

METASTABILITY CASCADES AND PREWETTING IN THE SOS MODEL

REZA GHEISSARI AND EYAL LUBETZKY

ABSTRACT. We study Glauber dynamics for the low temperature $(2 + 1)$ D Solid-On-Solid model on a box of side-length n with a floor at height 0 (inducing entropic repulsion) and a competing bulk external field λ pointing down (the prewetting problem). In 1996, Cesi and Martinelli showed that if the inverse-temperature β is large enough, then along a decreasing sequence of critical points $(\lambda_c^{(k)})_{k=0}^{K_\beta}$ the dynamics is torpid: its inverse spectral gap is $O(1)$ when $\lambda \in (\lambda_c^{(k+1)}, \lambda_c^{(k)})$ whereas it is $\exp[\Theta(n)]$ at each $\lambda_c^{(k)}$ for each $k \leq K_\beta$, due to a coexistence of rigid phases at heights $k + 1$ and k . Our focus is understanding (a) the onset of metastability as $\lambda_n \uparrow \lambda_c^{(k)}$; and (b) the effect of an unbounded number of layers, as we remove the restriction $k \leq K_\beta$, and even allow for $\lambda_n \rightarrow 0$ towards the $\lambda = 0$ case which has $O(\log n)$ layers and was studied by Caputo et al. (2014). We show that for any k , possibly growing with n , the inverse gap is $\exp[\tilde{\Theta}(1/|\lambda_n - \lambda_c^{(k)}|)]$ as $\lambda \uparrow \lambda_c^{(k)}$ up to distance $n^{-1+o(1)}$ from this critical point, due to a metastable layer at height k on the way to forming the desired layer at height $k + 1$. By taking $\lambda_n = n^{-\alpha}$ (corresponding to $k_n \asymp \log n$), this also interpolates down to the behavior of the dynamics when $\lambda = 0$. We complement this by extending the fast mixing to all λ uniformly bounded away from $(\lambda_c^{(k)})_{k=0}^\infty$. Together, these results provide a sharp understanding of the predicted infinite sequence of dynamical phase transitions governed by the layering phenomenon.

1. INTRODUCTION

Consider the low-temperature Solid-On-Solid (SOS) model on $\Lambda_n = \llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2$, with zero boundary conditions, a floor at height zero, and an external field $\lambda = \lambda_n \geq 0$ pointing downward: The model assigns to a height function $\varphi : \Lambda_n \rightarrow \llbracket 0, n \rrbracket$ (viewed as a surface inside a cube of side length n) the probability

$$\mu_{n,\lambda}(\varphi) \propto \exp\left(-\beta \sum_{u \sim v} |\varphi_u - \varphi_v| - \lambda \sum_v \varphi_v\right), \quad (1.1)$$

where $\beta > 0$ is the inverse-temperature, we let $u \sim v$ denote nearest-neighbor adjacency between the vertices u, v , and φ_x is taken to be 0 for all $x \notin \Lambda_n$ as per the boundary conditions. The study of this family of models in the statistical physics literature goes back to Burton, Cabrera and Frank [4] in 1951 in dimension $(2 + 1)$ and to Temperley [46] in 1952 in dimension $(1 + 1)$ as models of crystal formation/growth and approximations to the low-temperature plus/minus interface in the Ising model.

The model at $\lambda = 0$ and no floor (the integer height function φ can also take negative values) is associated with a *roughening phase transition* exclusive to dimension $2 + 1$: for some $\beta_R > 0$ (numerical simulations suggest $\beta_R \approx 0.806$) the SOS surface is rough at $\beta \leq \beta_R$ (e.g., the variance of φ_o diverges with n) whereas it is rigid (the variance is $O(1)$) for $\beta > \beta_R$ (see [20, 21] and [2] for $\beta \ll 1$ and $\beta \gg 1$, resp., and [32] proving they form a dichotomy). When setting the surface above a floor (a hard constraint $\varphi \geq 0$) at the low temperature regime (β large), it exhibits *entropic repulsion* [3]: the average height is propelled from $O(1)$ to $\Theta(\log n)$, despite the energetic cost charged along the zero boundary conditions, in order to gain entropy via downward-pointing spikes. When $\lambda = \lambda_n > 0$, the field induces an additional downwards force on the interface, inducing a non-trivial competition with the entropic repulsion, known as the *prewetting regime*.

We study Glauber dynamics for the SOS model from Eq. (1.1)—a continuous-time single-site Markov chain that, on one hand, gives a simple local recipe for sampling $\mu_{n,\lambda}$, and on the other hand, serves as a natural physical model for the evolution of a random surface towards the SOS measure at equilibrium. When the SOS surface is concentrated at height h (e.g., $h = \Theta(\log n)$ when $\lambda = 0$), understanding the rate of convergence of this dynamics to $\mu_{n,\lambda}$ is closely related to the time it takes the surface to successively create h macroscopic layers (each grown out of local droplets) starting, e.g., from the all-zero configuration.

