed there due to the role of κ_0 in the upper bound (2.9), which we do not include here—for the sake of consistency.)

Proposition 3.3. *There exist $\beta_0, C, \kappa_0 > 0$ so the following holds for all $\beta > \beta_0$. Let $S_n \subset \mathcal{L}_{0,n}$ be a simply-connected set such that $|S_n| \rightarrow \infty$ with n , and let $\mathbf{W}_n = \{W_z : z \notin S_n\}$ be such that $\rho(\mathbf{W}_n) \subset S_n^c$.*

(1) *For every $h = h_n \geq 1$ and every η ,*

$$\mu_n^\mp(\bar{M}_{S_n}^\downarrow < h \mid \mathbf{I}_{\mathbf{W}_n}, \mathcal{E}_{S_n}^\eta) \geq \exp(-|S_n|e^{-(4\beta-C)h}), \quad (3.3)$$

(2) *Further suppose that $\dim_{\text{ip}}(S_n) \leq \sqrt{\beta}$. For every $\sqrt{\log|S_n|} \leq h \leq \frac{1}{\sqrt{\beta}} \log|S_n|$ and every η satisfying that $\eta_z \geq 5h$ for all $z \in S_{n,h}^\circ := \{x \in S_n : d(x, \partial S_n) \geq e^{2\kappa_0 h}\}$,*

$$\mu_n^\mp(\bar{M}_{S_n}^\downarrow < h \mid \mathbf{I}_{\mathbf{W}_n}, \mathcal{E}_{S_n}^\eta) \geq \exp(-(1 + \varepsilon_\beta)|S_n|e^{-\alpha h}). \quad (3.4)$$

Proof. The proof will follow the same arguments used to establish (2.6), and (2.8) in [23, Proofs of (6.1) and (6.3)]; hence we will only describe the modifications required to adapt it to the conditional space given $\mathcal{E}_{S_n}^\eta$.

The proof of (2.6) used the fact that $\{\bar{M}_{S_n}^\downarrow < h\} \supset \{\mathcal{G}_{S_n}^{\text{m}}(4h)\}$ (defined in (2.5)) and shows that

$$\mu_n^\mp(\mathcal{G}_{S_n}^{\text{m}}(4h) \mid \mathbf{I}_{\mathbf{W}_n}) \geq \exp(-|S_n|e^{-(4\beta-C)h})$$

by iteratively exposing \mathfrak{G}_{x, S_n} (the sequence of nested walls \mathfrak{W}_{x, S_n} which nest x in S_n , as well as the walls nested in this sequence of walls) along a corresponding filtration \mathcal{F}_i . One uses the fact that

$$\mathcal{G}_{S_n}^{\text{m}}(4h) = \bigcap_{x \in S_n} \hat{G}_x, \quad \text{where} \quad \hat{G}_x = \bigcap_{u \in \rho(\mathfrak{G}_{x, S_n})} \{\mathbf{m}(\mathfrak{W}_{u, S_n}) < 4h\}, \quad (3.5)$$

and that, when revealing \mathfrak{G}_{x, S_n} , either $\text{diam}(\mathfrak{G}_{z_i, S_n}) \geq 4h$, which has probability at most $e^{-(4\beta-C)h}$ under $\mu_n^\mp(\cdot \mid \mathcal{F}_i)$, or there are $O(h^2)$ sites $z \in \rho(\mathfrak{G}_{x, S_n})$, and $\mu_n^\mp(\mathbf{m}(\mathfrak{W}_{z, S_n}) \geq 4h \mid \mathcal{F}_i) \leq e^{-(4\beta-C)h}$ for any given z , both bounds due to Corollary 2.25. Combined, $\mu_n^\mp(\hat{G}_x \mid \mathcal{F}_i) \geq 1 - O(h^2 e^{-(4\beta-C)h}) \geq \exp(-e^{-(4\beta-C)h})$. Applying the exact same argument while appealing to Corollary 3.2 in lieu of Corollary 2.25 now shows that (3.3) holds conditionally on $\mathcal{E}_{S_n}^\eta$.

For the proof of (2.8), we first tiled $S_{n,h}^\circ$ by $L \times L$ boxes Q_i with $L \sim \frac{1}{2}e^{2\kappa_0 h}$, with Q denoting their union. We then applied the above analysis to bound the probabilities of \hat{G}_z for $z \in S \setminus Q$. Thereafter, breaking Q_i into its bulk $Q'_i \subset Q_i$ and remainder $Q_i \setminus Q'_i$, one gives a lower bound on an event

$$D_i := \bigcap_{z \in Q'_i} \hat{H}_z \cap \bigcap_{z \in Q_i \setminus Q'_i} \hat{G}_z^{\text{m}}(4h)$$

where \hat{H}_z is some event involving the pillar of z in S_n and is a subset of $\{\mathbf{m}(\mathfrak{W}_{z, S_n}) < 5h\}$. In light of this fact, D_i is already a subset of the event $\bigcap_{z \in Q_i} \{\mathbf{m}(\mathfrak{W}_{z, S_n}) < 5h\}$. Since $\eta_z \geq 5h$ for all $z \in S_{n,h}^\circ$, this means D_i is already a subset of the event $\mathcal{E}_{S_n}^\eta$, and thus

$$\mu_n^\mp(D_i \mid \mathcal{F}_i, \mathcal{E}_{S_n}^\eta) \geq \mu_n^\mp(D_i \mid \mathcal{F}_i).$$

With this observation, the bounds from the proof in [23] can be applied as is. \blacksquare

4. BASIC BOUNDS USING THE ISING MEASURE WITHOUT A FLOOR

The aim of this section is to establish a priori estimates on the number of total wall faces in the interface, the number of wall faces in any collection of $N \leq n$ distinct walls, and the number of faces in ceilings that are small, say smaller than $n^{1.9}$, or not thick in that they have isoperimetric dimension much larger than 2. To prove these we first lower bound the probability that the interface lies above a floor at height 0, which gives one side of the bound of Proposition 1.6.

4.1. A map to lift the entire interface by k . We begin by introducing a simple but recurring tool we will use: a map to lift the entirety of the interface up by a height of k . Unlike height function models, even this simple map cannot be done in a bijective manner at the level of Ising interfaces, but the multiplicity of the map and the interactions it induces in the interface via the \mathbf{g} term in Theorem 2.16 are mild.

Definition 4.1. For any $k \geq 1$, define the map Φ_k^\uparrow as the following map: for an interface \mathcal{I} ,

- (1) Let $\theta_\uparrow \mathcal{I}$ be the shift of \mathcal{I} by the vector $(0, 0, k)$;
- (2) Let $B_k = \mathcal{F}(\partial \mathcal{L}_{0,n} \times \llbracket 0, k \rrbracket)$, and let $\mathcal{I}' = \theta_\uparrow \mathcal{I} \oplus B_k$ (where \oplus denotes the symmetric difference);
- (3) Let $\Phi_k^\uparrow(\mathcal{I})$ be given by taking \mathcal{I}' and removing all finite $*$ -connected components of $\mathcal{I}' \cup (\mathcal{L}_0 \setminus \mathcal{L}_{0,n})$.

The main use of this map is the following estimate on its effect on the mass of any set of interfaces when passed through Φ_k^\uparrow , showing it only costs roughly $\exp(4\beta kn)$ to lift the interface by k .

Proposition 4.2. Fix $\beta > \beta_0$. Consider any set of interfaces A , and any $k \geq 1$. Then,

$$\frac{\mu_n^\mp(A)}{\mu_n^\mp(\Phi_k^\uparrow(A))} \leq \exp(4(\beta + C)kn).$$

Moreover, this holds under $\hat{\mu}_n^{\mathbf{h}}$ for any $\mathbf{h} \geq 0$.

We begin by arguing that the map Φ_k^\uparrow is well-defined and describing some of its key properties.

Lemma 4.3. For every $k \geq 1$, Φ_k^\uparrow is well-defined on the set of all interfaces in Λ_n . For every \mathcal{I} it has $\mathbf{m}(\Phi_k^\uparrow(\mathcal{I}); \mathcal{I}) \leq 4kn$. Furthermore, if $\mathcal{I} \in \mathbf{I}_h^{\text{fl}}$, then $\Phi_k^\uparrow(\mathcal{I}) \setminus B_k \subset \mathcal{L}_{\geq k-h}$, and for every ceiling $\mathcal{C} \in \mathfrak{C}(\mathcal{I})$, its shift up by height k is a subset of a ceiling of $\Phi_k^\uparrow(\mathcal{I})$.

Proof. We begin with the well-definedness of Φ_k^\uparrow on the set of all interfaces on Λ_n . Consider the spin configuration obtained by taking $\sigma(\mathcal{I})$ and shifting it by $(0, 0, k)$; this is exactly the spin configuration whose set of separating faces are $\mathcal{I}' \oplus B_k$. Taking that spin configuration and flipping the spins in all *bubbles*, i.e., finite $*$ -connected sets of minus or plus spins, then yields a configuration $\Phi_k^\uparrow(\sigma(\mathcal{I}))$, say, whose set of separating faces is exactly $\Phi_k^\uparrow(\mathcal{I})$. Moreover, that set of separating faces corresponds to an interface per Remark 2.2 as the spin configuration has no finite $*$ -connected plus or minus components.

Next we turn to the excess energy generated by the map. Notice that the map consists of the addition of at most $4kn$ faces in step (2), and then removes some faces in step (3). As such, we have for every \mathcal{I} that $\mathbf{m}(\Phi_k^\uparrow(\mathcal{I}); \mathcal{I}) = |\Phi_k^\uparrow(\mathcal{I})| - |\mathcal{I}| \leq 4kn$.

Turning now to the properties of the range of the map, consider any $\mathcal{I} \in \mathbf{I}_h^{\text{fl}}$. The map Φ_k^\uparrow raises all faces of \mathcal{I} by k , and only *adds* faces of B_k . Namely, $\Phi_k^\uparrow(\mathcal{I}) \subset \theta_\uparrow \mathcal{I} \cup B_k$, and $\theta_\uparrow \mathcal{I} \in \mathbf{I}_{h-k}^{\text{fl}}$.

To prove the last property, recall that a horizontal face f is a ceiling face in \mathcal{I} if and only if all spins of $\sigma(\mathcal{I})$ in the column through f are plus below f , and minus above f (see [21, Observation 4.10]). Consider a ceiling face $f \in \mathcal{I}$; from earlier in this proof, in the spin configuration $\sigma(\Phi_k^\uparrow(\mathcal{I}))$, this entire column is shifted up vertically by k (shifting f to $f + (0, 0, k)$) and then all bubbles in the resulting spin configuration are deleted. However, no bubbles can intersect the column through $f + (0, 0, k)$ because if they did, that would violate the property that all spins below $f + (0, 0, k)$ are plus and all spins above are minus. As such, no spins are flipped in this column, and the face $f + (0, 0, k)$ must be a ceiling face of $\Phi_k^\uparrow(\mathcal{I})$. If the set of ceiling faces of \mathcal{I} are a subset of those of $\Phi_k^\uparrow(\mathcal{I})$, then evidently any connected component of them in \mathcal{I} will be a subset of a connected component of them in $\Phi_k^\uparrow(\mathcal{I})$, concluding the proof. \blacksquare

We now turn to the proof of Proposition 4.2.

Proof of Proposition 4.2. The proof of the proposition consists of taking into account the weight change from application of the map Φ_k^\uparrow , as well as its multiplicity.

We begin with considering the change in the weight of an interface in A when mapped through Φ_k^\uparrow . More precisely, we first show that for every \mathcal{I} and every $k \geq 1$, we have

$$\left| \log \frac{\mu_n^\mp(\Phi_k^\uparrow(\mathcal{I}))}{\mu_n^\mp(\mathcal{I})} + \beta \mathbf{m}(\Phi_k^\uparrow(\mathcal{I}); \mathcal{I}) \right| \leq C(kn + |\mathbf{m}(\Phi_k^\uparrow(\mathcal{I}); \mathcal{I})|). \quad (4.1)$$

To show (4.1), fix any \mathcal{I} and for ease of notation let $\mathcal{J} = \Phi_k^\uparrow(\mathcal{I})$. For $f \in \mathcal{I}$, let $\theta_\uparrow f$ be its vertical shift by k . By Theorem 2.16, the left-hand side of (4.1) is at most

$$\begin{aligned} \left| \sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I}) - \sum_{f' \in \mathcal{J}} \mathbf{g}(f', \mathcal{J}) \right| &\leq \sum_{f: \theta_\uparrow f \in \theta_\uparrow \mathcal{I} \setminus \mathcal{J}} |\mathbf{g}(f, \mathcal{I})| + \sum_{f' \in \mathcal{J} \setminus \theta_\uparrow \mathcal{I}} |\mathbf{g}(f', \mathcal{J})| \\ &\quad + \sum_{f: \theta_\uparrow f \in \theta_\uparrow \mathcal{I} \cap \mathcal{J}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\theta_\uparrow f, \mathcal{J})|. \end{aligned}$$

By (2.2), the first sum on the right-hand side is at most $\bar{K}|\theta_\uparrow \mathcal{I} \setminus \mathcal{J}|$ and the second sum is at most $\bar{K}|\mathcal{J} \setminus \theta_\uparrow \mathcal{I}|$, so that together, they contribute at most $\bar{K}|\theta_\uparrow \mathcal{I} \oplus \mathcal{J}|$. Turning to the last term, by (2.3), and the fact that the radius $\mathbf{r}(f, \mathcal{I}; \theta_\uparrow f, \mathcal{J})$ is the same as $\mathbf{r}(\theta_\uparrow f; \theta_\uparrow \mathcal{I}; \theta_\uparrow f, \mathcal{J})$ which must be attained at a face of $\theta_\uparrow \mathcal{I} \oplus \mathcal{J}$,

$$\sum_{f: \theta_\uparrow f \in \theta_\uparrow \mathcal{I} \cap \mathcal{J}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\theta_\uparrow f, \mathcal{J})| \leq \sum_{f' \in \theta_\uparrow \mathcal{I} \cap \mathcal{J}} \bar{K} e^{-\bar{c}\mathbf{r}(f'; \theta_\uparrow \mathcal{I}; f', \mathcal{J})} \leq \sum_{f' \in \theta_\uparrow \mathcal{I} \cap \mathcal{J}} \sum_{g \in \theta_\uparrow \mathcal{I} \oplus \mathcal{J}} \bar{K} e^{-\bar{c}d(f', g)}.$$

This is evidently at most $C\bar{K}|\theta_\uparrow \mathcal{I} \oplus \mathcal{J}|$. To conclude, notice that

$$|\theta_\uparrow \mathcal{I} \oplus \mathcal{J}| = |\mathcal{J} \setminus \theta_\uparrow \mathcal{I}| + |\theta_\uparrow \mathcal{I} \setminus \mathcal{J}|, \quad \text{and} \quad \mathbf{m}(\mathcal{J}; \mathcal{I}) = |\mathcal{J} \setminus \theta_\uparrow \mathcal{I}| - |\theta_\uparrow \mathcal{I} \setminus \mathcal{J}|.$$

Rearranging the latter equality and taking absolute values, we see that $|\theta_\uparrow \mathcal{I} \setminus \mathcal{J}| \leq |\mathcal{J} \setminus \theta_\uparrow \mathcal{I}| + |\mathbf{m}(\mathcal{J}; \mathcal{I})|$. Since $|\mathcal{J} \setminus \theta_\uparrow \mathcal{I}| \leq 4kn$, we then get $|\theta_\uparrow \mathcal{I} \oplus \mathcal{J}| \leq 8kn + |\mathbf{m}(\mathcal{J}; \mathcal{I})|$, yielding (4.1).

We next consider the multiplicity of the map Φ_k^\uparrow . The key claim here is that for every $M \geq 0$, for every \mathcal{J} in the range of Φ_k^\uparrow , we have

$$|\{\mathcal{I} \in (\Phi_k^\uparrow)^{-1}(\mathcal{J}) : |\mathbf{m}(\mathcal{J}; \mathcal{I})| = M\}| \leq \exp(C(kn + M)). \quad (4.2)$$

To show this, consider the set of faces $\mathcal{B} = \mathcal{I}' \setminus \mathcal{J}$, i.e., all bubbles deleted in step (3) of Definition 4.1. We first claim that this serves as a witness to the pre-image \mathcal{I} , i.e., for every fixed \mathcal{J} , given $\mathcal{B} = \mathcal{B}(\mathcal{I})$ one can uniquely reconstruct \mathcal{I} . This is evidently done by noticing that $\mathcal{I}' = \mathcal{J} \cup \mathcal{B}$, then taking $\mathcal{I}' \oplus B_k$, and shifting it by the vector $(0, 0, -k)$ to obtain \mathcal{I} .

Moreover, for an interface \mathcal{I} having $|\mathbf{m}(\mathcal{J}; \mathcal{I})| = M$, we must have $|\mathcal{B}| \leq 4kn + M$ since

$$|\mathcal{I}'| - |\mathcal{I}| \leq 4kn, \quad \text{and} \quad \mathbf{m}(\mathcal{J}; \mathcal{I}) = |\mathcal{I}'| - |\mathcal{I}| - |\mathcal{B}|.$$

It therefore suffices for us to enumerate over possible face sets \mathcal{B} having at most $4kn + M$ many faces. We first note that every $*$ -connected component of \mathcal{B} must intersect B_k ; indeed since $\theta_\uparrow \mathcal{I} \oplus \mathcal{I}' \subset B_k$, if a $*$ -connected component of \mathcal{B} doesn't intersect B_k , then its shift by $(0, 0, -k)$ was a finite $*$ -connected component of \mathcal{I} bounding a bubble, and therefore \mathcal{I} would not have been an admissible interface.