More precisely, the Glauber dynamics assigns every site in Λ_n a rate-1 Poisson clock, and every time such a clock rings, it updates the height φ_v at the corresponding site v to one of $\{\varphi_v - 1 \vee 0, \varphi_v, \varphi_v + 1 \wedge n\}$, weighted by the conditional probabilities induced on these by $\mu_{n,\lambda}$ given $\{\varphi_u : u \neq v\}$. Said probabilities (and by extension, the update rule at v) depend only on $\{\varphi_u : u = v \text{ or } u \sim v\}$ due to the nearest-neighbor interactions in Eq. (1.1), and by construction, the chain is reversible w.r.t. its stationary distribution $\mu_{n,\lambda}$. A standard way to measure its rate of convergence to equilibrium is the spectral-gap of its generator, denoted $\text{gap}_{n,\lambda}$, which governs the mixing time in L^2 -distance, and its dependence on the system size n as $n \rightarrow \infty$.

With no external field ($\lambda = 0$), Caputo et al. [9] showed that the inverse gap of Glauber dynamics for the (2 + 1)D SOS model above a floor satisfies $\text{gap}_{n,\lambda}^{-1} = \exp[\Theta(n)]$, as the surface encounters a sequence of metastable states towards equilibrium: starting from the all-zero state, it takes it time doubly-exponential in k —approximately $\exp[ce^{4\beta k}]$ —to create a new macroscopic layer at height k , with the bottlenecks at the final levels of height $\sim \frac{1}{4\beta} \log n$ dominating the mixing time and costing $\exp[cn]$.

In the presence of a bulk external field $\lambda > 0$ pointing downward, the competition between the field and the entropic repulsion may lower the typical surface height. For $\lambda > 0$ fixed, this lowers the SOS surface to height $k = k(\lambda)$, where as we vary $\lambda \downarrow 0$, the preferred height $k(\lambda)$ of the surface will grow as $\frac{1}{4\beta} \log \frac{1}{\lambda}$. Due to the rigidity of the surface, this induces a sequence of critical points $\lambda_c^{(k)} \approx e^{-4\beta k}$ at which the two heights $k + 1, k$ are both (equally) stable. (A similar infinite sequence of critical points along which the interface height diverges occurs in the *wetting problem*, where in lieu of the bulk external field λ the surface is tilted by $\lambda \sum_v \mathbf{1}_{\varphi_v=0}$, rewarding only sites that are pinned to height 0. A detailed understanding of this was developed by Lacoïn [30, 31] confirming the predicted phenomena [13]. See also [1, 17] and the excellent survey by Ioffe and Velenik [27] for more on the layering phenomena associated with wetting/pinning.) In our prewetting framework, at sufficiently low temperatures (large enough β), an extensive study of the statics and Glauber dynamics for this model¹ by Cesi and Martinelli [11, 12] in 1996 established a finite number of dynamical phase transitions $(\lambda_c^{(k)})_{k=0}^{K_\beta}$.

Theorem ([12]). Fix $\beta > 0$ large enough, and consider Glauber dynamics w.r.t. $\mu_{n,\lambda}$. There is a sequence of critical points $(\lambda_c^{(k)})_{k=0}^{K_\beta}$, where $K_\beta = \lfloor \exp(\frac{\beta}{20000}) \rfloor$, such that for every $k = 0, \dots, K_\beta$:

1. If $\lambda \in (\lambda_c^{(k+1)}, \lambda_c^{(k)})$ then $\text{gap}_{n,\lambda}^{-1} = \Theta(1)$ (moreover, under any boundary condition);
2. If $\lambda = \lambda_c^{(k)}$ then $\text{gap}_{n,\lambda}^{-1} = \exp[\Theta(n)]$ under free boundary conditions;

where the implicit constants depends on both β and λ .

Our aim here is to (a) understand the interpolation between the two behaviors given by [11, 12]—exponentially slow mixing at $\lambda = \lambda_c^{(k)}$ as opposed to fast mixing at $\lambda \neq \lambda_c^{(k)}$, as we take $\lambda = \lambda_n \rightarrow \lambda_c^{(k)}$; and (b) obtain the complete picture of an *infinite* sequence of dynamical phase transitions by removing the restriction on K_β , and en route, show the interpolation from the λ fixed case to the $\lambda = 0$ case studied in [9].

It ought to be noted that the analysis of the model in [11, 12] is quite involved, and itself was built on (and refined) a highly nontrivial study of the model in infinite-volume by Dinaburg and Mazel [15] (where a full infinite sequence of critical values $(\lambda_c^{(k)})_{k=0}^\infty$ for limiting free energies was derived). As mentioned in [11] as well as in the survey [27, §4.1.1], it was unclear whether extending the results of [11, 12] to $K_\beta \uparrow \infty$ (ideally up to $k = \Theta(\log n)$ where the $\lambda = 0$ behavior is observed) was technical or would required different ideas. (An outline of how some of the equilibrium results in [11] may be pushed beyond the restriction $k \leq K_\beta$ was sketched by Lebowitz and Mazel [33], albeit still for λ and k fixed independently of n .)

Our main theorem establishes the full infinite sequence of metastability windows of λ about $(\lambda_c^{(k)})_{k=0}^\infty$, with the distance of λ_n to the nearest critical point determining the rate of the exponential slowdown. Namely, the rate of metastability is determined by the time it takes a surface that is predominantly at height k to generate a critical droplet at the desired height $k + 1$. The size needed for the droplet to be stable

¹The setup of [11, 12] technically considers $\varphi : \Lambda_n \rightarrow \mathbb{Z}_{\geq 0}$, i.e., no upper bound on the maximal height. For λ bounded away from 0 uniformly in n , as they were considering, results with and without a ceiling readily transfer back and forth.