In order to enumerate over the choice of \mathcal{B} , we first choose some $N \leq M$ for the number of $*$ -connected components of \mathcal{B} and choose a subset of N faces f_1, \dots, f_N amongst B_k for some representative face of B_k for each of those N components. We then choose corresponding values M_1, \dots, M_N such that $\sum M_i \leq 4kn + M$ dictating how many faces belong to each of those components, and finally for each f_i , we enumerate over the possible $*$ -connected components of faces in $\mathcal{F}(\Lambda_n)$ of size M_i containing f_i . In total, this enumeration counts at most

$$M \binom{4kn}{M} 2^{N+4kn+M} C^{4kn+M}$$

many such choices, which is at most C^{4kn+M} for some other C , establishing (4.2).

We are now in position to conclude the proof of the proposition. For a set A of interfaces, we can rewrite

$$\mu_n^\mp(A) = \sum_{\mathcal{I} \in A} \mu_n^\mp(\mathcal{I}) = \sum_{\mathcal{J} \in \Phi_k^\uparrow(A)} \mu_n^\mp(\mathcal{J}) \sum_{-\infty < M \leq 4kn} \sum_{\substack{\mathcal{I} \in A \cap (\Phi_k^\uparrow)^{-1}(\mathcal{J}) \\ \mathfrak{m}(\mathcal{J}; \mathcal{I}) = M}} \frac{\mu_n^\mp(\mathcal{I})}{\mu_n^\mp(\mathcal{J})}.$$

Applying (4.1) and (4.2), we find

$$\mu_n^\mp(A) \leq \mu_n^\mp(\Phi_k^\uparrow(A)) \sum_{-\infty < M \leq 4kn} e^{\beta M + C(kn + |M|)} \leq e^{(4\beta + C)kn} \mu_n^\mp(\Phi_k^\uparrow(A)).$$

Dividing both sides out by $\mu_n^\mp(\Phi_k^\uparrow(A))$ yields the desired result. The fact that this holds under $\hat{\mu}_n^h$ follows from the fact that if $A \subset \mathbf{I}_h^{\text{fl}}$, then $\Phi_k^\uparrow(A) \subset \mathbf{I}_h^{\text{fl}}$ per Lemma 4.3. \blacksquare

4.2. Lower bound on the probability of lying above the floor. With the map of Definition 4.1 in hand, we can now prove the following simple lemma giving a lower bound on the probability of the interface under μ_n^\mp to lie above height 0. This will give the lower bound from Proposition 1.6, whereas the matching upper bound up to a factor of $1 + \varepsilon_\beta$ will be a consequence of our estimates in §5 (namely, Theorem 5.2).

Lemma 4.4. *There exists $\beta_0 > 0$ and a sequence $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow \infty$ such that, for every $\beta > \beta_0$,*

$$\mu_n^\mp(\mathbf{I}_0^{\text{fl}}) \geq \exp(-(1 + \varepsilon_\beta)n \log n).$$

Proof. Let $h_0 = (4\beta - C_0)^{-1} \log n$, where $C_0 > 0$ is the absolute constant from Proposition 2.27. Recalling that $\mathbf{I}_{h_0}^{\text{fl}} = \{\mathcal{I} : M_n^\downarrow \leq h_0\}$, we have by (2.6) (with $S_n = \mathcal{L}_{0,n}$) that

$$\mu_n^\mp(\mathbf{I}_{h_0}^{\text{fl}}) \geq \exp(-n^2 e^{-(4\beta - C_0)h_0}) = e^{-n}.$$

At the same time, we can compare the weight of $\mathbf{I}_{h_0}^{\text{fl}}$ to that of \mathbf{I}_0^{fl} by application of the map $\Phi_{h_0}^\uparrow$ since we have $\Phi_{h_0}^\uparrow(\mathbf{I}_{h_0}^{\text{fl}}) \in \mathbf{I}_0^{\text{fl}}$ by Lemma 4.3. In particular, by Proposition 4.2,

$$\mu_n^\mp(\mathbf{I}_{h_0}^{\text{fl}}) \leq e^{(4\beta + C')h_0 n} \mu_n^\mp(\mathbf{I}_0^{\text{fl}}).$$

Combining this with the aforementioned lower bound of $\mu_n^\mp(\mathbf{I}_{h_0}^{\text{fl}}) \geq e^{-n}$ concludes the proof. \blacksquare

4.3. Wall faces in total and in linearly many walls. Our first application of Lemma 4.4 will establish that the interface $\mathcal{I} \sim \hat{\mu}_n^h$ must contain at most $e^{-2\beta} n^2$ wall faces w.h.p., as given by the following lemma. (N.B. it is easy to infer the weaker upper bound of $(C/\beta)n^2$ on the number of such faces from the cluster expansion representation of Theorem 2.16. To do so, one compares all such interfaces to the completely flat one (that is, analyze the energy gain and multiplicity loss incurred in the map that deletes every wall), with a multiplicity of $\exp(C(n^2 + m))$ for interfaces with m excess faces competing with a probability gain of $\exp(-(\beta - C)m)$.)

Lemma 4.5. *There exist $\beta_0, C > 0$ such that, for every $\beta > \beta_0$,*

$$\mu_n^\mp \left(\sum \{\mathfrak{m}(W) : W \text{ is a wall in } \mathcal{I}\} \geq e^{-2\beta} n^2 \right) \leq \exp(-C e^{-4\beta} n^2).$$

Consequently, the same bound holds true under $\hat{\mu}_n^h$ for any $h \geq 0$.

Proof. Order the faces of $\mathcal{L}_{0,n}$ using labels $1, \dots, n^2$ by scanning $\mathcal{L}_{0,n}$ in a connected manner, row by row. Roughly put, we will reveal $\{W_x\}_{x \in \mathcal{L}_{0,n}}$ in this order, skipping any x that belongs to $\rho(\dot{W}_{x'})$ for some already revealed $W_{x'}$. Formally:

1. Initialize $F = \mathcal{L}_{0,n}$ as the set of unexplored faces.

2. While $F \neq \emptyset$, repeatedly reveal W_x for $x \in F$ having the smallest label in F , and delete $\rho(\dot{W}_x)$ from F .

(Note that, in this way, our set of revealed walls remains always connected to $\partial\Lambda_n$.) This process terminates at $F = \emptyset$ with a collection of nonempty walls $W_x \neq \emptyset$ whose hull is yet unexplored; processing these walls in an arbitrary order, apply to each such W_x the same exposure procedure (Step 2) from an initial $F = \rho(\dot{W}_x)$. (As before, the set of revealed walls is always connected to $\partial\Lambda_n$.) Let $(\mathcal{F}_j)_{j=1}^{n^2}$ be the corresponding filtration.

If x_j is about to be revealed in step j of the process, and S_{j-1} is the connected component of faces that contains x_j in $\mathcal{L}_{0,n} \setminus \rho(\bigcup_{k < j} (W_{x_k} \cup \{x_k\}))$, then the fact that the projections of the walls revealed thus far are connected to $\partial\mathcal{L}_{0,n}$ supports an application of Corollary 2.25, which implies that

$$\mu_n^\mp(\mathbf{m}(W_{x_j}) \geq r \mid \mathcal{F}_{j-1}) \leq \mu_n^\mp(\mathbf{m}(\mathfrak{W}_{x_j, S_{j-1}}) \geq r \mid \mathcal{F}_{j-1}) \leq e^{-(\beta-C)r}.$$

It follows that $\sum_i \mathbf{m}(W_{x_i})$ is stochastically dominated by $\sum_{i=1}^{n^2} (\xi_i - 1)$ where the ξ_i 's are i.i.d. Geometric(p) for $p = 1 - e^{-(4\beta-C)}$, using that every $W_x \neq \emptyset$ has $\mathbf{m}(W_x) \geq 4$. Thus, bounding $\sum_i \xi_i$ via the correspondence between the negative binomial and binomial distributions,

$$\mu_n^\mp \left(\sum \{ \mathbf{m}(W) : W \text{ is a wall in } \mathcal{I} \} > \lfloor e^{-2\beta} n^2 \rfloor \right) \leq \mathbb{P}(\text{Bin}(\lfloor (1 + e^{-2\beta}) n^2 \rfloor, p) < n^2).$$

The latter binomial random variable has mean $\mu \in [(1 + \frac{1}{2}e^{-2\beta})n^2, (1 + e^{-2\beta})n^2]$ provided β is large enough, whence the probability that it is less than $\mu - a$ for $a = \frac{1}{2}e^{-2\beta}n^2$ is at most $\exp(-\frac{1}{2}a^2/\mu)$. This establishes the claimed bound for $\sum_W \mathbf{m}(W)$ under μ_n^\mp , and the analogous bound under $\hat{\mu}_n^h$ follows from Lemma 4.4. ■

The next application of Lemma 4.4 rules out having n distinct walls, each with diameter at least $\log n$.

Lemma 4.6. *There exist $\beta_0, C > 0$ so that the following holds for every $\beta > \beta_0$. Let \mathfrak{E}_n be the event that there exist $N \leq n$ distinct faces $x_1, \dots, x_N \in \mathcal{L}_{0,n}$ such that $\mathbf{m}(\bigcup_i \mathfrak{W}_{x_i}) \geq (6/\beta)n \log n$ in the interface \mathcal{I} . Then $\mu_n^\mp(\mathfrak{E}_n) \leq \exp(-(5 - \varepsilon_\beta)n \log n)$, and consequently, $\hat{\mu}_n^h(\mathfrak{E}_n) \leq \exp(-3n \log n)$.*

Proof. Fix a set A of at most n distinct faces in $\mathcal{L}_{0,n}$. We infer from Corollary 2.26 that

$$\mu_n^\mp(\mathbf{m}(\bigcup_{x \in A} \mathfrak{W}_x) \geq (6/\beta)n \log n) \leq e^{-(\beta-C)(6/\beta)n \log n} \leq e^{-(6-\varepsilon_\beta)n \log n}$$

provided β_0 (thus also β) is large enough. Substituting the fact that $\binom{n^2}{|A|} \leq \binom{n^2}{n} \leq \exp(n \log n + n)$ yields

$$\mu_n^\mp(\mathfrak{E}_n) \leq \sum_{N \leq n} \binom{n^2}{N} e^{-(6-\varepsilon_\beta)n \log n} \leq e^{-(5-\varepsilon'_\beta)n \log n},$$

and the conclusion for $\hat{\mu}_n^h$ follows from Lemma 4.4. ■

4.4. Ceilings with a sub-linear area or a large isoperimetric dimension. We will next use Lemma 4.4 to rule out an excessive total area in ceilings whose hull satisfies $|\mathcal{C}| \leq c_\beta n^2 / \log^2 n$ (Lemma 4.7), as well as in ceilings with $|\mathcal{C}| \geq n$ and $\dim_{\text{ip}} \geq 4$ (Lemma 4.8), by showing the corresponding events have probabilities smaller than $\exp(-(1 + \varepsilon_\beta)n \log n)$. (N.B. that the threshold $|\mathcal{C}| = O(n^2 / \log^2 n)$ is correct when aiming at an error probability of $\exp(-cn \log n)$ for a total excess area of cn^2 , as the probability of finding a single such ceiling is some $\exp(-cn / \log n)$, and $c \log^2 n$ many such ceilings would be needed for a total area of cn^2 .)

Lemma 4.7. *There exist $\beta_0, C > 0$ so that for every fixed $\beta > \beta_0$ and every $1 \leq A \leq 2 \log_2 n - \sqrt{\log n}$,*

$$\mu_n^\mp \left(\sum \{ |\mathcal{C}| : \mathcal{C} \text{ is a ceiling with } 2^{A-1} \leq |\mathcal{C}| \leq \left(\frac{n}{e^\beta \log n} \right)^2 \} \geq \frac{C}{e^\beta A} n^2 \right) < e^{-5n \log n},$$

and the same bound holds under $\hat{\mu}_n^h$.

Proof. We begin by ruling out the contribution of ceilings \mathcal{C} with $|\mathcal{C}| \leq L_1$ for $L_1 := n^2 / \log^8 n$ via the bound

$$\mu_n^\mp \left(\sum \{ |\mathcal{C}| : \mathcal{C} \text{ is a ceiling with } 2^{A-1} \leq |\mathcal{C}| \leq L_1 \} \geq \frac{C}{e^\beta A} n^2 \right) \leq \exp(-\varepsilon_\beta n \log^2 n), \quad (4.3)$$

for some sequence ε_β vanishing as $\beta \rightarrow \infty$. Partitioning the set of ceilings in \mathcal{I} into sets $\mathfrak{A}_1, \mathfrak{A}_2, \dots$ given by

$$\mathfrak{A}_k = \{ \mathcal{C} : 2^{k-1} \leq |\mathcal{C}| \leq 2^k \},$$

we will argue that, for a suitable absolute constant $C_0 > 0$, for each $k = 1, \dots, \lceil \log_2 L_1 \rceil$ we have

$$\mu_n^\mp \left(\sum_{\mathcal{C} \in \mathfrak{A}_k} |\mathcal{C}| \geq (\varepsilon_{\beta,k} n)^2 \right) \leq \exp \left(-(\beta - C) \frac{(\varepsilon_{\beta,k} n)^2}{2^{k/2+1}} \right) \quad \text{for } \varepsilon_{\beta,k} := \frac{C_0}{e^{\beta/2k}}. \quad (4.4)$$

This, as we will later see, will readily imply (4.3) by a union bound over k .

To establish (4.4), denote the ceilings of \mathfrak{A}_k by $\mathcal{C}_1, \dots, \mathcal{C}_N$ for some N , and suppose that the event under consideration holds, i.e., $\sum_{i=1}^N |\mathcal{C}_i| \geq (\varepsilon_{\beta,k} n)^2$. Further let

$$\chi := \frac{\sum_{i=1}^N |\mathcal{C}_i|}{(\varepsilon_{\beta,k} n)^2} \quad (\geq 1).$$

It then follows, by the definition of \mathfrak{A}_k , that

$$N \leq \chi (\varepsilon_{\beta,k} n)^2 2^{1-k}.$$

In addition, for each i we have $|\partial \dot{\mathcal{C}}_i| \geq 4|\dot{\mathcal{C}}_i|^{1/2} \geq 4|\mathcal{C}_i|^{1/2}$ by the isoperimetric inequality in \mathbb{Z}^2 , whence again the definition of \mathfrak{A}_k implies that

$$\sum_{i=1}^N |\partial \dot{\mathcal{C}}_i| \geq 4 \sum_{i=1}^N |\mathcal{C}_i|^{1/2} \geq \chi (\varepsilon_{\beta,k} n)^2 2^{2-k/2}.$$

For any prescribed sequence of faces x_1, \dots, x_N such that $x_i \in \rho(\mathcal{C}_i)$ for each i , we can appeal to Corollary 2.26 to find that for any $r \geq N$,

$$\mu_n^\mp \left(\mathfrak{m} \left(\bigcup_{i \leq N} \mathfrak{W}_{x_i} \right) \geq r \right) \leq \exp(-(\beta - C)r).$$

Since $\mathfrak{m}(\bigcup_{i=1}^N \mathfrak{W}_{x_i}) \geq \frac{1}{2} \sum_{i=1}^N |\partial \dot{\mathcal{C}}_i|$ (each face of $\partial \dot{\mathcal{C}}_i$ corresponds to a distinct face in its supporting wall W , which may correspond in this way to a face in at most one other $\partial \dot{\mathcal{C}}_j$), we can take $r = \chi (\varepsilon_{\beta,k} n)^2 2^{1-k/2}$ (which is at least $N 2^{k/2} \geq N$ by the above bound on N) and conclude that

$$\mu_n^\mp \left(\sum_{\mathcal{C} \in \mathfrak{A}_k} |\mathcal{C}| \geq (\varepsilon_{\beta,k} n)^2 \right) \leq \sum_{\chi} \sum_{N \leq \chi (\varepsilon_{\beta,k} n)^2 2^{1-k}} \binom{n^2}{N} \exp \left(-(\beta - C) \chi (\varepsilon_{\beta,k} n)^2 2^{1-k/2} \right). \quad (4.5)$$

Using the bound $\sum_{j \leq pm} \binom{m}{j} \leq \exp(\mathbf{H}(p)m)$ where $\mathbf{H}(p) = p \log \frac{1}{p} + (1-p) \log \frac{1}{1-p}$ is the binary entropy function (with natural base)—valid for every m and $p \leq \frac{1}{2}$ (here $p = \varepsilon_{\beta,k}^2 2^{1-k} < \frac{1}{2}$ if β_0 is large enough)—yields

$$\mu_n^\mp \left(\sum_{\mathcal{C} \in \mathfrak{A}_k} |\mathcal{C}| \geq (\varepsilon_{\beta,k} n)^2 \right) \leq \sum_{\chi} \exp \left(\left(\mathbf{H}(\chi \varepsilon_{\beta,k}^2 2^{1-k}) - 2(\beta - C) \chi \varepsilon_{\beta,k}^2 2^{-k/2} \right) n^2 \right).$$

Hence, in order to establish (4.4) it suffices to show that if the absolute constant $C_0 > 0$ is large enough then

$$\mathbf{H}(\chi \varepsilon_{\beta,k}^2 2^{1-k}) \leq \frac{3}{2} (\beta - C) \chi \varepsilon_{\beta,k}^2 2^{-k/2} \quad \text{for every } k \geq 1, \quad (4.6)$$

noting the sum over at most n^2 values of $\chi \in [1, \varepsilon_{\beta,k}^{-2}]$ adds a $2 \log n$ term to the exponent, which is readily absorbed in the constant C from (4.4) since $(\varepsilon_{\beta,k} n)^2 2^{-(k/2+1)}$ is of order at least $n \log^2 n$ for all $k \leq \lceil \log_2 L_1 \rceil$. Indeed, using that $\mathbf{H}(p) \leq p(\log \frac{1}{p} + 1)$ for any $0 < p < 1$, it suffices to show that, for every $k \geq 1$,

$$2 \log(1/\varepsilon_{\beta,k}) + \log(1/\chi) + (\log 2)(k-1) + 1 \leq \frac{3}{4} (\beta - C) 2^{k/2}. \quad (4.7)$$

By the definition of $\varepsilon_{\beta,k}$ and the fact that $\chi \geq 1$, it suffices to show, for every $k \geq 1$,

$$\beta + (\log 2)(k-1) + 2 \log k + 1 - 2 \log C_0 \leq \frac{3}{4} (\beta - C) 2^{k/2}. \quad (4.8)$$

For every $k \geq 1$ we have $\frac{3}{4} 2^{k/2} \beta \geq 3 \cdot 2^{-3/2} \beta \geq \frac{21}{20} \beta$; thus, in order to establish (4.6) it suffices to show that

$$(\log 2)(k-1) + 2 \log k + 1 - 2 \log C_0 \leq \frac{1}{20} (\beta - C) 2^{k/2},$$

and indeed this easily holds for every large enough k , whereas selecting a large enough C_0 allows us to assume that k is sufficiently large. (We note in passing that the β term in the left-hand of (4.8) is the crux of our final threshold $C e^{-\beta n^2}$ in (4.3), which corresponds to $(\varepsilon_{\beta,1} n)^2$. A term of 2β on the left-hand of (4.8), for instance, would violate this inequality for $k=1$ and a large enough β .)

Having established (4.4), we observe that $\varepsilon_{\beta,k}^2 2^{-k/2}$ is decreasing in k , whence a union bound over k shows that $\sum_{\mathcal{C} \in \mathfrak{A}_k} |\mathcal{C}| \leq (\varepsilon_{\beta,k} n)^2$ holds for all $k = 1, \dots, \lceil \log_2 L_1 \rceil$ except with probability at most

$$\lceil \log_2 L_1 \rceil \exp \left(-(\beta - C) C_0^2 e^{-\beta} \lceil \log_2 L_1 \rceil^{-2} \frac{1}{2\sqrt{2}} L_1^{-1/2} n^2 \right) \leq \exp(-e^{-\beta} n \log^2 n),$$

where the $O(\log L_1)$ prefactor as well as the absolute constants in the exponent were offset by the $\beta - C$ term. On this event, the fraction of n^2 in ceilings with $2^{A-1} \leq |\mathcal{C}| \leq L_1$ is at most $\sum_{k \geq A} \varepsilon_{\beta,k}^2 \leq C_0^2 e^{-\beta} \sum_{k \geq A} k^{-2}$, thereby giving (4.3).

Letting $L_2 := e^{-2\beta} n^2 / \log^2 n$, we will treat ceilings with $L_1 \leq |\mathcal{C}| \leq L_2$ by establishing that

$$\mu_n^\mp \left(\left| \bigcup \{ \rho(\mathcal{C}) : \mathcal{C} \text{ is a ceiling with } L_1 \leq |\mathcal{C}| \leq L_2 \} \right| \geq C e^{-\beta} n^2 \right) \leq \exp(-c_0 \beta n \log n) \quad (4.9)$$

for some absolute constants $C, c_0 > 0$. This will follow from showing that, for all $k = \lceil \log_2 L_1 \rceil, \dots, \lceil \log_2 L_2 \rceil$,

$$\mu_n^\mp \left(\sum_{\mathcal{C} \in \mathfrak{A}_k} |\mathcal{C}| \geq (\bar{\varepsilon}_{\beta,k} n)^2 \right) \leq \exp \left(-(\beta - C) \frac{(\bar{\varepsilon}_{\beta,k} n)^2}{2^{k/2+1}} \right) \quad \text{for } \bar{\varepsilon}_{\beta,k} := \frac{e^{-\beta/2}}{\lceil \log_2 L_2 \rceil - k + 8}. \quad (4.10)$$

As argued above (4.10) will follow once we show the analog of (4.7) w.r.t. the modified quantity $\bar{\varepsilon}_{\beta,k}$, which here translates into showing that for every $\lceil \log_2 L_1 \rceil \leq k \leq \lceil \log_2 L_2 \rceil$,

$$\beta + 2 \log(\lceil \log_2 L_2 \rceil - k + 8) + (\log 2)(k - 1) + 1 \leq \frac{3}{4}(\beta - C)2^{k/2}.$$

This indeed holds as the left-hand is $O(k) = O(\log n)$ while the right-hand has order $2^{k/2} \geq \sqrt{L_1} = n^{1-o(1)}$ for each such k , and (4.10) is obtained. To identify the dominant term in the union bound over k , note that

$$\frac{\bar{\varepsilon}_{\beta,k}^2 2^{-k/2}}{\bar{\varepsilon}_{\beta,k+1}^2 2^{-(k+1)/2}} = \sqrt{2} \left(1 - \frac{1}{\lceil \log_2 L_2 \rceil - k + 8} \right)^2 \geq \sqrt{2} \cdot \left(\frac{7}{8} \right)^2 > 1,$$

whence $\sum_{\mathcal{C} \in \mathfrak{A}_k} |\mathcal{C}| \leq (\bar{\varepsilon}_{\beta,k} n)^2$ holds for all $k = \lceil \log_2 L_1 \rceil, \dots, \lceil \log_2 L_2 \rceil$ except with probability

$$\lceil \log_2 L_2 \rceil \exp \left(-(\beta - C) \frac{1}{64} e^{-\beta} \frac{1}{2\sqrt{2}} L_2^{-1/2} n^2 \right) \leq \exp(-c_0 \beta n \log n),$$

using the definition of L_2 . This establishes (4.9), thereby completing the proof. \blacksquare

Lemma 4.8. *There exist $\beta_0, C > 0$ so that for every fixed $\beta > \beta_0$,*

$$\mu_n^\mp \left(\bigcup \{ \rho(\dot{\mathcal{C}}) : \mathcal{C} \text{ is a ceiling with } |\dot{\mathcal{C}}| \geq n \text{ and } \dim_{\text{ip}}(\dot{\mathcal{C}}) \geq 4 \} \geq n^{5/3} \right) \leq \exp(-(\beta - C)n^{7/6}), \quad (4.11)$$

and the same bound holds under $\hat{\mu}_n^h$.

Proof. Let $\mathcal{C}_1, \dots, \mathcal{C}_N$ be the outermost ceilings satisfying $|\dot{\mathcal{C}}| \geq n$ and $\dim_{\text{ip}}(\dot{\mathcal{C}}) \geq 4$. Clearly, $N \leq n$ by the lower bound on the area of each of these pairwise disjoint ceilings. Furthermore, by the assumption on the isoperimetric dimension, $|\partial \dot{\mathcal{C}}_i| \geq |\dot{\mathcal{C}}_i|^{3/4}$ for every i ; thus, whenever $\sum_{i=1}^N |\dot{\mathcal{C}}_i| \geq n^{5/3}$ we can conclude that

$$\sum_i |\partial \dot{\mathcal{C}}_i| \geq n^{-1/2} \sum_i |\dot{\mathcal{C}}_i| \geq n^{7/6}.$$

By the same argument the led to (4.5), it now follows that the probability in the left-hand of (4.11) is at most

$$\sum_{N \leq n} \binom{n^2}{N} \exp \left(-(\beta - C)n^{7/6} \right) \leq \exp \left(-(\beta - C')n^{7/6} \right),$$

where we absorbed $\sum_{N \leq n} \binom{n^2}{N} \leq \exp((1 + o(1))n \log n)$ into the larger constant C' , as required. \blacksquare

5. BOUNDING THE INTERFACE HISTOGRAM BELOW A GIVEN HEIGHT

Our objective in this section is to obtain the following bound on the height histogram of the interface in the presence of a floor at height $-h$, capturing the essence of entropic repulsion.

Theorem 5.1. *Let $L_{n,k}$ be the set of $x \in \mathcal{L}_{0,n}$ such that $\mathcal{I} \cap (\{x\} \times \mathbb{R})$ intersects $\mathcal{L}_{<h_n^* - h - k}$ or is not a singleton. There exist $\beta_0, C > 0$ so that for all $\beta > \beta_0$ and every $h < h_n^* - 1$ and $k \geq 1$,*

$$\hat{\mu}_n^h (|L_{n,k}| \geq C e^{-\beta} n^2) \leq \exp \left(-e^{(\beta - C)k} n \right).$$

This will be a consequence of the following estimate on the total area in macroscopic and near macroscopic ceilings below height $h_n^* - h - 1$.

Theorem 5.2. *There exist $\beta_0, C > 0$ so that for every fixed $\beta > \beta_0$, if $h < h_n^* - 1$ and*

$$\mathfrak{C}_j(\mathcal{I}) = \{\mathcal{C} \text{ is a ceiling in } \mathcal{I} \text{ with } \text{ht}(\mathcal{C}) = j\}$$

then for every $1 \leq k \leq h_n^ - 1$,*

$$\widehat{\mu}_n^h \left(\sum \{|\mathcal{C}| : \mathcal{C} \in \mathfrak{C}_{h_n^* - h - k - 1}(\mathcal{I}), |\mathcal{C}| \geq n^{1.9}\} \geq e^{-\beta k n^2} \right) \leq e^{-(e^{(\beta-C)^k} \wedge 3 \log n)n}.$$

Remark 5.3. The threshold of $e^{-\beta k n^2}$ on $\sum |\mathcal{C}|$ in Theorem 5.2 could be extended via the same proof into $e^{-a\beta k n^2}$ for any fixed $1 < a < 2$ (with a bottleneck at $a = 2$ due to our reliance on $e^{-a\beta h_n^*} = O(n^{3/2})$, whereas $e^{-2\beta h_n^*}$ is only guaranteed to be $O(n^{3/2+\varepsilon_\beta})$ by the best upper bound we have for h_n^*).

Observe that the proof of Theorem 5.1 follows directly from Theorem 5.2 along with our results from §4: By Lemma 4.7, the total area in $\rho(\mathcal{C})$ for ceilings \mathcal{C} with $|\mathcal{C}| \leq n^2 / \log^3 n$ is at most $Ce^{-\beta n^2}$ except with probability $\exp(-5n \log n)$; thus, for fixed $\bar{k} \geq 1$, the size of $\bigcup \{\rho(\mathcal{C}) : \mathcal{C} \in \mathfrak{C}_{<h_n^* - h - \bar{k}}\}$ is at most $e^{-\beta \bar{k} n^2}$ except with probability $\exp(-e^{(\beta-C)\bar{k}n})$ via Theorem 5.2 and a union bound over $k = \bar{k}, \dots, h_n^* - 1$; and there are at most $e^{-2\beta} n^2$ wall faces in \mathcal{I} except with probability $\exp(-c_\beta n^2)$ by Lemma 4.5.

The general principle behind the proof is the standard competition between energy and entropy which propels the interface to height $h_n^* - 1$. It will be illuminating to consider the case $h = 0$ and show, via a straightforward Peierls argument which served as the basis of the sharp results for the SOS model in [9, 10], that the interface in $\widehat{\mu}_n^0$ is propelled to height $(1 - \varepsilon_\beta)h_n^*$. We do so in the next claim, and then explain why one cannot hope for such an argument, albeit effective for SOS, to be applicable for the Ising model.

Claim 5.4. *There exist $\beta_0, C_0 > 0$ such that, if $h_0 = \lfloor (4\beta + C_0)^{-1} \log n \rfloor$, then for every $\beta > \beta_0$ and $k \geq 2$,*

$$\widehat{\mu}_n^0 \left(\sum \{|\mathcal{C}| : \mathcal{C} \in \mathfrak{C}_{h_0 - k}(\mathcal{I})\} \geq e^{-2\beta k n^2} \right) \leq \exp(-e^{\beta k} n).$$

Proof. Let $k \geq 1$. For $\mathcal{I} \in \mathbf{I}_0^{\text{fl}}$, let

$$A_0(\mathcal{I}) = \bigcup \{\rho(\mathcal{C}) : \mathcal{C} \in \mathfrak{C}_{h_0 - k - 1}(\mathcal{I})\}, \quad A_1(\mathcal{I}) = \bigcup \{\rho(\mathcal{C}) : \mathcal{C} \in \mathfrak{C}_{h_0 - k}(\mathcal{I})\},$$

and let

$$\mathbf{I} = \{\mathcal{I} : |A_0(\mathcal{I})| \geq e^{-2\beta k n^2}\}.$$

Further let $\mathbf{I}' = \Phi_1^\uparrow(\mathbf{I})$, where we recall from Definition 4.1 which effectively shifts the interface up by 1. In particular, $|A_1(\mathcal{I}')| \geq e^{-2\beta k n^2}$ for every $\mathcal{I}' \in \mathbf{I}'$.

For $\mathcal{I}' \in \mathbf{I}'$ and a subset $A \subset A_1(\mathcal{I}')$, let $\mathbf{J}_{\mathcal{I}'}$ be the interfaces obtained from \mathcal{I}' by setting all spins of $\sigma(\mathcal{I}')$ in $A \times (\frac{1}{2} + \mathbb{Z}_+)$ to be minus (modifying $h_0 - k$ spins above each $x \in A$). As $\mathcal{I}' \cap ((\mathcal{L}_{0,n} \setminus \partial \mathcal{L}_{0,n}) \times \mathbb{R}) \subset \mathcal{L}_{\geq 1}$ for any $\mathcal{I}' \in \mathbf{I}_0^{\text{fl}}$, by Lemma 4.3, we can recover \mathcal{I}' from $\mathcal{J} \in \mathbf{J}_{\mathcal{I}'}$ (reading A off from $\mathcal{J} \cap \mathcal{L}_0$). This implies that the interface sets $\{\mathbf{J}_{\mathcal{I}'} : \mathcal{I}' \in \mathbf{I}'\}$ are pairwise disjoint.

Moreover, if $|A| = m$ then it is easy to verify that for some absolute constant $C > 1$,

$$\widehat{\mu}_n^0(\mathcal{J}) \geq \exp(-(\beta + C)4(h_0 - k)m) \widehat{\mu}_n^0(\mathcal{I}'),$$

since any such \mathcal{J} has $\mathbf{m}(\mathcal{J}; \mathcal{I}') \leq 4(h_0 - k)m$ by construction. Taking $C_0 = 4C$ in the definition of h_0 , this translates into $\widehat{\mu}_n^0(\mathcal{J}) \geq p^m \widehat{\mu}_n^0(\mathcal{I}')$ where $p = p_{k,n} = e^{(4\beta + C_0)k} / n$ (noting that $p_{k,n} \leq p_{h_0,n} \leq 1$), and so,

$$\begin{aligned} 1 &\geq \sum_{\mathcal{I}' \in \mathbf{I}'} \sum_{\mathcal{J} \in \mathbf{J}_{\mathcal{I}'}} \widehat{\mu}_n^0(\mathcal{J}) \geq \sum_{\mathcal{I}' \in \mathbf{I}'} \widehat{\mu}_n^0(\mathcal{I}') \sum_{A \subset A_1(\mathcal{I}')} p^{|A|} = \sum_{\mathcal{I}' \in \mathbf{I}'} \widehat{\mu}_n^0(\mathcal{I}') (1 + p)^{|A_1(\mathcal{I}')|} \\ &\geq \exp\left(\frac{1}{2} p e^{-2\beta k n^2}\right) \widehat{\mu}_n^0(\mathbf{I}') = \exp\left(\frac{1}{2} e^{(2\beta + C_0)k} n\right) \widehat{\mu}_n^0(\mathbf{I}'), \end{aligned}$$

using $1 + p \geq \exp(p/2)$ for $p \leq 1$ in the inequality between the lines. By Proposition 4.2 we have $\widehat{\mu}_n^0(\mathbf{I}) \leq \widehat{\mu}_n^0(\mathbf{I}') \exp(4(\beta + C')n)$ for another $C' > 0$, completing the proof provided β_0 is large enough. \blacksquare

The proof of Claim 5.4 captured the effect of entropic repulsion: if an interface \mathcal{I} is such that $A_0(\mathcal{I})$, the projections of the ceilings at height $j = h_0 - k$, is too large, one can compare it to the set of (many) interfaces \mathcal{J} obtained by elevating all heights via a single $n \times n$ slab of (+) spins (a cost of $\exp(-(4\beta - C)n)$ in energy) while replacing any subset of A_0 by spikes reaching down to height 0 (a gain of $\exp(e^{-(4\beta - C)j} |A_0|)$ in entropy); comparing the two competing terms, the entropy term wins when $j < h_0$, giving the claim. Unfortunately, the correct rate α of large deviations in the 3D Ising model exceeding height j is realized not

by a deterministic shape such as a spike of height j but rather by a *distribution* over pillars \mathcal{P}_x that behave as decorated random walks (e.g., typically $\text{diam}(\rho(\mathcal{P}_x))$ has order \sqrt{j} for such \mathcal{P}_x). For one to be able to modify \mathcal{I} to have such downward pillars \mathcal{P}_x at a subset $A \subset A_0$, they would need to all fit in the corresponding ceilings and with one another; then, one would need to account for their randomness and interactions among themselves and with existing walls. Even more problematic is the fact that comparing the interface \mathcal{I} with a given pillar \mathcal{P}_x of height j to an interface \mathcal{I} without it via cluster expansion (Theorem 2.16) incurs a cost of $\exp(\bar{K}j)$ due the interaction term \mathbf{g} . This would result in an error of $n^{\varepsilon\beta}$ —thus not improving on Claim 5.4.

One could instead begin by examining the interfaces \mathcal{I}' obtained from \mathcal{I} by elevating all height by 1 yet without planting any spikes or pillars. As mentioned above, $\hat{\mu}_n^0(\mathcal{I}')$ is at least $\exp(-4(\beta - C)n)\hat{\mu}_n^0(\mathcal{I})$, due to the additional $n \times n$ slab of (+) spins. The latter interfaces have their minima within $A_0 \times \mathbb{R}$ not reaching to height 0, which one would expect to have probability at most $\exp(-e^{-\alpha h}|A_0|)$ via the estimate (2.9) in Proposition 2.28 (proved via the approximate Domain Markov property for ceilings in [23]). If this probability outweighs the $\exp(-4(\beta - C)n)$ factor, we will obtain the desired lower bound.