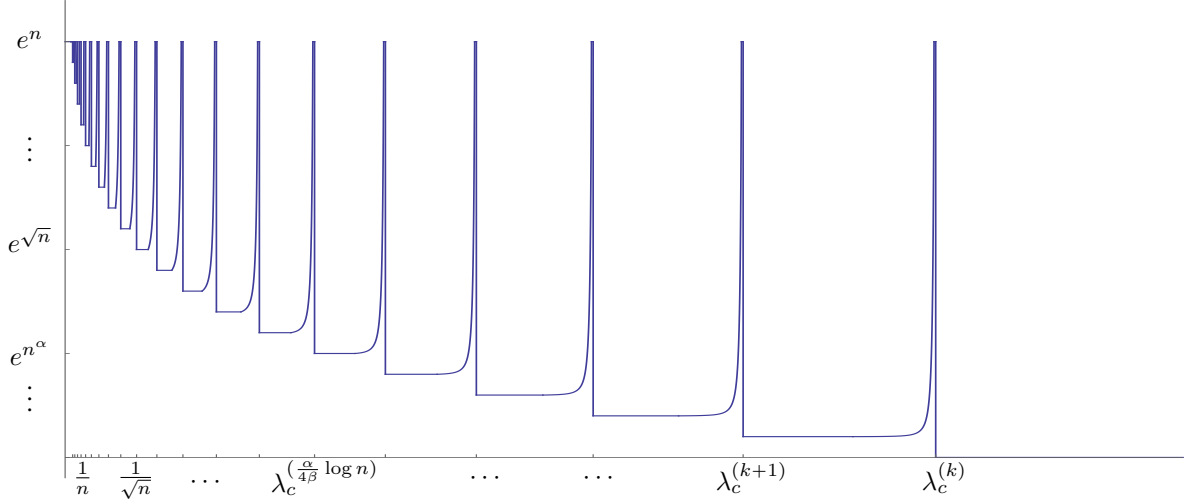


FIGURE 1. The inverse gap $\text{gap}_{n, \lambda_n}^{-1}$ under zero boundary conditions plotted against λ , demonstrating the diverging sequence of metastability cascades as $\lambda \downarrow 0$. The slowdowns near the critical points are expected to be asymmetrical due to the boundary condition.

is in turn determined by an isoperimetric tradeoff and the difference of λ to $\lambda_c^{(k)}$. (A similar metastability induced by the time to grow a critical droplet is seen in the dynamics of a low-temperature 2D Ising model with (+)-boundary and a competing small (−) external field, as was studied in [45]: see also [43, 44]).

Theorem 1.1. *Consider the SOS Glauber dynamics at $\beta > \beta_0$ large enough and a bulk external field $\lambda_n > 0$ with zero boundary as per Eq. (1.1). There exists a sequence of critical points $(\lambda_c^{(k)})_{k=0}^\infty$ such that the following holds. Let*

$$\mathbf{d}^+(x) := \min_k \{(\lambda_c^{(k)} - x) : \lambda_c^{(k)} \geq x\}, \quad \text{and} \quad \mathbf{d}(x) := \min_k \{|\lambda_c^{(k)} - x|\}. \quad (1.2)$$

For every $\varepsilon > 0$ there exist constants $c(\varepsilon), C(\beta, \varepsilon) > 0$ such that for all n ,

(1) If λ_n is such that $\mathbf{d}(\lambda_n) > \varepsilon$ then $C^{-1} \leq \text{gap}_{n, \lambda_n}^{-1} \leq C$.

(2) If λ_n is such that $\frac{c\beta \log n}{n} \leq \mathbf{d}^+(\lambda_n) \leq \varepsilon$ then²

$$\frac{1}{C} \vee \frac{1}{n^2} \exp \left[\frac{1}{C \mathbf{d}^+(\lambda_n)} \right] \leq \text{gap}_{n, \lambda_n}^{-1} \leq \exp \left[\frac{C(\log n)^3}{\mathbf{d}(\lambda_n)} \right]. \quad (1.3)$$

The following corollary sharply characterizes the onset of metastability in the two regimes which served as our earlier motivation: (a) as $\lambda_n \uparrow \lambda_c^{(k_n)}$, and (b) as $\lambda_n \downarrow 0$.

Corollary 1.2. *In the setting of Theorem 1.1, for any fixed $0 < \alpha < 1$, in either of the following situations:*

(a) $\lambda_n = \lambda_c^{(k)} - \Theta(n^{-\alpha})$ for any $1 \leq k \leq o(\log n)$,

(b) $\lambda_n = \Theta(n^{-\alpha})$ such that $\mathbf{d}(\lambda_n) \gtrsim n^{-\alpha}$ (e.g., the sequence $\lambda_n = e^{-\beta - 4\beta k_n}$ for $k_n = \lfloor \frac{\alpha}{4\beta} \log n \rfloor$),

we have $\text{gap}_{n, \lambda_n}^{-1} = \exp[\tilde{\Theta}(n^\alpha)]$.

In Item (a) of Corollary 1.2, the typical interface will be localized at height $k + 1$, but under k boundary conditions, only height- $(k + 1)$ layers of size greater than n^α are desirable (smaller ones do not have enough room for the entropic repulsion to propel them to height $k + 1$). As the dynamics climbs from below, say initialized at all-zero, this generates a bottleneck, whereby there is an exponential in n^α waiting time to randomly form such an atypically large droplet. In Item (b), initialized from all-zero, there is a cascade of

²This restriction can in fact be replaced with $\frac{c\beta k}{n} \leq \mathbf{d}^+(\lambda_n) \leq \varepsilon$ where $k = \text{argmin}_k \{(\lambda_c^{(k)} - x) : \lambda_c^{(k)} > x\}$.

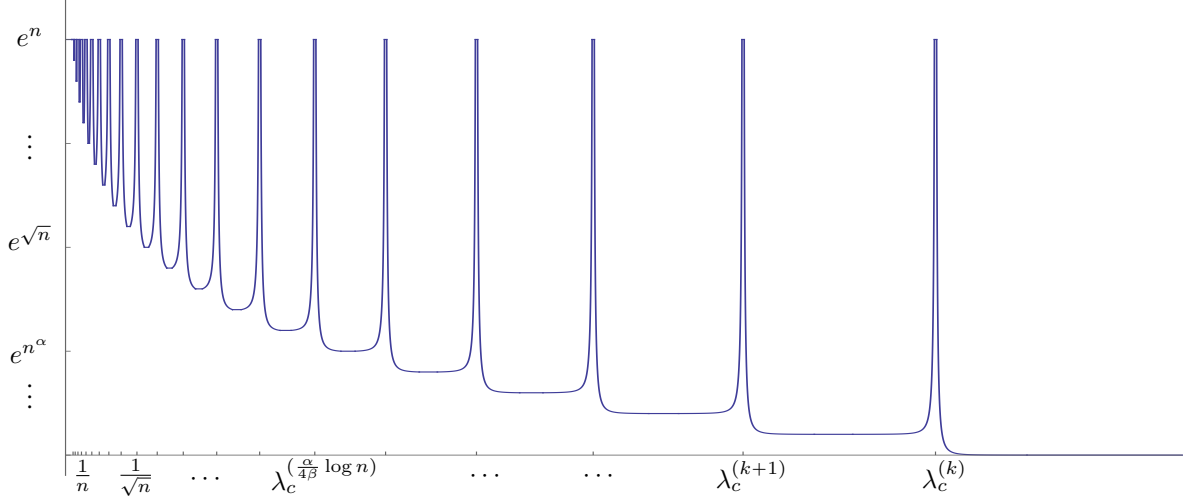


FIGURE 2. The inverse gap $\text{gap}_{n, \lambda_n}^{-1}$ for the model under periodic boundary conditions plotted against λ . The critical slowdowns are symmetric, as per the bounds of Theorem 1.3.

metastable transitions with times exponential in $(\lambda_c^{(k)} - \lambda_n)_{k \geq 0}$ to go from k to $k + 1$, with the final one to get to $k_n = \lfloor \frac{\alpha}{4\beta} \log n \rfloor$ dominating the mixing time. See Fig. 1 for a rough depiction of the above.

As evident from the fact that Eq. (1.3) features $d^+(\lambda_n)$ in the lower bound yet $d(\lambda_n)$ in the upper bound (measuring the distance to the closest critical point $\lambda_c^{(k)}$ in the latter yet only to the closest $\lambda_c^{(k)}$ above λ_n in the former), we expect the mixing time to be asymmetric in the window about each $\lambda_c^{(k)}$ (see Fig. 1). Indeed, when $d(\lambda_n) \ll d^+(\lambda_n)$, e.g., $\lambda_n \downarrow \lambda_c^{(k)}$, the typical interface is rigid about height k , but in a region with boundary conditions $k + 1$, only height- k layers of size greater than $d(\lambda_n)^{-1}$ will be desirable. This would induce an exponential slowdown if the boundary conditions were above height k , but the height-0 boundary conditions enable the demolition of the height- k layer from the outside inwards (via droplet shrinkage). This dynamical picture—in particular, if we identify heights $\{\leq k\}$ as minus spins and heights $\{> k\}$ as plus spins—is in analogy with (and more complicated than) the dynamics of the low-temperature Ising model with minus boundary conditions, initialized from all-plus (possibly with a small negative field of $d(\lambda_n)$). The latter is a notoriously hard open problem in the mixing time literature: see [36] for the best known bounds. In particular, for fixed k , exactly at $\lambda_c^{(k)}$, and with zero boundary conditions, one would expect the mixing time to be polynomially fast, with a $\Theta(n^2)$ inverse gap. We discuss the delicate expected behavior in this regime in more detail in Section 1.2.

When considering the SOS model under *periodic* (instead of zero) boundary conditions—that is, Λ_n in Eq. (1.1) is replaced by $(\mathbb{Z}/n\mathbb{Z})^2$ —the above asymmetry is no longer present, as illustrated by the next theorem. (See Fig. 2 for a depiction of the bounds of Theorem 1.3 symmetric about $(\lambda_c^{(k)})_{k=0}^\infty$.)