However, the event $\mathcal{G}_{S_{n,h}^{\circ}}$ in (2.9) complicates matters, and is a real (rather than a technical) obstacle: Mesoscopic walls nested in a ceiling may encapsulate additional ceilings and modify their heights in a complicated way (whereby a single wall supports multiple ceilings in different heights), making some of them more favorable, by the exact same entropic repulsion mechanism that propels the interface to height $h_n^* - 1$. Thus, we cannot preclude the existence of such walls, and must resort to an analysis of their subtle effect.

Our approach for handling the delicate effect of mesoscopic walls involves conditioning on them—possibly revealing large ceilings at different heights in the process—then reevaluating the new landscape of ceilings, and applying the aforementioned entropic repulsion argument for a specific height in this conditional space.

Consider an integer $1 \leq k < h_n^*$. With $\kappa_0 > 0$ the absolute constant from Proposition 2.28, define

$$\mathbf{W}^\dagger(\mathcal{I}) = \left\{ W : W \text{ is a wall of } \mathcal{I} \text{ with } \text{diam}(\rho(W)) \geq (e^{\kappa_0(h_n^* - k)} \vee \log n) \right\} \quad (5.1)$$

(note the dependence on k) and associate to every \mathcal{I} the interface \mathcal{I}^\dagger consisting only of these walls, i.e.,

$$\mathcal{I}^\dagger := \mathcal{I}_{\mathbf{W}^\dagger(\mathcal{I})}. \quad (5.2)$$

Remark 5.5. If $W' \ni W$ for some $W \in \mathbf{W}^\dagger$ then $W' \in \mathbf{W}^\dagger$ since in that case $\text{diam}(\rho(W')) > \text{diam}(\rho(W))$. (This would not be the case were our inclusion criterion for \mathbf{W}^\dagger instead been phrased e.g. in terms of $\mathbf{m}(W)$.) Consequently, every face of a ceiling $\mathcal{C} \in \mathcal{I}$ supported by some $W \in \mathbf{W}^\dagger(\mathcal{I})$ is also a ceiling face in \mathcal{I}^\dagger at the exact same height (whereas \mathcal{C} may be contained in a strictly larger ceiling $\mathcal{C}^\dagger \in \mathcal{I}^\dagger$ with $\text{ht}(\mathcal{C}^\dagger) = \text{ht}(\mathcal{C})$).

With this in mind, we will later use that if the set of ceilings $\mathfrak{C}_{h_n^* - h - k - 1}(\mathcal{I})$ with $|\dot{\mathcal{C}}| \geq n^{1.9}$ is $\{\mathcal{C}_1, \dots, \mathcal{C}_N\}$, then $\mathfrak{C}_{h_n^* - h - k - 1}(\mathcal{I}^\dagger)$ includes ceilings $\{\mathcal{C}_1^\dagger, \dots, \mathcal{C}_{N'}^\dagger\}$ for some $N' \leq N$ where every \mathcal{C}_j must have $\mathcal{C}_j \subset \mathcal{C}_{i_j}^\dagger$ for some i_j . (N.B. that $N' \leq N$ as a single ceiling in \mathcal{I}^\dagger may include multiple \mathcal{C}_j 's, as well as extra horizontal wall faces from \mathcal{I} which were deleted in the transition to \mathcal{I}^\dagger , but distinct ceilings in \mathcal{I}^\dagger must correspond to disjoint such sets of ceilings in \mathcal{I} by definition of a ceiling as a $*$ -connected component of ceiling faces.)

The key reduction in the proof of Theorem 5.2 is the next bound on the total area of a certain class of ceilings in $\mathfrak{C}_{h_n^* - h - 1}(\Phi_k^\dagger(\mathcal{I}^\dagger))$, which we will thereafter use to bound the total area of ceilings in $\mathfrak{C}_{h_n^* - h - k - 1}(\mathcal{I})$.

Lemma 5.6. *There exist absolute constants $\beta_0, C > 0$ so that, for all $\beta > \beta_0$ and $0 \leq h < h_n^* - 1$, the following holds. Let $1 \leq k < h_n^*$, define \mathcal{I}^\dagger as in (5.2), and let*

$$\mathbf{C}^\dagger(\mathcal{I}) = \left\{ \mathcal{C}^\dagger \in \mathfrak{C}_{h_n^* - h - 1}(\mathcal{I}^\dagger) : |\mathcal{C}^\dagger| \geq n^{1.9}, \bar{M}_{\mathcal{C}^\dagger}^\downarrow(\mathcal{I}) < h_n^* - k \right\}.$$

Then for every sufficiently large enough n we have

$$\hat{\mu}_n^h \left(\sum_{\mathcal{C}^\dagger \in \mathbf{C}^\dagger(\mathcal{I})} |\mathcal{C}^\dagger| \geq e^{-\beta k} n^2, \mathbf{m}(\mathbf{W}^\dagger(\mathcal{I})) \leq n \log n \right) \leq e^{-(e^{(\beta - C)k} \wedge n^{1/5})n}.$$

Furthermore, with $\gamma := ne^{-\alpha h_n^*}$ this remains true when replacing $e^{(\beta - C)k}$ in the right-hand by $\gamma e^{(3\beta - C)k}$, and moreover, it holds conditionally on the event $\{\mathbf{W}^\dagger(\mathcal{I}) = \mathbf{W}_n\}$ for any \mathbf{W}_n such that $\mathbf{m}(\mathbf{W}_n) \leq n \log n$.

(We emphasize that the definition of \mathbf{C}^\dagger looked at $\bar{M}_{\mathcal{C}^\dagger}$ in \mathcal{I} as opposed to in \mathcal{I}^\dagger .)

We postpone the proof of this lemma in order to show how to derive Theorem 5.2 from it.

5.1. **Proof of Theorem 5.2 modulo Lemma 5.6.** Consider some $\mathcal{I} \in \mathbf{I}_h^\dagger$. Since every $\mathbf{W}^\dagger(\mathcal{I})$ consists of distinct walls, each of which has size at least $\log n$, if $\sum_{W \in \mathbf{W}^\dagger(\mathcal{I})} \mathfrak{m}(W) \geq \frac{1}{2}n \log n$, then any subset of $|\mathbf{W}^\dagger(\mathcal{I})| \wedge n$ many of those walls also has total \mathfrak{m} at least $\frac{1}{2}n \log n$. Thus, by Lemma 4.6 with $N = |\mathbf{W}^\dagger(\mathcal{I})| \wedge n$,

$$\widehat{\mu}_n^h \left(\sum \{ \mathfrak{m}(W) : W \in \mathbf{W}^\dagger(\mathcal{I}) \} \geq \frac{1}{2}n \log n \right) \leq e^{-3n \log n}.$$

It will therefore suffice to show that

$$\widehat{\mu}_n^h(\mathbf{I}_k) \leq \exp(-e^{(\beta-C)kn}) \quad (5.3)$$

for

$$\mathbf{I}_k = \left\{ \mathcal{I} \in \mathbf{I}_h^\dagger : \sum_{\mathcal{C} \in \mathfrak{C}_{h_n^* - h - k - 1}(\mathcal{I})} |\mathcal{C}| \mathbf{1}_{\{|\mathcal{C}| \geq n^{1.9}\}} \geq e^{-\beta k} n^2, \mathfrak{m}(\mathbf{W}^\dagger(\mathcal{I})) \leq \frac{1}{2}n \log n \right\}.$$

(N.B. we are aiming to bound $\sum |\mathcal{C}|$, throughout this theorem, as opposed to $\sum |\tilde{\mathcal{C}}|$; indeed, the latter can include an array of ceilings at different heights, which would be treated by sets \mathbf{I}_k 's as above for different k 's.) For each $\mathcal{I} \in \mathbf{I}_k$, define

$$\tilde{\mathcal{I}} = \Phi_k^\dagger(\mathcal{I})$$

where Φ_k^\dagger is the map given in Definition 4.1, which we recall translates \mathcal{I} by $(0, 0, k)$, adds a subset of the $4kn$ wall faces $\mathcal{F}(\partial \mathcal{L}_{0,n} \times \llbracket 0, k \rrbracket)$, and then deletes some other subset of wall faces. By the definition of $\mathcal{I} \in \mathbf{I}_k$,

$$\sum_{W \in \mathbf{W}^\dagger(\tilde{\mathcal{I}})} \mathfrak{m}(W) \leq \sum_{W \in \mathbf{W}^\dagger(\mathcal{I})} \mathfrak{m}(W) + 4kn \leq \frac{1}{2}n \log n + 4kn < n \log n, \quad (5.4)$$

using here that $4kn \leq (1 + \varepsilon_\beta)\beta^{-1}n \log n$.

Denote by \mathbf{C} the set of every ceiling $\mathcal{C} \in \mathfrak{C}_{h_n^* - h - k - 1}(\mathcal{I})$ that satisfies $|\mathcal{C}| \geq n^{1.9}$, and further let $\tilde{\mathbf{C}}$ be the analogous subset of the ceilings $\mathfrak{C}_{h_n^* - h - 1}(\tilde{\mathcal{I}})$. As no ceiling face was deleted from $\theta_1 \mathcal{I}$ in forming $\tilde{\mathcal{I}}$, we see that every $\mathcal{C} \in \mathbf{C}$ must satisfy $\rho(\mathcal{C}) \subset \rho(\tilde{\mathcal{C}})$ for some $\tilde{\mathcal{C}} \in \mathfrak{C}_{h_n^* - h - 1}(\tilde{\mathcal{I}})$, and so $|\tilde{\mathcal{C}}| \geq n^{1.9}$ as well, thus $\tilde{\mathcal{C}} \in \tilde{\mathbf{C}}$. In particular, if $\mathcal{I} \in \mathbf{I}_k$ then $\tilde{\mathcal{I}}$ satisfies

$$\sum_{\tilde{\mathcal{C}} \in \tilde{\mathbf{C}}} |\tilde{\mathcal{C}}| \geq e^{-\beta k} n^2.$$

Now, if $\tilde{\mathcal{C}} \in \tilde{\mathbf{C}}$ then its supporting wall W must satisfy $\text{diam}(\rho(W)) \geq \text{diam}(\tilde{\mathcal{C}}) \geq \frac{1}{2}|\tilde{\mathcal{C}}|^{1/2} \geq \frac{1}{2}n^{0.95}$, implying that $W \in \mathbf{W}^\dagger(\tilde{\mathcal{I}})$ for every sufficiently large n (as we have $\exp(\kappa_0(h_n^* - k)) \leq n^{\varepsilon_\beta}$ whereas $\varepsilon_\beta < \frac{1}{2}$, say, provided β_0 is suitably large). In particular (invoking Remark 5.5), every face in $\tilde{\mathcal{C}} \in \tilde{\mathbf{C}}$ must be part of some ceiling $\mathcal{C}^\dagger \in \tilde{\mathcal{I}}^\dagger$ at the same height of $h_n^* - h - 1$ (which, once again, satisfies $|\mathcal{C}^\dagger| \geq |\tilde{\mathcal{C}}| \geq n^{1.9}$). Moreover, by Lemma 4.3 we know that $\tilde{\mathcal{I}} \cap ((\mathcal{L}_{0,n} \setminus \partial \mathcal{L}_{0,n}) \times \mathbb{R}) \subset \mathcal{L}_{\geq k-h}$; thus, for every ceiling $\mathcal{C}^\dagger \in \mathfrak{C}_{h_n^* - h - 1}(\tilde{\mathcal{I}}^\dagger)$,

$$M_{\mathcal{C}^\dagger}^\dagger(\tilde{\mathcal{I}}) \leq (h_n^* - h - 1) - (k - h) = h_n^* - k - 1,$$

which, when combined, imply that $\mathcal{C}^\dagger \in \mathbf{C}^\dagger(\tilde{\mathcal{I}})$ as defined in Lemma 5.6. In conclusion,

$$\sum_{\mathcal{C}^\dagger \in \mathbf{C}^\dagger(\tilde{\mathcal{I}})} |\mathcal{C}^\dagger| \geq e^{-\beta k} n^2 \quad \text{and} \quad \mathfrak{m}(\mathbf{W}^\dagger(\tilde{\mathcal{I}})) < n \log n$$

holds for all $\tilde{\mathcal{I}} \in \Phi_k^\dagger(\mathbf{I}_k)$. We may thus bound $\widehat{\mu}_n^h(\Phi_k^\dagger(\mathbf{I}_k))$ via Lemma 5.6, and combine it with Proposition 4.2 to get that

$$\widehat{\mu}_n^h(\mathbf{I}_k) \leq e^{(4\beta+C)kn} \widehat{\mu}_n^h(\Phi_k^\dagger(\mathbf{I}_k)) \leq \exp \left[\left((4\beta + C)k - (\gamma e^{(3\beta-C)k} \wedge n^{1/5}) \right) n \right].$$

Since $k = O(\log n)$, the case where $\gamma e^{(3\beta-C)k} \geq n^{1/5}$ immediately leads to $\widehat{\mu}_n^h(\mathbf{I}_k) \leq \exp[-(1 - o(1))n^{6/5}]$, thus it remains to treat the converse case. Recalling that γ satisfies

$$\exp(-2\beta - e^{-4\beta}) \leq \gamma < \exp(2\beta),$$

we see that $\gamma e^{(3\beta-C)k} \geq e^{(\beta-C)k}$ (with room to spare: we could have taken $C' \exp[\beta + (3\beta - C)(k - 1)]$), so

$$\widehat{\mu}_n^h(\mathbf{I}_k) \leq \exp \left[(4\beta + C)kn - e^{(\beta-C)k} n \right] \leq \exp \left[-e^{(\beta-C')k} n \right],$$

where replacing the constant $C > 0$ by some larger absolute constant $C' > 0$ in the last transition allowed us to absorb the term $(4\beta + C)kn$. We have thus established (5.3), as required. \blacksquare

5.2. Proof of Lemma 5.6. Let $\mathcal{I} \sim \hat{\mu}_n^h$, and reveal the wall collection $\mathbf{W}^\dagger(\mathcal{I})$, which by our assumption satisfies $\mathbf{m}(\mathbf{W}^\dagger(\mathcal{I})) \leq n \log n$. We would like to obtain an upper bound on the probability that there are $\varepsilon_\beta n^2$ faces in ceilings of $\mathbf{C}^\dagger(\mathcal{I})$ (for an appropriate ε_β). As our large deviation estimates in Proposition 2.28 are only applicable to simply-connected regions, extra care must be taken to modify the ceilings of \mathcal{I}^\dagger into a simply-connected subset. To do so, we will expose a minimal additional collection of walls, indexed by \mathcal{P}^\dagger below.

Towards that purpose, define the random variable \mathcal{P}^\dagger , measurable given $\mathbf{W}^\dagger(\mathcal{I})$, to be

$$\mathcal{P}^\dagger = \mathcal{P}^\dagger(\mathcal{I}) := \text{minimum face subset of } \mathcal{L}_{0,n} \text{ such that } \mathcal{P}^\dagger(\mathcal{I}) \cup \rho(\mathbf{W}^\dagger(\mathcal{I})) \cup \partial\mathcal{L}_{0,n} \text{ is } * \text{-connected}$$

and the collection \mathcal{S}_h^\dagger (also measurable given $\mathbf{W}^\dagger(\mathcal{I})$, as the latter determines \mathcal{I}^\dagger) to be

$$\mathcal{S}_h^\dagger = \mathcal{S}_h^\dagger(\mathcal{I}) := \{S^\dagger = \rho(\mathcal{C}^\dagger) \setminus \mathcal{P}^\dagger : \mathcal{C}^\dagger \in \mathfrak{C}_{h_n^* - h - 1}(\mathcal{I}^\dagger), |\mathcal{C}^\dagger| \geq n^{1.9}\}.$$

(For $\mathcal{P}^\dagger(\mathcal{I})$ to be well-defined, if there exist multiple face subsets of $\mathcal{L}_{0,n}$, all with the same cardinality, that turn $\rho(\mathbf{W}^\dagger(\mathcal{I})) \cup \partial\mathcal{L}_{0,n}$ to be $*$ -connected, choose the identity of $\mathcal{P}^\dagger(\mathcal{I})$ according to a predefined arbitrary lexicographic ordering of all subsets of $\mathcal{L}_{0,n}$.) We will show that

$$\hat{\mu}_n^h \left(\sum_{S^\dagger \in \mathcal{S}_h^\dagger} |S^\dagger| \mathbf{1}_{\{\bar{M}_{S^\dagger}^\dagger(\mathcal{I}) < h_n^* - k\}} \geq e^{-\beta k} n^2 \mid \mathbf{W}^\dagger(\mathcal{I}) \right) \mathbf{1}_{\{\mathbf{m}(\mathbf{W}^\dagger(\mathcal{I})) \leq n \log n\}} \leq e^{-(\gamma e^{(3\beta - C)k} \wedge n^{1/5})n}, \quad (5.5)$$

where here we used $\bar{M}_{S^\dagger}^\dagger$ to denote the minimum height of \mathcal{I} within S^\dagger after centering it by $\text{ht}(\mathcal{C}^\dagger) = h_n^* - h - 1$.

It is important to stress that for every ceiling \mathcal{C}^\dagger in \mathcal{I}^\dagger , the only obstacles preventing it from being simply-connected would be walls from $\mathbf{W}^\dagger(\mathcal{I})$ nested in it; hence, by the minimality of \mathcal{P}^\dagger ,

$$\rho(\mathcal{C}^\dagger) \setminus \mathcal{P}^\dagger \text{ is simply-connected for every ceiling } \mathcal{C}^\dagger \text{ in } \mathcal{I}^\dagger.$$

(If \mathcal{C}^\dagger is supported by $W \in \mathbf{W}^\dagger(\mathcal{I})$ and $P = \mathcal{P}^\dagger \cap \mathcal{C}^\dagger$ connects $x, y \in W$ and some $W' \in \mathbf{W}^\dagger(\mathcal{I})$ nested in \mathcal{C}^\dagger , one can delete the part of P connecting x to W , retaining connectivity of W, W' via the other part of P .)

Recall that each \mathcal{C}^\dagger has $|\mathcal{C}^\dagger| \geq n^{1.9}$ and $|\partial\mathcal{C}^\dagger| \leq n \log n$ (the latter due to our assumption on $\mathbf{W}^\dagger(\mathcal{I})$). To address the effect of subtracting the set \mathcal{P}^\dagger from the ceilings \mathcal{C}^\dagger of \mathcal{I}^\dagger , we will appeal to the following classical result concerning Euclidean Traveling Salesman Problems (TSPs). (For the d -dimensional analog of this result, see, e.g., the references in [37, §2], as well as Lemma 1 there.)

Fact 5.7. *[[41]; see also, e.g., [37, §2]] For every m points x_1, \dots, x_m in $[0, 1]^2$ there exists a path connecting them of length at most $\sqrt{2.8m} + 2$.*

For completeness, we include the short classical argument giving the bound $2\lceil\sqrt{m}\rceil$: one divides $[0, 1]^2$ into strips of height $1/\sqrt{m}$ each, and proceeds from the bottom strip to the top via switchbacks, connecting the points within each strip from left-to-right, then right-to-left, etc. If we were to use only horizontal and vertical lines (as an upper bound), and there are k_i points in strip i , then the total length associated with strip i would be at most $1 + k_i/\sqrt{m}$, whence the overall total length is at most $\lceil\sqrt{m}\rceil + \sum k_i/\sqrt{m} \leq 2\lceil\sqrt{m}\rceil$.

By our assumption that $\mathbf{m}(\mathbf{W}^\dagger(\mathcal{I})) \leq n \log n$, we can apply Fact 5.7 to the faces of $\rho(\mathbf{W}^\dagger(\mathcal{I})) \cup \partial\mathcal{L}_{0,n}$ with $m = O(n \log n)$, rescaling $[0, 1]^2$ to $\mathcal{L}_{0,n}$. This yields that there exists a face subset $\mathcal{P}^\dagger \subset \mathcal{L}_{0,n}$ connecting $\rho(\mathbf{W}^\dagger(\mathcal{I})) \cup \partial\mathcal{L}_{0,n}$ of size

$$|\mathcal{P}^\dagger| = O(n\sqrt{m}) = O(n^{3/2}\sqrt{\log n}) < n^{1.6}$$

for large enough n . In the first equality, the extra n came from rescaling of $[0, 1]^2$ to $\mathcal{L}_{0,n}$. Also the discretization effect of faces amounts to at most $2m = O(n \log n)$ additional faces at the endpoints of each of the paths.

With this bound in hand, we may proceed to show (5.5). Write $\mathbf{C}^\dagger(\mathcal{I}) = \{\mathcal{C}_1^\dagger, \dots, \mathcal{C}_N^\dagger\}$ for some N , as well as $S_i^\dagger = \rho(\mathcal{C}_i^\dagger) \setminus \mathcal{P}^\dagger$, and note that we may assume that

$$\sum_{i=1}^N |S_i^\dagger| \geq e^{-\beta k} n^2, \quad (5.6)$$

or else there is nothing left to show. For each $z \in \mathcal{P}^\dagger$, let U_z be $\rho(\mathcal{C}^\dagger)$ for the ceiling \mathcal{C}^\dagger in \mathcal{I}^\dagger with $z \in \rho(\mathcal{C}_i^\dagger)$, and reveal the walls in \mathfrak{G}_{z, U_z} (where we recall $\mathfrak{G}_{x, S}$ is the sequence of nested walls $\mathfrak{W}_{x, S}$ which nest x in S ,

as well as the every wall nested in this sequence of walls), denoting the set of all walls revealed so far by \mathbf{W}_0 . For every $i = 1, \dots, N$, let

$$\tilde{S}_i = S_i^\dagger \setminus \bigcup_{z \in \mathcal{P}^\dagger} \rho(\mathfrak{G}_{z, U_z}),$$

which is simply-connected since $\tilde{S}_i^\dagger = S_i^\dagger$ (and \tilde{S}_i was obtained by deleting from S_i^\dagger , at the very most, a collection of simply-connected regions each adjacent to its boundary). Observe that every W revealed as part of \mathfrak{G}_{z, U_z} for $z \in \mathcal{P}^\dagger$ necessarily has $\text{diam}(\rho(W)) \leq \exp(\kappa_0(h_n^* - 1 - k)) \vee \log n \leq \exp(\kappa_0 h_n^*) \vee \log n \leq n^{\varepsilon_\beta}$ (otherwise it would have been part of $\mathbf{W}^\dagger(\mathcal{I})$). Thus, when obtaining \tilde{S}_i from S_i^\dagger , we subtracted at most $|\mathcal{P}^\dagger| n^{\varepsilon_\beta} \leq n^{1.6+\varepsilon_\beta}$ from its area, and similarly added at most $n^{1.6+\varepsilon_\beta}$ to its perimeter. Recalling that S_i^\dagger itself was obtained by deleting \mathcal{P}^\dagger from \mathcal{C}_i^\dagger which had $|\mathcal{C}_i^\dagger| \geq n^{1.9}$ and $|\partial \mathcal{C}_i^\dagger| \leq n \log n$, it follows that

$$|\tilde{S}_i| \geq |\mathcal{C}_i^\dagger| - |\mathcal{P}^\dagger| n^{\varepsilon_\beta} \geq n^{1.9} - n^{1.6+\varepsilon_\beta} \geq \frac{1}{2} n^{1.9},$$

and

$$|\partial \tilde{S}_i| \leq |\partial \mathcal{C}_i^\dagger| + n^{1.6+\varepsilon_\beta} \leq (1 + o(1)) n^{1.6+\varepsilon_\beta}.$$

Combining these inequalities, while recalling that \tilde{S}_i is simply-connected, shows that for large enough n ,

$$\dim_{\text{ip}}(\tilde{S}_i) \leq 7$$

provided that β_0 is large enough (whence ε_β becomes small enough). Finally, note that

$$\sum_{i=1}^N |S_i^\dagger \setminus \tilde{S}_i| \leq |\mathcal{P}^\dagger| (e^{\kappa_0 h_n^*} \vee \log n) \leq n^{1.6+\varepsilon_\beta} < \frac{1}{2} e^{-\beta k} n^2, \quad (5.7)$$

where the last inequality is by the fact that $e^{-\beta k} n^2 \geq n^{7/4-\varepsilon_\beta}$ for large n , since $h_n^* \leq \frac{1}{4\beta-C} \log n + C'$ and so $e^{-\beta k} \geq e^{-\beta h_n^*} \geq n^{-1/4-\varepsilon_\beta}$. Combining this with the (deterministic) fact that

$$\sum_{i=1}^N |S_i^\dagger| \mathbf{1}_{\{\bar{M}_{S_i^\dagger}^\downarrow < h_n^* - k\}} \leq \sum_{i=1}^N \left(|\tilde{S}_i| \mathbf{1}_{\{\bar{M}_{\tilde{S}_i}^\downarrow < h_n^* - k\}} + |S_i^\dagger \setminus \tilde{S}_i| \right),$$

it therefore suffices to show that

$$\hat{\mu}_n^h(\mathfrak{E} \mid \mathbf{W}_0) \leq e^{-(\gamma e^{(\beta-C)k} \wedge n^{1/5})n} \quad \text{where} \quad \mathfrak{E} = \left\{ \sum_{i=1}^N |\tilde{S}_i| \mathbf{1}_{\{\bar{M}_{\tilde{S}_i}^\downarrow(\mathcal{I}) < h_n^* - k\}} \geq \frac{1}{2} e^{-\beta k} n^2 \right\}.$$

To this end, we now wish to successively reveal the walls in \tilde{S}_i . Let \mathcal{F}_i be the filtration corresponding to initially revealing \mathbf{W}_0 , then proceeding to reveal $\mathbf{W}_i = \{W_z : z \in \tilde{S}_i\}$ at step $i = 1, \dots, N$. Recalling the definition of the event \mathcal{G} playing a role in Proposition 2.28, given by

$$\mathcal{G}_A(r) = \bigcap_{z \in A} \{\text{diam}(\rho(W_z)) \leq r \vee \log n\},$$

we stress that every $z \in \mathcal{L}_{0,n}$ whose W_z was not revealed as part of \mathbf{W}_0 is conditioned to satisfy $\mathcal{G}_z(e^{\kappa_0 \bar{h}})$ for

$$\bar{h} = h_n^* - k.$$

By taking a supremum at every step i over the walls outside \tilde{S}_i , it therefore remains to bound the conditional probability that the event \mathfrak{E} occurs given on \mathbf{W}_0 , via

$$\hat{\mu}_n^h(\mathfrak{E} \mid \mathbf{W}_0) \leq \prod_{i=1}^N \hat{\mu}_n^h(\bar{M}_{\tilde{S}_i}^\downarrow < \bar{h} \mid \mathcal{G}_{\mathcal{L}_{0,n} \setminus \{\rho(W) : W \in \mathbf{W}^\dagger\}}(e^{\kappa_0 \bar{h}}), \mathcal{F}_{i-1}) \leq \prod_{i=1}^N \Psi_i \quad (5.8)$$

for

$$\Psi_i := \sup_{\substack{\mathbf{W} = \{W_z : z \notin \tilde{S}_i\} \\ \rho(\{\dot{W} : W \in \mathbf{W}\}) \subset \tilde{S}_i}} \hat{\mu}_n^h(\bar{M}_{\tilde{S}_i}^\downarrow < \bar{h} \mid \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \bar{h}}), \mathbf{I}_{\mathbf{W}}).$$

Recall that if $S_i^\dagger = \rho(\mathcal{C}_i^\dagger) \setminus \mathcal{P}^\dagger$ for some ceiling $\mathcal{C}_i^\dagger \in \mathfrak{C}_{h_n^* - h - 1}(\mathcal{I}^\dagger)$ then by definition $\text{ht}(\mathcal{C}_i^\dagger) = h_n^* - h - 1$. Therefore, when looking at Ψ_i , the implicit conditioning in $\hat{\mu}_n^h$ that $\mathcal{I} \in \mathbf{I}_n^{\text{fl}}$ amounts to having $\bar{M}_{\tilde{S}_i}^\downarrow < h_n^*$ via μ_n^\mp . With this in mind, for every $\mathbf{W} = \{W_z : z \notin \tilde{S}_i\}$ with $\rho(\{\dot{W} : W \in \mathbf{W}\}) \subset \tilde{S}_i$ we have

$$\begin{aligned} \hat{\mu}_n^h \left(\bar{M}_{\tilde{S}_i}^\downarrow < \hbar \mid \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}), \mathbf{I}_W \right) &= \mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < \hbar \mid \bar{M}_{\tilde{S}_i}^\downarrow < h_n^*, \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}), \mathbf{I}_W \right) \\ &= \frac{\mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < \hbar, \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}) \mid \mathbf{I}_W \right)}{\mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < h_n^*, \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}) \mid \mathbf{I}_W \right)}. \end{aligned} \quad (5.9)$$

First consider $1 \leq k \leq h_n^* - 1 - \sqrt{2 \log n}$, so that $\sqrt{2 \log n} \leq \hbar \leq h_n^* - 2$, whence the criterion for $\mathcal{G}_z(e^{\kappa_0 \hbar})$ is dominated by the $\exp(\kappa_0 \hbar)$ term (exceeding, e.g., $\log^2 n$ for every large n). As \tilde{S}_i is simply-connected, has $\dim_{\text{ip}}(\tilde{S}_i) \leq 7$ and $\hbar \geq \sqrt{2 \log n} \geq (\log |\tilde{S}_i|)^{1/2}$, we may use (2.9) from Proposition 2.28 for $h = \hbar$, yielding

$$\mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < \hbar, \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}) \mid \mathbf{I}_W \right) \leq \exp(- (1 - \varepsilon_\beta) |\tilde{S}_i| e^{-\alpha \hbar}),$$

where we used that the event $\mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar})$ implies $\mathcal{G}_{\tilde{S}_i^\circ}(e^{\kappa_0 \hbar})$ for the subset $\tilde{S}_i^\circ \subset \tilde{S}_i$ specified in that proposition. At the same time, invoking (2.8) from the same proposition, this time for $h = h_n^*$, we get

$$\mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < h_n^*, \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}) \mid \mathbf{I}_W \right) \geq \exp(- (1 + \varepsilon_\beta) |\tilde{S}_i| e^{-\alpha h_n^*}),$$

here using that $\mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar})$ is implied by $\mathcal{G}_{\tilde{S}_i}^m(5h_n^*)$ (and hence also by $\mathcal{G}_{\tilde{S}_i}^m(4h_n^*)$) since $e^{\kappa_0 \hbar} > \log^2 n > 5h_n^*$ for every sufficiently large n . Substituting the last two displays in (5.9), we find that

$$\Psi_i \leq \exp \left(- \left((1 - \varepsilon_\beta) e^{-\alpha \hbar} - (1 + \varepsilon_\beta) e^{-\alpha h_n^*} \right) |\tilde{S}_i| \right).$$

We know from [21, Corollary 5.2] that

$$\alpha_{\hbar} + \alpha_k - \varepsilon_\beta \leq \alpha_{\hbar+k} \quad \text{and} \quad \alpha_k \geq (4\beta - C)k,$$

which, since $\hbar + k = h_n^*$ and we defined $\gamma := n \exp(-\alpha h_n^*)$, implies that

$$\frac{\gamma}{n} = e^{-\alpha h_n^*} \leq e^{-\alpha_k + \varepsilon_\beta} e^{-\alpha \hbar} \leq e^{-(4\beta - C)k} e^{-\alpha \hbar}.$$

Hence, we may absorb the term $(1 + \varepsilon_\beta) e^{-\alpha h_n^*} = (1 + \varepsilon_\beta) \gamma / n$ from the above upper bound on Ψ_i into the constant $C > 0$ from the lower bound $e^{-\alpha \hbar} \geq e^{(4\beta - C)k} \gamma / n$ (with $k \geq 1$ and β_0 large) and obtain that

$$\Psi_i \leq \exp \left(-\gamma e^{(4\beta - C)k} |\tilde{S}_i| / n \right); \quad (5.10)$$

thus, by plugging in the fact that

$$\sum_i |\tilde{S}_i| \geq \frac{1}{2} e^{-\beta k} n^2, \quad (5.11)$$

obtained from (5.6) and (5.7), we infer that

$$\prod_{i=1}^N \Psi_i \leq \exp \left(-\gamma e^{(4\beta - C)k} \sum_{i=1}^N \frac{|\tilde{S}_i|}{n} \right) \leq \exp \left(-\frac{1}{2} \gamma e^{(3\beta - C)k} n \right). \quad (5.12)$$

Next, consider $\frac{1}{3} h_n^* \leq k \leq h_n^* - 1$. Treating this regime is significantly easier. (The overlap between the regimes is indicative: the latter can be used to yield a lower bound of $(1 - \varepsilon_\beta) h_n^*$ on the typical height of \mathcal{I} .) With the same starting point (5.9), we invoke (2.7) from Proposition 2.27 to deduce that

$$\mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < \hbar, \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}) \mid \mathbf{I}_W \right) \leq \mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < \hbar \mid \mathbf{I}_W \right) \leq \exp(-|\tilde{S}_i| e^{-(4\beta - C)\hbar}),$$

whereas (2.6) from Proposition 2.27 implies that

$$\mu_n^\mp \left(\bar{M}_{\tilde{S}_i}^\downarrow < h_n^*, \mathcal{G}_{\tilde{S}_i}(e^{\kappa_0 \hbar}) \mid \mathbf{I}_W \right) \geq \mu_n^\mp \left(\mathcal{G}_{\tilde{S}_i}(4h_n^*) \mid \mathbf{I}_W \right) \geq \exp(-|\tilde{S}_i| e^{-(4\beta + C)h_n^*}),$$

since $(e^{\kappa_0 \hbar} \vee \log n) \geq (4h_n^* \vee \log n) = \log n$. Combining these estimates, and using that $\hbar = h_n^* - k$, yields

$$\Psi_i \leq \exp \left(-|\tilde{S}_i| e^{(4\beta - C)k - 2Ch_n^*} \right) \leq \exp \left(-e^{(4\beta - C)k} |\tilde{S}_i| n^{-1 - \varepsilon_\beta} \right).$$

Together with the bound (5.11) and the fact that $k \geq \frac{1}{3}h_n^*$, this shows that

$$\prod_{i=1}^N \Psi_i \leq \exp\left(-\frac{1}{4}e^{(3\beta-C)k}n^{1-\varepsilon_\beta}\right) \leq \exp\left(-cn^{5/4-\varepsilon_\beta}\right). \quad (5.13)$$

Combining (5.12) and (5.13) with (5.8) concludes the proof. \blacksquare

Remark 5.8. One can extend Lemma 5.6—followed by Theorem 5.2, and in turn, Theorem 5.1—to the case $h = h_n^* - 1$ if either $k \geq 2$ or the quantity $\gamma = n \exp(-\alpha_{h_n^*})$ —which we recall satisfies $e^{-2\beta-\varepsilon_\beta} \leq \gamma < e^{2\beta}$ —is in a given range, for a suitably modified ε_β^k replacing $e^{-\beta k}$. Consider for instance the regime where

$$\gamma > \beta^2 \quad \text{or} \quad k \geq 2.$$

One then defines $\mathbf{C}^\dagger(\mathcal{I})$ from Lemma 5.6 to be

$$\mathbf{C}^\dagger(\mathcal{I}) = \left\{ \mathcal{C}^\dagger \in \mathfrak{C}_{h_n^*-h}(\mathcal{I}^\dagger) : |\mathcal{C}^\dagger| \geq n^{1.9}, \bar{M}_{\mathcal{C}^\dagger}^\downarrow(\mathcal{I}) \leq h_n^* - k \right\},$$

and argues that, in place of its conclusion, one has

$$\hat{\mu}_n^h \left(\sum_{\mathcal{C}^\dagger \in \mathbf{C}^\dagger(\mathcal{I})} |\mathcal{C}^\dagger| \geq c_0 \beta^{-k} n^2, \mathfrak{m}(\mathbf{W}^\dagger(\mathcal{I})) \leq n \log n \right) \leq \exp\left(- (5\gamma \beta^{-k} e^{(4\beta-C)(k-1)} \wedge n^{1/5}) n\right).$$

Indeed, this follows by showing that, in lieu of (5.5), one has

$$\hat{\mu}_n^h \left(\sum_{S^\dagger \in \mathcal{S}_h^\dagger} |S^\dagger| \mathbf{1}_{\{\bar{M}_{S^\dagger}^\downarrow(\mathcal{I}) \leq h_n^* - k\}} \geq e^{-\beta k} n^2 \mid \mathbf{W}^\dagger(\mathcal{I}) \right) \mathbf{1}_{\{\mathfrak{m}(\mathbf{W}^\dagger(\mathcal{I})) \leq n \log n\}} \leq e^{-(5\gamma \beta^{-k} e^{(4\beta-C)(k-1)} \wedge n^{1/5}) n}. \quad (5.14)$$

Following the same argument used to prove Lemma 5.6 yet with $h = h_n^* + 1 - k$, instead of the bound (5.10) one arrives at $\Psi_i \leq \exp(-\gamma C e^{(4\beta-C)(k-1)} |\tilde{S}_i|/n)$. Taking $c_0 := 5C$ and plugging in $\sum_{i=1}^N |\tilde{S}_i|/n \geq c_0 \beta^{-k} n$ implies (5.14). In the application of this lemma towards proving Theorem 5.2, the exponent on the right of (5.14) competes with $(4\beta + C)kn$. For $k \geq 2$, the former is the dominant term regardless of γ , whereas for $k = 1$ it is the dominant term if $5\gamma/\beta > 4\beta + C$, as in our assumption.