Theorem 1.3. *In the setting of Theorem 1.1 but with periodic boundary conditions, so long as $\lambda_n \geq \frac{\log n}{n}$, for every $\varepsilon > 0$ there exists $C(\beta, \varepsilon)$ such that for all n ,*

- (1) *If λ_n is such that $d(\lambda_n) \geq \varepsilon$, then $C^{-1} \leq \text{gap}_{n, \lambda_n}^{-1} \leq C$.*
- (2) *If λ_n is such that $d(\lambda_n) \leq \varepsilon$, then*

$$\frac{1}{C} \vee \frac{1}{n^2} \exp \left[\frac{1}{C} \left(n \wedge \frac{1}{d(\lambda_n)} \right) \right] \leq \text{gap}_{n, \lambda_n}^{-1} \leq \exp \left[C \left(n \log n \wedge \frac{(\log n)^3}{d(\lambda_n)} \right) \right].$$

The case of free boundary conditions, which was the setting in which [12] showed $\exp[\Theta(n)]$ mixing at $(\lambda_c^{(k)})_{k=0}^{K_\beta}$, is similar to the torus in that the boundary conditions do not favor one of $\{k, k + 1\}$. If λ_n is uniformly bounded away from 0, the factor of $\log n$ in the upper bound of Item 2 of Theorem 1.3 can similarly be dropped (see Remark 6.5), thus recovering for a fixed λ , bounded away from 0, the same bounds

for the torus as obtained in [12] for free boundary conditions. Moreover, Theorem 1.3 allows λ to be taken arbitrarily close to 0, as well as λ_n approaching 0 or a critical point as $n \rightarrow \infty$. While we did not pursue this, it should be possible to modify the arguments of this paper to the case of free boundary conditions.

We conclude this part with a brief comment on other possible boundary conditions. In analogy to the low-temperature 2D Ising model, the metastable windows can be quite sensitive to the choice of boundary conditions. If the boundary conditions were high, rather than at 0, they would favor height $k + 1$ over k near $\lambda_c^{(k)}$, and Item 2 of Theorem 1.1 would apply with d^+ replaced by $d^- = \min_k \{(x - \lambda_c^{(k)}) : \lambda_c^{(k)} \leq x\}$. Dobrushin-type boundary conditions, say $\{\leq k\}$ on half the boundary and $\{> k\}$ on the other half, should induce interfaces and rich static and dynamic behaviors: see the discussion of these interfaces as $d(\lambda_n) \downarrow 0$ in Section 1.2. Finally, the $O(1)$ inverse gap bound of Item 1 in Theorem 1.1 applies for all fixed boundary conditions, so long as n is sufficiently large as a function of ε , extending the fast mixing results of [12] (which needed $\lambda > \lambda_c^{(K_\beta)}$) to hold for λ arbitrarily close to 0 (while kept $\Omega(1)$ in n).

1.1. Proof ideas. The proofs in our paper are divided into two parts, the equilibrium portion comprising Sections 3 to 5 and the dynamical portion comprising Sections 6 and 7.

Our analysis of the equilibrium estimates differs from the works [11, 15] first and foremost via its focus on *level lines* (viewing the SOS configuration as a contour ensemble, each associated to a level using appropriate boundary conditions). The aforementioned previous works on the $\lambda > 0$ setting studied *cylinders* as the main objects—maximal stacks of contours (all having the same projection on \mathbb{Z}^2). Operations on contours (such as Peierls-type maps that shift a single level line up or down) can in general allow a much more refined analysis than operations on cylinders, as was demonstrated in the work [9] for the $\lambda = 0$ case. This different perspective does introduce additional complications, yet is essential in our extension of the previous results to any height k (as opposed to $k \leq K_\beta$). Namely, by combining probabilistic arguments on the set of contours (based on coarse graining, monotonicity, and Peierls maps) with the (already quite sophisticated) cluster-expansion based techniques from [11, 15] (adapted from cylinders to our setting of contours in Sections 3 and 4), we arrive at the following understanding of the shape of an SOS surface with zero boundary conditions in the presence of an external field λ . Note the subtle relation between the height k the surface reaches, the distance d^+ of the external field to its nearest critical point, and the domain size n .

Theorem 1.4. *For every $\beta > \beta_0$ there exist some $C > 0$, an ε_β (going to zero as $\beta \rightarrow \infty$), and an infinite sequence of critical points $(\lambda_c^{(k)})_{k=0}^\infty$ such that, if $\lambda_n \geq \frac{\log n}{n}$ satisfies $\lambda_n \in (\lambda_c^{(k_n+1)}, \lambda_c^{(k_n)})$ then with μ_{n, λ_n} -probability at least $1 - e^{-n/C}$,*

- (1) *if $n \geq C\beta k_n/d^+(\lambda_n)$, all but an ε_β fraction of the sites will be at height $k_n + 1$.*
- (2) *If $n \leq 1/(Cd^+(\lambda_n))$, all but an ε_β fraction of the sites in Λ_n will be at height k_n .*

We in fact prove stronger equilibrium results in Section 5, showing *spatial mixing* estimates between SOS models with different boundary conditions that hold on boxes of side-length $n_0 \gtrsim k_n/d(\lambda_n)$, yet fail when $n_0 \lesssim 1/d(\lambda_n)$. Note that the analysis of contours, beyond its effectiveness in capturing the behavior of the surface in regimes of λ unaddressed by [11], opens the door to a host of questions addressing the refined geometry of the individual level lines (notably the top one): e.g., understanding their macroscopic scaling limits and the random fluctuations around those as λ approaches $\lambda_c^{(k)}$; see Section 1.2 for concrete examples.

Turning to the dynamical arguments, our proofs of the exponential lower and upper bounds are guided by the following mechanism for equilibration, which we describe here for the case of zero boundary condition and $\lambda_n = \lambda_c^{(k_n)} - d^+(\lambda_n)$. If we let $n \geq n_0 = C\beta k_n/d^+(\lambda_n)$, then on a domain of size n , the equilibrium surface is predominantly at height $k_n + 1$ per Theorem 1.4. On the other hand, if the dynamics is initialized at heights $\{\leq k_n\}$ it needs to grow a layer at height $k_n + 1$ of size at least n_0 until that layer would become thermodynamically stable and could grow to encompass a $(1 - \varepsilon_\beta)$ -fraction of Λ_n (if the layer has diameter smaller than n_0 , the exponential cost in its boundary dominates the free energy benefit of the area it confines being at height $k_n + 1$ over k_n). Since the Glauber dynamics makes single-site updates, it in particular needs to first form a layer of diameter between $n_0/2$ and n_0 , an event that has probability $e^{-n_0/C}$. Once such a

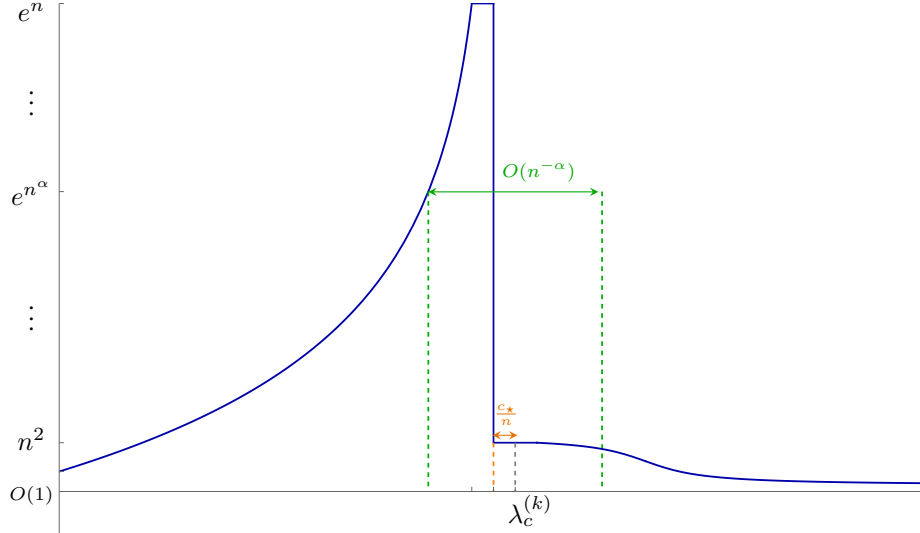


FIGURE 3. The predicted behavior of the inverse gap with zero boundary conditions when $\lambda_n = \lambda_c^{(k)} \pm o(1)$ for fixed $k \geq 0$. Theorem 1.1 establishes the curve left of the discontinuity (up to poly-logarithmic factors); the right side of the curve should be governed by stable layers shrinking via mean-curvature flow, leading to an $O(n^2)$ relaxation time at $\lambda_c^{(k)}$.

droplet has been formed, without waiting much longer, it will have formed around every site in Λ_n , and therefore the $k_n + 1$ layer will in fact encompass almost all of Λ_n , leading to rapid equilibration.

For the lower bounds on $\text{gap}_{n,\lambda_n}^{-1}$, the above intuition is used to construct a bottleneck event for Glauber dynamics (see Section 7). For the upper bound, in Section 6, we guide the dynamics towards equilibrium, using censoring and monotonicity to reduce the mixing time on Λ_n to the mixing time on graphs of cut-width at most n_0 , where n_0 is the smallest scale at which the model exhibits weak spatial mixing. While reductions of mixing times to weak spatial mixing are common in the literature (e.g., the classical work [39]), boundary conditions pose difficulties when only weak spatial mixing, but not strong spatial mixing, holds. We handle this difficulty by also showing weak spatial mixing estimates in annuli of thickness n_0 , whose mixing time will also only be exponential in n_0 , and which is used to couple near the boundary.

Finally, when λ is uniformly bounded away from $(\lambda_c^{(k)})_{k=0}^\infty$ (i.e., $d(\lambda_n) \geq \varepsilon > 0$) the $O(1)$ bound on $\text{gap}_{n,\lambda_n}^{-1}$ follows similarly to [12]; first, we use that in 2D systems, weak spatial mixing implies strong spatial mixing [40], then we use that strong spatial mixing implies an $O(1)$ inverse gap. Note that the dependence of this bound on $1/d(\lambda_n)$ blows up in a doubly exponential manner whereas the upper bounds of Item 2 of Theorem 1.1 and Item 2 of Theorem 1.3 are only exponential.