5.3. Proof of Proposition 1.6. The lower bound in the proposition is precisely the statement of Lemma 4.4, and it remains to show how to derive the upper bound via Theorem 5.2. If \mathbf{C} is the set of ceilings in \mathcal{I} with $|\mathcal{C}| \geq n^{1.9}$ and height at least $h_n^* - 1$, and \mathfrak{E} is the event that $\sum_{\mathcal{C} \in \mathbf{C}} |\mathcal{C}| \geq A := (1 - 2e^{-\beta})n^2$, then Theorem 5.2 together with Lemma 4.7 imply that $\hat{\mu}_n^0(\mathfrak{E}) = 1 - o(1)$, implying in particular that

$$\mu_n^\mp(\mathbf{I}_0^{\text{fl}}) = (1 - o(1))\mu_n^\mp(\mathfrak{E}).$$

We have at most $n^{0.1}$ such ceilings, and if $\mathfrak{W}_1, \dots, \mathfrak{W}_N$ ($N \leq n^{0.1}$) are the nested sequences of walls supporting them, then

$$\mathfrak{m}\left(\bigcup_{i=1}^N \mathfrak{W}_i\right) \geq 4\sqrt{A}(h_n^* - 1) \geq \frac{1 - e^{-\beta/2}}{\beta + 1/4} n \log n = (1 - \varepsilon_\beta)\beta^{-1} n \log n$$

since $h_n^* \geq (4\beta + e^{-\beta})^{-1} \log n - C$ for large enough n , and an area of A ceiling faces must be supported by a perimeter of at least $4\sqrt{A}$ vertical wall faces at heights $1, \dots, h_n^* - 1$. Using $\sum_{N \leq n^{0.1}} \binom{n^2}{N} \leq \exp(n^{0.1+o(1)})$, we deduce from Theorem 2.26 that

$$\mu_n^\mp(\mathfrak{E}) \leq \exp\left(-(\beta - C)(1 - \varepsilon_\beta)\beta^{-1} n \log n\right) \sum_{N \leq n^{0.1}} \binom{n^2}{N} \leq \exp\left(-(1 - \varepsilon'_\beta)n \log n\right),$$

and the proof is complete. \blacksquare

6. BOUNDING THE INTERFACE HISTOGRAM ABOVE A GIVEN HEIGHT

Our goal in this section is to establish the following bound on the fraction of the sites in $\mathcal{L}_{0,n}$ above/below which the interface is not a singleton at height 0 whenever $\mathbf{h} \geq h_n^*$.

Theorem 6.1. *Let U_n be the set of $x \in \mathcal{L}_{0,n}$ such that $\mathcal{I} \cap (\{x\} \times \mathbb{R}) \neq \{0\}$. There exist $\beta_0, C > 0$ so that for all $\beta > \beta_0$ and every $\mathbf{h} \geq h_n^*$, we have*

$$\hat{\mu}_n^{\mathbf{h}}(|U_n| \geq Ce^{-\beta}n^2) \leq \exp(-e^{-\beta}n).$$

This will be a consequence of the following estimate on encountering walls with a near linear excess energy.

Theorem 6.2. *There exist $\beta_0, C > 0$ so the following holds for all $\beta > \beta_0$. For every $\mathbf{h} \geq h_n^*$ and $r \geq n^{9/10}$,*

$$\hat{\mu}_n^{\mathbf{h}}(\exists \text{ a wall } W \text{ in } \mathcal{I} \text{ with } |\rho(\dot{W})| \geq r^2) \leq e^{-(\beta-C)(r \wedge n \log n)}.$$

Furthermore, for any $k \geq 1$,

$$\hat{\mu}_n^{\mathbf{h}}(\exists \text{ walls } \{W_i\}_{i=1}^k \text{ in } \mathcal{I} \text{ with disjoint hulls and } \min_i |\rho(\dot{W}_i)| \geq r^2) \leq e^{-(\beta-C)(kr \wedge n \log n)}.$$

Theorem 6.1 readily follows from Theorem 6.2 and our results from Section 4; to see this, argue as follows: Lemma 4.7 guarantees that the total area in $\rho(\dot{\mathcal{C}})$ for ceilings \mathcal{C} with $|\dot{\mathcal{C}}| \leq n^2/\log^3 n$ is at most $Ce^{-\beta}n^2$ except with probability $\exp(-5n \log n)$. To treat the remaining ceilings having $|\dot{\mathcal{C}}| > n^2/\log^3 n$, define

$$\mathbf{W}_\ell(\mathcal{I}) = \{\text{outermost walls } W \text{ in } \mathcal{I} \text{ with } e^{-\ell} \leq |\rho(\dot{W})|/n^2 \leq e^{1-\ell}\} \quad (\ell \in \llbracket 2\beta, 3 \log \log n \rrbracket).$$

An application of Theorem 6.2 with $k = e^{\ell/2}$ and $r = ne^{-\ell/2}$ shows that $|\mathbf{W}_\ell(\mathcal{I})| \leq e^{\ell/2}$ (reflecting a total area of at most $n^2e^{1-\ell/2}$) except with probability $\exp(-(\beta-C)n)$. A union bound then shows that, except with probability $\exp(-(\beta-C)n)$, we have $\bigcup_\ell \bigcup_{W \in \mathbf{W}_\ell} |\rho(\dot{W})| \leq Cn^2e^{-\beta}$. The probability that a single outermost W exists with $|\rho(\dot{W})| \geq e^{-2\beta}n^2$ (and hence $\mathbf{m}(W) \geq e^{-\beta}n$) is, by another application of Theorem 6.2, at most $\exp(-(\beta-C)e^{-\beta}n)$, the dominant term in our final estimate. Finally, by Lemma 4.5, the total number of wall faces is at most $e^{-2\beta}n^2$ except with probability $\exp(-c_\beta n^2)$.

It will be illuminating to describe how the vanilla SOS approach can yield an estimate on the total area in ceilings at height at least 1 when $\mathbf{h} \geq (1 + \varepsilon_\beta)h_n^*$ once we combine it with the approximate Domain Markov property for ceilings established in [23]. More generally, we have the following claim, applicable to any $\mathbf{h} \geq 0$ (the above mentioned estimate corresponds to the special case $\mathbf{h} = h_1$, which in turn is $(1 + \varepsilon_\beta)h_n^*$; at the other end, the case $\mathbf{h} = 0$ gives a bound on $\sum\{|\mathcal{C}| : \mathcal{C} \in \mathfrak{C}_{>h_1}(\mathcal{I})\}$ for $\mathcal{I} \in \hat{\mu}_n^0$, a counterpart for Claim 5.4).

Claim 6.3. *There exist $\beta_0, C_0 > 0$ such that for every $h_1 \geq (1 + \frac{1}{\sqrt{\beta}})\frac{1}{4\beta - C_0} \log n$ and all $\beta > \beta_0$ and $\mathbf{h} \geq 0$,*

$$\hat{\mu}_n^{\mathbf{h}}\left(\sum\{|\mathcal{C}| : \mathcal{C} \in \mathfrak{C}_{>h_1-\mathbf{h}}(\mathcal{I})\} \geq (1/\sqrt{\beta})n^2\right) \leq O(\exp(-3n \log n)).$$

Proof. Suppose $\mathbf{h} \leq h_1$ (otherwise there is nothing to prove), and let $\beta_0, C_0 > 0$ be given by Proposition 2.27. Further set $h_0 = \lfloor \frac{1}{4\beta - C_0} \log n \rfloor$. For every interface \mathcal{I} , let $\{\mathcal{C}_1, \dots, \mathcal{C}_N\}$ denote the set of outermost ceilings in $\mathfrak{C}_{>h_1-\mathbf{h}}(\mathcal{I})$ with $|\dot{\mathcal{C}}| \geq n^{1.9}$ (note that $N \leq n^{0.1}$). Lemma 4.7 allows us to ignore ceilings with smaller areas, whereas if x_i is any face in the wall supporting \mathcal{C}_i then Lemma 4.6 allows us to preclude any \mathcal{I} in which $\mathbf{m}(\bigcup_i \mathfrak{W}_{x_i}) > (6/\beta)n \log n$. With this in mind, let \mathbf{I} be the set of interfaces where $\sum_i |\mathcal{C}_i| \geq (1/\sqrt{\beta})n^2$ and in addition $\mathbf{m}(\bigcup_i \mathfrak{W}_{x_i}) \leq 30nh_0$.

Reveal the outermost walls of $\mathcal{I} \in \mathbf{I}$, and continue the process inductively in supported ceilings, while not revealing any walls nested in the \mathcal{C}_i 's. Applying Proposition 2.27 iteratively on $S_i = \rho(\dot{\mathcal{C}}_i)$ for each \mathcal{C}_i , we see that for $S = \bigcup_i S_i$ and any realization \mathbf{W} of the walls outside of S (subject to having $\mathcal{I} \in \mathbf{I}_n^{\text{fl}}$) we have

$$\hat{\mu}_n^{\mathbf{h}}(\bar{M}_S^\downarrow \leq h_0 \mid \mathbf{I}\mathbf{W}) \geq \mu_n^{\mp}(\bar{M}_S^\downarrow \leq h_0 \mid \mathbf{I}\mathbf{W}) \geq \exp(-|S|(e^{-(4\beta-C_0)(h_0+1)}) \geq e^{-|S|/n} \geq e^{-n},$$

using that, given the wall set \mathbf{W} , the event $\mathcal{I} \in \mathbf{I}_n^{\text{fl}}$ is implied by $\bar{M}_S^\downarrow \leq h_1$ (which is implied by $\bar{M}_S^\downarrow \leq h_0$) in the first inequality. So, if \mathbf{I}' is the set of $\mathcal{I} \in \mathbf{I}$ such that $\bar{M}_S^\downarrow \leq h_0$ for $S = \bigcup_i \rho(\dot{\mathcal{C}}_i)$, then $\hat{\mu}_n^{\mathbf{h}}(\mathbf{I}') \geq \exp(-n)\hat{\mu}_n^{\mathbf{h}}(\mathbf{I})$.

Now consider the map Φ on \mathbf{I}' which, for every \mathcal{C}_i —recalling x_i is some face in the wall supporting it—deletes the entire nested sequence \mathfrak{W}_{x_i} and replaces it by cylinders of the form $\rho(\partial\mathcal{C}) \times [h^-, (\text{ht}(\mathcal{C}) \wedge (h_0 - \mathbf{h}))]$ for every ceiling \mathcal{C} that was supported by one of the walls $W \in \mathfrak{W}_{x_i}$, with h^- the height of its nesting ceiling. (Notice that $\Phi(\mathcal{I}) \in \mathbf{I}_n^{\text{fl}}$ holds for all $\mathcal{I} \in \mathbf{I}'$.)

The ceilings $\{\mathcal{C}_i\}$ account for at least $(\beta^{-1/2}n^2)^{1/2}$ vertical wall faces in each of the slabs h_0+1, \dots, h_1 in \mathcal{I} ; thus, in our cluster expansion (namely, Theorem 2.16), we have $\beta \mathbf{m}(\mathcal{I}; \Phi(\mathcal{I})) \geq \beta^{3/4}n(h_1 - h_0) \geq \beta^{1/4}nh_0$, vs. the contribution of the interaction terms \mathbf{g} in that theorem, which is at most $\bar{K}|\bigcup_i \mathfrak{W}_{x_i}| \leq 30\bar{K}nh_0$. Altogether, if β_0 is large enough then $\hat{\mu}_n^0(\mathbf{I}) \leq \exp(n)\hat{\mu}_n^0(\mathbf{I}') \leq \exp(-(1 - \varepsilon_\beta)\beta^{1/4}nh_0)$, as required. ■

Remark 6.4. We now explain why one cannot hope for the proof of Claim 6.3 to hold all the way to $h_1 = (1 + o(1))h_n^*$, while pointing out a subtle but important difference from the Peierls argument used in the SOS model. Whereas in the SOS model (being a distribution over height functions) one can always lower a level line by 1 via decreasing the heights in its interior, the analog of a level line in the Ising model—a ceiling \mathcal{C}_0 supported by a wall W_0 —might not be (in fact often will not be) consistent with such an operation. Indeed, even if we suppose \mathcal{C}_0 is the only ceiling supported by W_0 (to simplify matters), it may be (and in fact often will be the case) that the ceiling \mathcal{C}_0 is a part of a larger connected set $F \subset \mathcal{I}$ of horizontal faces at height $\text{ht}(\mathcal{C}_0)$, which are categorized as wall (rather than ceiling) faces due to the shape of the interface far (say, at distance $\varepsilon \log n$) below them. It is the set F that one would want to trim so as to gain $|\partial\mathcal{C}|$ in energy (since, locally around $\text{ht}(\mathcal{C}_0)$, that set is the analogue of an SOS level line). However, such an operation might shift the thermal fluctuations of the interface above $\rho(F \setminus \mathcal{C}_0)$ and have them clash with other wall faces. Further complicating matters is the fact that the wall W_0 might be tilted, whence its horizontal fluctuations (on account of which faces in $F \setminus \mathcal{C}_0$ are not ceiling faces) are not well understood. To bypass these issues, the proof of Claim 6.3 “straightened” the walls supporting the ceilings under consideration into cylinders; e.g., in the case described above with a single \mathcal{C}_0 , the energetic cost of modifying W_0 in $k \sim h_n^*$ slabs would have order at least $k|\partial\mathcal{C}_0|$. To offset this cost and rule out a ceiling \mathcal{C}_0 with $|\partial\mathcal{C}_0| \asymp n$, one would need to gain at least $(1/\beta)kn$ deleted walls from the trimming operation; with $\partial\mathcal{C}_0$ contributing $O(n)$ such wall faces in every slab being shrunk, this argument would need a leeway of $h_1 \geq (1 + c/\beta)k \geq (1 + \varepsilon_\beta)h_n^*$ slabs.

Whereas these obstacles are highly nontrivial, in the setting of Theorem 6.2 there is a single wall W_0 whose entire wall cluster may be successfully deleted in the regime $h \geq h_n^*$ —with the caveat that it will require the extension given in §3 of the bounds on the maximum within a ceiling obtained in [23].

6.1. Proof of Theorem 6.2. Let us first consider the situation of a single wall W_0 ; the case of multiple walls $W_0^{(1)}, \dots, W_0^{(\ell)}$ will be obtained by iteratively applying this argument.

We will show the following stronger statement, applicable to any $h \geq 0$. Recalling that the supporting ceiling of a wall W_0 is the ceiling whose faces are adjacent to \dot{W}_0 (which for outermost walls necessarily belongs to \mathfrak{C}_0), and letting

$$\mathfrak{I}_r := \left\{ \mathcal{I} : \exists \text{ a wall } W_0 \text{ in } \mathcal{I} \text{ supported by } \hat{\mathcal{C}}_0 \in \mathfrak{C}_{\geq h_n^* - h}(\mathcal{I}) \text{ such that } \mathbf{m}(W_0) \geq r \right\},$$

we will argue that

$$\hat{\mu}_n^h(\mathfrak{I}_r) \leq e^{-(\beta - C)(r \wedge n \log n)} \quad \text{for every } r \geq n^{9/10}.$$

In the case $h \geq h_n^*$, this will immediately imply the required result, as every W with $\mathbf{m}(\rho(\dot{W})) \geq r^2$ is contained in some outermost wall W_0 with $\mathbf{m}(W_0) \geq |\rho(\dot{W}_0)|^{1/2} \geq |\rho(\dot{W})|^{1/2} \geq r$, for which the above estimate holds.