1.2. Open problems.

The critical window under zero boundary conditions. Consider the dynamical phase transition about $\lambda_c^{(k)}$ for some fixed $k \geq 0$. Our main results established the full dynamical phase transition in $\log \text{gap}_{n,\lambda_n}^{-1}$ (up to a poly-log(n) factor) in the torus (Theorem 1.3). However, as noted above, the case of zero boundary conditions is significantly more involved, and our results in Theorem 1.1 are only sharp for λ below $\lambda_c^{(k)}$; we expect that the difference between the two settings would be apparent already when viewing $\lambda = \lambda_c^{(k)}$, where the inverse-gap under zero boundary conditions should no longer be $\exp[\Theta(n)]$, but instead have order n^2 (the conjectured [25] order of the inverse-gap in Ising under all-plus boundary conditions). More precisely, we would expect that as $\lambda_n \uparrow \lambda_c^{(k)}$, the regime where the inverse-gap is $\exp[\Theta(n)]$ —which is reached as λ_n draws within distance $O(1/n)$ from $\lambda_c^{(k)}$ — would terminate discontinuously at $\lambda_n = \lambda_c^{(k)} - c_*/n$ for some $c_*(k, \beta)$, reminiscent of the critical point appearing in the work of Schonmann and Shlosman [45] on Glauber dynamics for the 2D Ising model under an external field. At this λ_n , and as λ_n further increases up

to $\lambda_n = \lambda_c^{(k)} + O(1/n)$, the inverse-gap should be polynomial in n , and thereafter should smoothly drop to $O(1)$ as $\lambda_n - \lambda_c^{(k)}$ increases to a constant order (see a possible depiction of this behavior in Fig. 3). A first step towards verifying this conjectured behavior would be to establish sub-exponential mixing at $\lambda = \lambda_c^{(k)}$.

Ferrari–Spohn limit in critical prewetting. Recall that the SOS model above a floor with $\lambda = 0$ is known to have a typical height of order $\log n$; it is also known [8–10] to be typically rigid about a single level, where the corresponding level-line contour has cube-root fluctuations around its deterministic scaling limit. In particular, as said limit coincides with the sides of the box near their centers, it is believed that in the setting of Eq. (1.1) with $\lambda = 0$, taking the intersection of the top level-line with an $n^{2/3} \times n^{2/3}$ box R_n about the center of the bottom boundary of Λ_n , and then rescaling its width by $n^{2/3}$ and its height by $n^{1/3}$, should give a nontrivial random process in the $n \rightarrow \infty$ limit. Unfortunately, the attraction/repulsion effects between the $\log n$ level lines, as well as against the boundary of Λ_n , pose a formidable challenge to the analysis of the model. So far it is only known that, in an $n^{2/3+\varepsilon} \times n^{2/3+\varepsilon}$ box R'_n about the center of the bottom boundary, the maximum vertical distance of the top level line from $\partial\Lambda_n$ is at least $n^{1/3-\varepsilon'}$ and at most $n^{1/3+\varepsilon'}$ (the aforementioned cube-root fluctuations). Continuous (1 + 1)D models approximating this ensemble of SOS level lines in R_n have been studied in detail, notably the ensemble of non-crossing Brownian excursions with an area tilt [5–7, 14], where the area A_k under the k -th curve tilts the probability by a factor of $\exp(-\mathfrak{b}^k A_k)$ for some $\mathfrak{b} > 1$ (in the SOS model with $\lambda = 0$, the entropic repulsion reward becomes stronger at the level k decreases, corresponding to a factor $\mathfrak{b} = e^{4\beta}$). Note that a single Brownian excursion with an area tilt has the law of a Ferarri–Spohn [18] diffusion. In the discrete setting, the longstanding prewetting problem in the 2D Ising model (akin to the SOS problem studied here, yet involving only a single level line) was very recently settled: the cube-root fluctuations of $n^{1/3+o(1)}$ at the critical scale $\lambda_n = \Theta(1/n)$ were established by Velenik [47] in 2004; Ganguly and the first author [22] established tightness of the fluctuations at scale $n^{1/3}$; and the Ferrari–Spohn limit law was finally established by Ioffe et al. [26].

In light of these results, and given the new results from this work, it would be interesting to examine the microscopic fluctuations of the top level line of the SOS model at $\lambda_n = \lambda_c^{(k)} - \Theta(1/n)$, corresponding to critical prewetting between levels k and $k + 1$. There, despite the existence of $k + 1$ interacting contours, one would expect the behavior to be effectively dictated by the top contour alone (as this value of λ is bounded away from the critical points corresponding to lower level lines), making it more amenable to analysis.

Question 1.5. Fix $\beta > \beta_0$, an integer $k \geq 0$ and some constant $c > 0$, and consider the (2 + 1)D SOS measure μ_{n,λ_n} from Eq. (1.1) for $\lambda_n = \lambda_c^{(k)} - c/n$. Let R_n be the square whose bottom edge overlaps the bottom boundary $\partial\Lambda_n$ along $\llbracket -n^{2/3}, n^{2/3} \rrbracket$, take the $(k + 1)$ -level line of SOS restricted to R_n , and rescale its width by $n^{2/3}$ and height by $n^{1/3}$. Does its law converge to a Ferrari–Spohn diffusion on $[-1, 1]$?