Consider a standard wall W_0 , as well as a realization of its complete wall cluster $\mathbf{F}_0 = \text{Clust}(W_0)$, such that $\mathbf{m}(\mathbf{F}_0) = r$ for some $r \geq n^{9/10}$. Further let

$$\mathfrak{I}_{\mathbf{F}_0} := \left\{ \mathcal{I} : \begin{array}{l} W_0 \in \text{standard wall collection of } \mathcal{I}, \text{ its wall cluster is } \mathbf{F}_0 \\ \text{and it is supported by a ceiling } \hat{\mathcal{C}}_0 \in \mathfrak{C}_{\geq h_n^* - h}(\mathcal{I}) \end{array} \right\},$$

emphasizing that we did not restrict this set to $\mathcal{I} \in \mathbf{I}_h^{\text{fl}}$. It will suffice to show that

$$\hat{\mu}_n^h(\mathfrak{I}_{\mathbf{F}_0}) \leq e^{-(\beta - C)(r \wedge n \log n)}, \tag{6.1}$$

since there are $O(n^2) = e^{o(r)}$ locations for the placement of W_0 , and at most $e^{\bar{c}r}$ wall clusters \mathbf{F} with $\mathbf{m}(\mathbf{F}) \leq r$ for some $\bar{c} > 0$ (independent of β), whence a union bound over (6.1) will conclude the proof.

If $r \geq n \log n$ then this is readily implied by [23, Lemma 3.10], as $\mu_n^{\mp}(\mathfrak{I}_{\mathbf{F}_0^x}) \leq \exp(-(\beta - C)\mathbf{m}(\mathbf{F}_0^x))$ which is then at most $\exp(-(\beta - C)n \log n)$, hence extends to $\hat{\mu}_n^h$ via Lemma 4.4. Assume therefore that $r \leq n \log n$.

Further assume—this time, an assumption that will require extra justification—that W_0 satisfies

$$\sum \{ |\mathcal{C}| : \mathcal{C} \text{ is a ceiling of } \mathcal{I}_{W_0} \text{ with } \mathcal{C} \Subset W_0, \dim_{\text{ip}}(\mathcal{C}) > 4 \} \leq n^{7/4}. \tag{6.2}$$

(We will later show that, except with probability $O(\exp(-3n \log n))$, every wall $W_z \in \mathcal{I}$ satisfies this.)

Let $\widehat{\mathcal{C}}_0$ denote the ceiling that supports W_0 , and let

$$\mathbf{C}_0 = \{\mathcal{C} \text{ is a ceiling of } \mathcal{J} = \mathcal{I}_{\mathbf{F}_0} \text{ with } \mathcal{C} \Subset W_0\} \quad \text{and} \quad S_0 = \bigcup_{\mathcal{C} \in \mathbf{C}_0} \rho(\mathcal{C}).$$

(Note that \mathbf{C}_0 consists not of ceilings of \mathcal{I}_{W_0} but rather those of the interface \mathcal{J} comprising the wall cluster \mathbf{F}_0 . Whereas the former are simply-connected, a ceiling $\mathcal{C} \in \mathbf{C}_0$ might not be, due to nested walls in $\text{Clust}(W_0)$.) Further let

$$\widetilde{\mathcal{J}}_{\mathbf{F}_0} := \left\{ \mathcal{I} \in \mathcal{J}_{\mathbf{F}_0} : \bar{M}_{S_0}^\downarrow \leq h_n^* \right\}.$$

(Observe that $\mathcal{I} \in \widetilde{\mathcal{J}}_{\mathbf{F}_0}$ need not belong to \mathbf{I}_h^{fl} even if it does satisfy that $\mathcal{I} \cap (\rho(\dot{W}_0)^c \times \mathbb{R}) \subset \mathcal{L}_{\geq -h}$; for instance, the interface will intersect $\mathcal{L}_{< -h}$ when a ceiling $\mathcal{C} \in \mathbf{C}_0$ has $\text{ht}(\mathcal{C}) < h_n^* - h$ while $\bar{M}_{\rho(\mathcal{C})}^\downarrow = h_n^*$.) Let $\mathbf{V}_0 = \{W_z : z \notin \rho(\dot{W}_0)\}$ be an arbitrary realization of the set of walls of $\rho(\dot{W}_0)^c$ which is compatible with W_0 and such that $\mathcal{I} \cap (\rho(\dot{W}_0)^c \times \mathbb{R}) \subset \mathcal{L}_{\geq -h}$ and the ceiling $\widehat{\mathcal{C}}_0$ supporting W_0 has $\text{ht}(\widehat{\mathcal{C}}_0) \geq h_n^* - h$. The key to the proof of (6.1) (and in turn, the entire theorem) is establishing that for any such \mathbf{V}_0 ,

$$\widehat{\mu}_n^h(\widetilde{\mathcal{J}}_{\mathbf{F}_0} \mid \mathcal{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}) \geq \exp(-Cr) \quad \text{for an absolute constant } C > 0. \quad (6.3)$$

It is important to stress the effect of conditioning on $\mathcal{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}$. Conditioning on the $\mathbf{I}_{\mathbf{V}_0}$ —namely on the walls \mathbf{V}_0 external to W_0 —will not introduce any complication as far as the behavior of \mathcal{I} within $\mathcal{C} \in \mathbf{C}_0$ is concerned, thanks to our results [23] on the rigidity and approximate Domain Markov property for ceilings. However, the conditioning on $\mathcal{J}_{\mathbf{F}_0}$ does need to be handled with care, as it is equivalent to conditioning on

$$\mathbf{m}(W_z) \leq d_\rho(z, \partial W_0) \quad \text{for every } z \in S_0,$$

on top having the walls in \mathbf{F}_0 belong to the interface; in other words, in terms of the event \mathcal{E}_S^η from (3.1),

$$\widehat{\mu}_n^h(\cdot \mid \mathcal{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}) = \widehat{\mu}_n^h(\cdot \mid \mathcal{E}_{S_0}^\eta, \mathbf{I}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}) \quad \text{with } \eta_z = d_\rho(z, \partial W_0) \text{ for all } z \in S_0. \quad (6.4)$$

Before analyzing \mathcal{I} in this delicate conditional space towards (6.3), let us show how to infer (6.1) from it. Let Φ be the map on interfaces $\mathcal{I} \in \widetilde{\mathcal{J}}_{\mathbf{F}_0}$ that deletes all of \mathbf{F}_0 from the standard wall representation of \mathcal{I} . Clearly, Φ is a bijection (\mathbf{F}_0 being fixed), and the crux of the definition of $\widetilde{\mathcal{J}}_{\mathbf{F}_0}$ is that $\mathcal{I}' := \Phi(\mathcal{I}) \in \mathbf{I}_h^{\text{fl}}$ holds for all $\mathcal{I} \in \widetilde{\mathcal{J}}_{\mathbf{F}_0}$ (in fact even if we had \mathcal{I} intersect $S_0 \cap \mathcal{L}_{< -h}$). Indeed, \mathcal{I} does not intersect $\rho(\dot{W}_0)^c \times \mathcal{L}_{< -h}$, and having deleted \mathbf{F}_0 from it to reach \mathcal{I}' , clearly \mathcal{I}' does not intersect $(S_0)^c \times \mathcal{L}_{< -h}$, whereas in $S_0 \times \mathbb{R}$, since we have $\text{ht}(\widehat{\mathcal{C}}_0) \geq h_n^* - h$ (by the definition of \mathbf{V}_0), it has a minimum height of at least $h_n^* - h - \bar{M}_{S_0}^\downarrow(\mathcal{I}) \geq -h$ by the definition of $\widetilde{\mathcal{J}}_{\mathbf{F}_0}$. Recalling that $\mathbf{m}(\mathbf{F}_0) = r$, we infer from [23, Lemma 3.10] that

$$\widehat{\mu}_n^h(\widetilde{\mathcal{J}}_{\mathbf{F}_0} \mid \mathbf{I}_{\mathbf{V}_0}) = \sum_{\mathcal{I}' : \Phi^{-1}(\mathcal{I}') \in \widetilde{\mathcal{J}}_{\mathbf{F}_0}} \widehat{\mu}_n^h(\Phi^{-1}(\mathcal{I}') \mid \mathbf{I}_{\mathbf{V}_0}) \leq \sum_{\mathcal{I}' : \Phi^{-1}(\mathcal{I}') \in \widetilde{\mathcal{J}}_{\mathbf{F}_0}} e^{-(\beta-C)r} \widehat{\mu}_n^h(\mathcal{I}' \mid \mathbf{I}_{\mathbf{W}_0}) \leq e^{-(\beta-C)r}.$$

(Here we only conditioned on $\mathbf{I}_{\mathbf{V}_0}$ rather than on $\mathcal{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}$.) Via the key inequality (6.3), this implies

$$\widehat{\mu}_n^h(\mathcal{J}_{\mathbf{F}_0} \mid \mathbf{I}_{\mathbf{V}_0}) = \frac{\widehat{\mu}_n^h(\widetilde{\mathcal{J}}_{\mathbf{F}_0} \mid \mathbf{I}_{\mathbf{V}_0})}{\widehat{\mu}_n^h(\widetilde{\mathcal{J}}_{\mathbf{F}_0} \mid \mathcal{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0})} \leq e^{-(\beta-C)r}, \quad (6.5)$$

which, given that this holds for every \mathbf{V}_0 compatible with interfaces in $\mathcal{J}_{\mathbf{F}_0}$, establishes (6.1).

To prove (6.3), let $\mathcal{P} \subset \mathcal{L}_{0,n}$ be a minimum set of faces that makes $\{\rho(W) : W \in \mathbf{F}_0\} \cup \rho(W_0)$ connected (we read \mathcal{P} deterministically from \mathbf{F}_0 via an arbitrarily predefined tie breaking between sets of equal size). As we next observe, the criterion for a wall W to be part of $\text{Clust}(W_0)$ readily implies that

$$|\mathcal{P}| \leq \mathbf{m}(\mathbf{F}_0) = r;$$

indeed, initializing $\widehat{\mathbf{F}} = \{W_0\}$ and $\widehat{\mathcal{P}} = \emptyset$, consider the process of adding to $\widehat{\mathbf{F}}$ walls $W \in \mathbf{F}_0 \setminus \widehat{\mathbf{F}}$ one at a time, in tandem with adding to $\widehat{\mathcal{P}}$ the shortest path between $\rho(W)$ and $\rho(\widehat{\mathbf{F}})$, until arriving at $\widehat{\mathbf{F}} = \mathbf{F}_0$ and $\widehat{\mathcal{P}}$ which connects every $\{\rho(W) : W \in \widehat{\mathbf{F}}\}$ to $\rho(W_0)$. Every $W \in \mathbf{F}_0 \setminus \widehat{\mathbf{F}}$ must satisfy $d_\rho(W, \widehat{\mathbf{F}}) \leq \mathbf{m}(W)$ by the definition of $\text{Clust}(W_0)$, hence incurs an addition of at most $\mathbf{m}(W)$ faces to $\widehat{\mathcal{P}}$. The process terminates with $|\widehat{\mathcal{P}}| \leq \mathbf{m}(\mathbf{F}_0)$, hence the same bound applies to the minimum face set \mathcal{P} connecting these walls.

Ordering the faces of \mathcal{P} as z_1, z_2, \dots , we proceed to expose \mathfrak{G}_{z_i, S_0} for $i = 1, 2, \dots$ conditionally on $\mathfrak{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}$ and the walls already revealed along this process for $j < i$, denoting the corresponding filtration by \mathcal{F}_i . Recalling (3.5) from the proof of Proposition 3.3, we again define

$$\widehat{G}_x = \bigcap_{u \in \rho(\mathfrak{G}_{z_i, S_0})} \{\mathfrak{m}(\mathfrak{W}_{u, S_0}) < 4h_n^*\},$$

and note that by the exact same argument that followed (3.5), now with $\eta_z = d_\rho(z, \partial W_0)$, we have

$$\mu_n^\mp(\widehat{G}_{z_i} \mid \mathcal{E}_{S_0}^\eta, \mathbf{I}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}, \mathcal{F}_i) \geq \exp(-e^{-(4\beta-C)h_n^*}),$$

and the same bound under $\widehat{\mu}_n^h$. The result of iterating this bound for all $z_i \in \mathcal{P}$, followed by the bound $1 - x \geq e^{-\frac{3}{2}x}$ for $x < \frac{1}{2}$, is that for

$$\mathcal{G}_{\mathcal{P}}^m(4h_n^*) = \bigcap_{z \in \mathcal{P}} \{\mathfrak{m}(\mathfrak{W}_{z, \mathbf{C}_0}) < 4h_n^*\}$$

we have, in light of (6.4), that

$$\widehat{\mu}_n^h(\mathcal{G}_{\mathcal{P}}^m(4h_n^*) \mid \mathfrak{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}) \geq \left(1 - e^{-(4\beta-C')h_n^*}\right)^{|\mathcal{P}|} \geq \exp\left(-|\mathcal{P}|e^{-(4\beta-C')h_n^*}\right) \geq \exp(-rn^{-1+\varepsilon\beta}), \quad (6.6)$$

where we used that $h_n^* \geq (4\beta + e^{-4\beta})^{-1} \log n - C$ (see [21, Corollary 5.2]). Letting $\mathbf{W}_{\mathcal{P}} = \bigcup \{\mathfrak{G}_z : z \in \mathcal{P}\}$, note that the interface cannot have $\bar{M}_{\rho(\mathcal{C})}^\downarrow \geq h_n^*$ at $\{z\} \times \mathbb{R}$ for $z \in \rho(\mathcal{C})$ without also having $\mathfrak{m}(\mathfrak{W}_{z, \rho(\mathcal{C})}) \geq 4h_n^*$.

Conditional on $\mathcal{G}_{\mathcal{P}}^m(4h_n^*)$, we next look at events of the form $\{\bar{M}_S^\downarrow \leq h_n^*\}$ for $S = \rho(\mathcal{C}) \setminus \rho(\mathbf{W}_{\mathcal{P}})$ with $\mathcal{C} \in \mathbf{C}_0$.

Partition $\mathcal{C} \in \mathbf{C}_0$ into two subsets $\mathbf{C}'_0, \mathbf{C}''_0$ as follows:

$$\mathbf{C}'_0 := \{\mathcal{C} \in \mathbf{C}_0 : |\mathcal{C}| \geq n^{7/4} \text{ and } \dim_{\text{ip}}(\rho(\mathcal{C})) \leq 8\} \quad , \quad \mathbf{C}''_0 := \mathbf{C}_0 \setminus \mathbf{C}'_0,$$

consider $S = \rho(\mathcal{C}) \setminus \rho(\mathbf{W}_{\mathcal{P}})$ for $\mathcal{C} \in \mathbf{C}_0$, and let \mathbf{W} be an arbitrary realization of the walls of S^c under $\widehat{\mu}_n^h$ that is compatible with $\mathbf{I}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}$, $\mathcal{G}_{\mathcal{P}}^m(4h_n^*)$ and $\mathcal{E}_{\mathbf{C}_0}^\eta$. Note that, by the minimality of \mathcal{P} , the set S is simply-connected. Furthermore, having $\mathcal{I} \in \mathbf{I}_h^{\text{fl}}$ within $S \times \mathbb{R}$ is equivalent to having $\bar{M}_S^\downarrow \leq \text{ht}(\mathcal{C}) + h$; hence,

$$\widehat{\mu}_n^h(\bar{M}_S^\downarrow \leq h_n^* \mid \mathcal{E}_S^\eta, \mathbf{I}_{\mathbf{W}}) = 1 \quad \text{whenever } \text{ht}(\mathcal{C}) \leq h_n^* - h. \quad (6.7)$$

Suppose now that \mathcal{C} is such that $\text{ht}(\mathcal{C}) > h_n^* - h$. In this case, $\bar{M}_S^\downarrow \leq h_n^*$ implies that $\mathcal{I} \in \mathbf{I}_h^{\text{fl}}$, so

$$\begin{aligned} \widehat{\mu}_n^h(\bar{M}_S^\downarrow \leq h_n^* \mid \mathcal{E}_S^\eta, \mathbf{I}_{\mathbf{W}}) &= \mu_n^\mp(\bar{M}_S^\downarrow \leq h_n^* \mid \mathcal{E}_S^\eta, \mathbf{I}_{\mathbf{W}}, \bar{M}_S^\downarrow \leq \text{ht}(\mathcal{C}) + h) \\ &\geq \mu_n^\mp(\bar{M}_S^\downarrow \leq h_n^* \mid \mathcal{E}_S^\eta, \mathbf{I}_{\mathbf{W}}). \end{aligned} \quad (6.8)$$

If $\mathcal{C} \in \mathbf{C}'_0$ then $|S| \geq |\mathcal{C}| - |\mathcal{P}| \geq (1 - o(1))|\mathcal{C}|$ (recalling that $|\mathcal{P}| \leq r \leq n \log n$ and $|\mathcal{C}| \geq n^{7/4}$), whereas $|\partial S| \leq |\partial \mathcal{C}| + |\mathcal{P}| \leq (1 + o(1))|\mathcal{C}|^{7/8}$ (as $|\partial \mathcal{C}| \leq |\mathcal{C}|^{7/8}$ due to $\dim_{\text{ip}}(\mathcal{C}) \leq 8$, while $r \leq n \log n$ and $|\mathcal{C}|^{7/8} > n^{3/2}$). In particular, $\dim_{\text{ip}}(S) \leq 8 + o(1)$ in this case. By Proposition 3.3, we may thus apply (3.4), deducing that

$$\mu_n^\mp(\bar{M}_S^\downarrow \leq h_n^* \mid \mathbf{I}_{\mathbf{W}}, \mathcal{E}_S^\eta) \geq \exp\left(-\frac{1}{2}|\mathcal{C}|e^{-\alpha h_n^*+1}\right) \geq \exp\left(-\frac{1}{2}|\mathcal{C}|e^{-\alpha h_n^*+1}\right), \quad (6.9)$$

noting that we meet the assumptions for every n large enough there since $h_n^* \geq (4\beta + \varepsilon\beta)^{-1} \log n \geq \sqrt{\log |S|}$, and for every z that satisfies $d(z, \partial S) \geq e^{2\kappa_0 h_n^*}$ our constraint in η is set to $\eta_z = d(z, \partial W_0) \geq e^{2\kappa_0 h_n^*} > 5h_n^*$. For the remaining ceilings of \mathbf{C}_0 , namely $\mathbf{C} \in \mathbf{C}''_0$, we may appeal to (3.3) and derive

$$\mu_n^\mp(\bar{M}_S^\downarrow \leq h_n^* \mid \mathbf{I}_{\mathbf{W}}, \mathcal{E}_S^\eta) \geq \exp\left(-|S|e^{-(4\beta-C)h_n^*}\right) \geq \exp\left(-|\mathcal{C}|e^{-(4\beta-C)h_n^*}\right). \quad (6.10)$$

Combining (6.6)–(6.10), we conclude that

$$\widehat{\mu}_n^h(\mathfrak{J}_{\mathbf{F}_0} \mid \mathfrak{J}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}) \geq \exp\left(-n^{-1+\varepsilon\beta}r - (1 + \varepsilon\beta) \sum_{\mathcal{C} \in \mathbf{C}'_0} |\mathcal{C}|e^{-\alpha h_n^*+1} - \sum_{\mathcal{C} \in \mathbf{C}''_0} |\mathcal{C}|e^{-(4\beta-C)h_n^*}\right), \quad (6.11)$$

and it remains to account for the total area of the ceilings in \mathbf{C}'_0 and \mathbf{C}''_0 . We will use the following straightforward deterministic fact.

Fact 6.5. *Let a, r be integers, let \mathbf{W} be a set of walls with a total of r faces, and suppose there are m distinct ceilings $\{\mathcal{C}_i\}_{i=1}^m$ in $\mathbf{I}_{\mathbf{W}}$ such that $\rho(\mathcal{C}_i) \subset \bigcup_{W \in \mathbf{W}} \rho(\dot{W})$ and $|\mathcal{C}_i| \leq a$ for all i . Then $\sum_{i=1}^m |\mathcal{C}_i| \leq 3r\sqrt{a}$.*

Proof. As every \mathcal{C}_i is bounded by wall faces in \mathbf{W} , and each face of $W \in \mathbf{W}$ bounds at most 2 distinct ceiling faces, we have $\sum_{i=1}^m |\partial\mathcal{C}_i| \leq 2r$. So, via the isoperimetric inequality $|\mathcal{C}_i| \leq |\partial\mathcal{C}_i|^2$ (with a factor 4 to spare),

$$\sum_{i=1}^m |\mathcal{C}_i| \leq \sqrt{a} \sum_{i:|\partial\mathcal{C}_i| \leq \sqrt{a}} |\partial\mathcal{C}_i| + |\{i : |\partial\mathcal{C}_i| \geq \sqrt{a}\}| a \leq \sqrt{a}r + \frac{2r}{\sqrt{a}}a = 3r\sqrt{a}. \quad \blacksquare$$

For \mathbf{C}'_0 we use (along the same vein as the proof of the fact above) that $|\mathcal{C}| \leq |\partial\mathcal{C}|n/4$ and $\sum |\partial\mathcal{C}| \leq 2r$ to deduce that

$$\sum_{\mathcal{C} \in \mathbf{C}'_0} |\mathcal{C}| \leq nr/2.$$

For \mathbf{C}''_0 , recall that $\mathcal{C}_0, \dots, \mathcal{C}_L$ are the ceilings of \mathcal{I}_{W_0} nested in W_0 (as opposed to \mathbf{C}_0 , which takes $\text{Clust}(W_0)$ into consideration). For every $\mathcal{C} \in \mathbf{C}''_0$ there is a unique $i \geq 1$ such that $\rho(\mathcal{C}) \subset \rho(\mathcal{C}_i)$, so

$$\sum_{\mathcal{C} \in \mathbf{C}''_0} |\mathcal{C}| \mathbf{1}_{\{\exists i: \rho(\mathcal{C}) \subset \rho(\mathcal{C}_i), \dim_{\text{ip}}(\mathcal{C}_i) > 4\}} \leq \sum_i |\mathcal{C}_i| \mathbf{1}_{\{\dim_{\text{ip}}(\mathcal{C}_i) > 4\}} \leq n^{7/4}$$

by construction of W_0 , whereas Fact 6.5 (and having $\mathfrak{m}(\mathbf{F}_0) = r \leq n \log n$) shows that

$$\sum_{\mathcal{C} \in \mathbf{C}''_0} |\mathcal{C}| \mathbf{1}_{\{|\mathcal{C}| \leq n^{7/4}\}} = O(n^{15/8} \log n).$$

Finally, if $|\mathcal{C}| > n^{7/4}$ yet $\rho(\mathcal{C}) \subset \rho(\mathcal{C}_i)$ for some i such that $\dim_{\text{ip}}(\rho(\mathcal{C}_i)) \leq 4$, then $|\partial\mathcal{C}| \leq |\partial\mathcal{C}_i| + \mathfrak{m}(\mathbf{F}_0)$ which is at most $|\mathcal{C}_i|^{3/4} + r \leq (1 + o(1))n^{3/2}$, so $\dim_{\text{ip}}(\rho(\mathcal{C})) \leq 7 + o(1)$. Therefore, $\mathcal{C} \in \mathbf{C}'_0$ for large n .

Substituting these bounds in (6.11), along with the facts $\alpha_{h_n^*+1} \geq \alpha_{h_n^*} + 4\beta - C \geq \log n + 2\beta - C$ and $h_n^* \geq (4\beta + e^{-4\beta})^{-1} \log n - C$ from [21, Corollary 5.2], it now follows that, for some other absolute $C > 0$,

$$\widehat{\mu}_n^{\text{h}}(\bar{\mathcal{J}}_{\mathbf{F}_0} | \bar{\mathcal{J}}_{\mathbf{F}_0} \cap \mathbf{I}_{\mathbf{V}_0}) \geq \exp\left(-e^{-2\beta+C}r - n^{7/8+\varepsilon\beta}\right) \geq \exp(-C'e^{-2\beta}r), \quad (6.12)$$

where the term $n^{-1+\varepsilon\beta}r$ was absorbed into the $e^{-2\beta+C}r$ term, and the final inequality used that $r \geq n^{9/10}$. This establishes (6.3), thereby concluding the proof for all W_0 satisfying (6.2).

It remains to justify the assumption (6.2) for W_0 ; recalling that we have $\mathfrak{m}(W_0) \leq n \log n$, this property will be obtained as a consequence of our results from Section 4, as we have the following:

- (a) if $\{\mathcal{C}_i\}$ is the set of ceilings of $\mathcal{I} \sim \widehat{\mu}_n^{\text{h}}$ with $|\dot{\mathcal{C}}_i| \geq n$ and $\dim_{\text{ip}}(\dot{\mathcal{C}}_i) \geq 4$, then $|\bigcup_i \rho(\dot{\mathcal{C}}_i)| < n^{5/3}$ with probability at least $1 - \exp(-(\beta - C)n^{7/6})$ by Lemma 4.8;
- (b) if W_0 is a wall with $\mathfrak{m}(W_0) \leq n \log n$, and $\{\mathcal{C}_i\}$ is the set of ceilings of \mathcal{I}_{W_0} that are nested in W_0 and satisfy $|\mathcal{C}_i| \leq n$, then Fact 6.5 shows that, deterministically, $\sum_i |\mathcal{C}_i| = O(n^{3/2} \log n)$.

Combining these (along with a union bound on W_0 and its location) we see that W_0 satisfies (6.2) (we arrived at an upper bound of $n^{5/3} + n^{3/2+o(1)} < n^{7/4}$) except with probability $O(\exp(-n^{7/6-o(1)}))$, as claimed.

To obtain the result for a collection of walls $W_0^{(1)}, \dots, W_0^{(\ell)}$, observe that our bound on the probability of finding a given such W_0 was conditional on \mathbf{V}_0 , an arbitrary realization of walls in $\rho(\dot{W}_0)^c$. Therefore, for any fixed collection of such walls, we may iteratively bound the probability of $W_0^{(i)}$ conditional on the occurrence of $\{W_0^{(j)}\}_{j < i}$, provided that the hulls of said walls are disjoint (guaranteed by our hypothesis), and obtain an overall bound of $\exp(-(\beta - C)(\ell r \wedge n \log n))$. A union bound over the location and structure of these walls thus concludes the proof. \blacksquare

Remark 6.6. One can extend Theorem 6.2—and in turn, Theorem 6.1—to the case of $\mathfrak{h} = h_n^* - 1$, in the same vein as Remark 5.8, provided that the quantity $\gamma = n \exp(-\alpha_{h_n^*})$ (which satisfies $e^{-2\beta-\varepsilon\beta} \leq \gamma < e^{2\beta}$) is in a given range, e.g.,

$$\gamma \leq \beta.$$

We claim that (6.1) holds for $\bar{\mathcal{J}}_{\mathbf{F}_0}$ defined w.r.t. a supporting ceiling $\widehat{\mathcal{C}}_0 \in \mathfrak{C}_{\geq 0}$. We accordingly modify $\bar{\mathcal{J}}_{\mathbf{F}_0}$ to feature a strict inequality: $\bar{\mathcal{J}}_{\mathbf{F}_0} = \{\mathcal{I} \in \bar{\mathcal{J}}_{\mathbf{F}_0} : \bar{M}_{S_0}^\downarrow < h_n^*\}$, and claim that the arguments hold true thereafter. Indeed, since $\text{ht}(\widehat{\mathcal{C}}_0) \geq 0$, having $\bar{M}_{S_0}^\downarrow \leq h_n^* - 1 = \mathfrak{h}$ would guarantee that $\Phi(\mathcal{I}) \in \mathbf{I}_n^{\text{fl}}$ for every $\mathcal{I} \in \bar{\mathcal{J}}_{\mathbf{F}_0}$, whence (6.5) remains true. For the analogue of (6.7), we now have that $\widehat{\mu}_n^{\text{h}}(\bar{M}_S^\downarrow < h_n^* | \mathcal{E}_S^\eta, \mathbf{I}_{\mathbf{W}}) = 1$ whenever $\text{ht}(\mathcal{C}) \leq 0$, as $\mathfrak{h} = h_n^* - 1$; and for (6.8), in case $\text{ht}(\mathcal{C}) > 0$, we can again bound $\widehat{\mu}_n^{\text{h}}(\bar{M}_S^\downarrow < h_n^* | \mathcal{E}_S^\eta, \mathbf{I}_{\mathbf{W}})$ from below by $\mu_n^\mp(\bar{M}_S^\downarrow < h_n^* | \mathcal{E}_S^\eta, \mathbf{I}_{\mathbf{W}})$. We conclude that (6.11) holds for this case when replacing $\alpha_{h_n^*+1}$ by $\alpha_{h_n^*}$,

and so (6.12) holds true if the term $e^{-2\beta+C}r$ is to be replaced by $(1 + \varepsilon_\beta)\gamma r/2$. Our assumption $\gamma \leq \beta$ thus supports the analogue of (6.3) with a lower bound of $\exp(-(1 + \varepsilon_\beta)(\beta/2)r)$, which is outweighed by the term $\exp(-(\beta - C)r)$ from (6.5) to give a final probability estimate of $\exp(-(\beta/2 - C)r \wedge n \log n)$.

Acknowledgements. R.G. thanks the Miller Institute for Basic Research in Science for its support. The research of E.L. was supported in part by NSF grants DMS-1812095 and DMS-2054833.

REFERENCES

- [1] D. B. Abraham. Solvable model with a roughening transition for a planar ising ferromagnet. *Phys. Rev. Lett.*, 44:1165–1168, May 1980.
- [2] K. S. Alexander, F. M. Dunlop, and S. Miracle-Sole. Layering in the Ising model. *Journal of Statistical Physics*, 141(2):217–241, Oct. 2010.
- [3] A. G. Basuev. Ising model in half-space: A series of phase transitions in low magnetic fields. *Theoretical and Mathematical Physics*, 153(2):1539–1574, 2007.
- [4] E. Bolthausen, J.-D. Deuschel, and G. Giacomin. Entropic repulsion and the maximum of the two-dimensional harmonic crystal. *Ann. Probab.*, 29(4):1670–1692, 2001.
- [5] E. Bolthausen, J.-D. Deuschel, and O. Zeitouni. Entropic repulsion of the lattice free field. *Communications in Mathematical Physics*, 170(2):417 – 443, 1995.
- [6] 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.
- [7] M. Campanino, D. Ioffe, and Y. v. Velenik. Ornstein-Zernike theory for finite range ising models above T_c . *Probability Theory and Related Fields*, 125(3):305–349, 2003.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] J. T. Chalker. The pinning of an interface by a planar defect. *Journal of Physics A: Mathematical and General*, 15(9):L481–L485, sep 1982.
- [12] J.-D. Deuschel. Entropic repulsion of the lattice free field. II. The 0-boundary case. *Communications in Mathematical Physics*, 181(3):647 – 665, 1996.
- [13] J.-D. Deuschel and G. Giacomin. Entropic repulsion for the free field: Pathwise characterization in $d \geq 3$. *Communications in Mathematical Physics*, 206:447–462, 1999.
- [14] R. Dobrushin, R. Kotecký, and S. Shlosman. *Wulff construction*, volume 104 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1992.
- [15] R. L. Dobrushin. Description of a random field by means of conditional probabilities and conditions for its regularity. *Theory of Probability & Its Applications*, 13(2):197–224, 1968.
- [16] R. L. Dobrushin. The Gibbs state that describes the coexistence of phases for a three-dimensional Ising model. *Teor. Veroyatnost. i Primenen.*, 17:619–639, 1972.
- [17] J. Fröhlich and C.-E. Pfister. Semi-infinite Ising model. I. Thermodynamic functions and phase diagram in absence of magnetic field. *Comm. Math. Phys.*, 109(3):493–523, 1987.
- [18] J. Fröhlich and C.-E. Pfister. Semi-infinite Ising model. II. The wetting and layering transitions. *Comm. Math. Phys.*, 112(1):51–74, 1987.
- [19] S. Ganguly and R. Gheissari. Local and global geometry of the 2D Ising interface in critical prewetting. *The Annals of Probability*, 49(4):2076 – 2140, 2021.
- [20] H. Georgii. *Gibbs Measures and Phase Transitions*. De Gruyter studies in mathematics. W. de Gruyter, 1988.
- [21] R. Gheissari and E. Lubetzky. Tightness and tails of the maximum in 3D Ising interfaces. *Ann. Probab.*, 49(2):732–792, 2021.
- [22] R. Gheissari and E. Lubetzky. Maximum and shape of interfaces in 3D Ising crystals. *Communications on Pure and Applied Mathematics*, 75(12):2575–2684, 2022.
- [23] R. Gheissari and E. Lubetzky. Approximate Domain Markov property for rigid Ising interfaces. *J. Stat. Phys.*, 190(5):Paper No. 99, 2023.
- [24] L. Greenberg and D. Ioffe. On an invariance principle for phase separation lines. *Annales de l’I.H.P. Probabilités et statistiques*, 41(5):871–885, 2005.
- [25] P. Holický and M. Zahradník. On entropic repulsion in low temperature Ising models. In N. Boccara, E. Goles, S. Martinez, and P. Picco, editors, *Cellular Automata and Cooperative Systems*, pages 275–287. Springer Netherlands, Dordrecht, 1993.
- [26] O. Hryniv. On local behaviour of the phase separation line in the 2D Ising model. *Probab. Theory Related Fields*, 110(1):91–107, 1998.
- [27] D. Ioffe, S. Ott, S. Shlosman, and Y. Velenik. Critical prewetting in the 2D Ising model, 2021.
- [28] D. Ioffe, S. Ott, Y. Velenik, and V. Wachtel. Invariance principle for a potts interface along a wall. *Journal of Statistical Physics*, 180(1):832–861, 2020.
- [29] D. Ioffe, S. Shlosman, and F. L. Toninelli. Interaction versus entropic repulsion for low temperature ising polymers. *Journal of Statistical Physics*, 158(5):1007–1050, 2015.

- [30] D. Ioffe and Y. Velenik. Low-temperature interfaces: prewetting, layering, faceting and Ferrari-Spohn diffusions. *Markov Process. Related Fields*, 24(3):487–537, 2018.
- [31] H. Lacoïn. Wetting and layering for solid-on-solid i: Identification of the wetting point and critical behavior. *Communications in Mathematical Physics*, 362(3):1007–1048, 2018.
- [32] E. Lubetzky, F. Martinelli, and A. Sly. Harmonic pinnacles in the discrete Gaussian model. *Comm. Math. Phys.*, 344(3):673–717, 2016.
- [33] I. Melicherčík. Entropic repulsion of two interfaces in ising model, 1991.
- [34] R. A. Minlos and J. G. Sinaĭ. The phenomenon of “separation of phases” at low temperatures in certain lattice models of a gas. I. *Mat. Sb. (N.S.)*, 73 (115):375–448, 1967.
- [35] C.-E. Pfister and Y. Velenik. Large deviations and continuum limit in the 2D Ising model. *Probab. Theory Related Fields*, 109(4):435–506, 1997.
- [36] C.-E. Pfister and Y. Velenik. Interface, surface tension and reentrant pinning transition in the 2D Ising model. *Comm. Math. Phys.*, 204(2):269–312, 1999.
- [37] J. M. Steele. Probabilistic and worst case analyses of classical problems of combinatorial optimization in Euclidean space. *Math. Oper. Res.*, 15(4):749–770, 1990.
- [38] H. van Beijeren. Interface sharpness in the Ising system. *Comm. Math. Phys.*, 40(1):1–6, 1975.
- [39] Y. Velenik. Entropic repulsion of an interface in an external field. *Probability Theory and Related Fields*, 129(1):83–112, May 2004.
- [40] Y. Velenik. Localization and delocalization of random interfaces. *Probab. Surv.*, 3:112–169, 2006.
- [41] S. Verblunsky. On the shortest path through a number of points. *Proc. Amer. Math. Soc.*, 2:904–913, 1951.

R. GHEISSARI

DEPARTMENT OF MATHEMATICS, NORTHWESTERN UNIVERSITY

Email address: `gheissari@northwestern.edu`

E. LUBETZKY

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY 10012, USA.

Email address: `eyal@courant.nyu.edu`