A variant of this question, which is somewhat simpler yet should still exhibit a Ferarri–Spohn limit, is to consider μ_{n,λ_n} under boundary conditions at height k (as opposed to zero) for $\lambda_n = \lambda_c^{(k)} - c/n$. That setting should be essentially the same as the case of $\lambda_c^{(0)}$, where there is typically at most 1 macroscopic level line. A useful first step towards these questions would be to understand the microscopic features of the k -to- $(k + 1)$ contours at $\lambda_c^{(k)}$, e.g., á la the tools developed for Ising interfaces in [16] and [42].

SOS near the roughening point and 3D Ising. We conclude with two challenging fronts. First, in light of the recently obtained [32] sharpness of the roughening transition for the SOS model, one would want to extend the above results to all $\beta > \beta_R$. For β so close to the roughening point, not only do cluster expansion techniques (used in this paper) become non-applicable, so do the Peierls maps on contours (which were the starting point of the entropic repulsion analysis of [9] for instance). Second, given recent progress on understanding the entropic repulsion phenomenon for the Ising interface [23, 24], one would hope to establish a similar metastability induced by layering for the Ising Glauber dynamics. To that end, a prerequisite would be to pinpoint the typical height of the interface in the prewetting problem for 3D Ising.

2. NOTATION AND PRELIMINARIES

In this section, we introduce much of the notation we will use throughout the paper. We also introduce the basic contour representation of SOS configurations, and include a few basic but important estimates that follow from the monotonicity of the model in the external field.

2.1. Notational disclaimers. Throughout the paper, we imagine β_0 to be a large absolute constant, so that $\beta > \beta_0$ indicates β sufficiently large (independent of everything else). We are only interested in $\lambda \leq 1$, say, since for β large, even the first critical point $\lambda_c^{(0)}$ will be less than 1, and monotonicity arguments can easily reduce analysis of larger λ to $\lambda = 1$. Finally, we use C in various place to indicate an absolute constant, and ε_β a sequence going to zero as $\beta \rightarrow \infty$, both of which may change from line to line.

2.2. Underlying geometry. The underlying geometry we work with are subsets of \mathbb{Z}^2 , with edge sets given by subsets of nearest neighbor edges $E(\mathbb{Z}^2)$. For a domain $V \subset \mathbb{Z}^2$, we use $E(V)$ to denote the edges in $E(\mathbb{Z}^2)$ between two vertices in V . When these subsets are squares, we use the notation $\Lambda_n = \llbracket -\frac{n}{2}, \frac{n}{2} \rrbracket^2$ with corresponding edge-set $[-\frac{n}{2}, \frac{n}{2}]^2 \cap E(\mathbb{Z}^2)$. For a general domain $V \subset \mathbb{Z}^2$, let

$$\partial_i V = \{x \in V : x \sim \mathbb{Z}^2 \setminus V\} \quad \text{and} \quad \partial_o V = \{x \in \mathbb{Z}^2 \setminus V : x \sim V\},$$

where the subscripts i, o are meant to indicate inner and outer boundary. The edge boundary is

$$\partial_e V = \{e = vw : v \in \partial_i V, w \in \partial_o V\},$$

and we can also denote the strict interior of a set V by $\overset{\circ}{V} = V \setminus \partial_i V$.

2.3. The SOS measure. We define the SOS model on general domains, with generic floors (and ceilings) and describe some important properties of the model.

Consider a finite domain $V \subset \mathbb{Z}^2$. A boundary condition is an assignment of heights to the vertices of $\partial_o V$. A general SOS distribution that we consider may also have floors and ceilings that vary depending on the vertex; for that, we let $(\mathbf{a}, \mathbf{b}) = ((a_v)_{v \in V}, (b_v)_{v \in V})$ be such that $a_v < b_v$ for all v , and we call a_v the floor and b_v the ceiling for a vertex.

For an inverse temperature $\beta > 0$, external field $\lambda \geq 0$, the SOS distribution on V with boundary conditions $\phi \in \mathbb{Z}^{\partial_o V}$ and ceiling/floor collection (\mathbf{a}, \mathbf{b}) is the distribution

$$\mu_{\lambda, \beta, V}^{\mathbf{a}, \mathbf{b}, \phi}(\varphi) = \frac{1}{Z_{\lambda, \beta, V}^{\mathbf{a}, \mathbf{b}, \phi}} e^{-H(\varphi)} \prod_{v \in V} \mathbf{1}_{a_v \leq \varphi_v \leq b_v}, \quad (2.1)$$

where $Z_{\lambda, \beta, V}^{\mathbf{a}, \mathbf{b}, \phi}$ is a normalizing factor to make it a probability measure, and

$$H(\varphi) = \beta \sum_{v \sim w} |\varphi_v - \varphi_w| + \lambda \sum_v \varphi_v. \quad (2.2)$$

where in the presence of boundary conditions ϕ , the first sum includes a sum over all edges between V and $\mathbb{Z}^2 \setminus V$ in which case φ_w is taken to be ϕ_w . If we write $\mu[X]$ for a random variable X , we mean the expectation under the event μ .

In order to simplify notation, we will tend to drop \mathbf{a}, \mathbf{b} and λ, β from the subscripts and superscripts of the measure to write $\mu_{\phi, V}$. We retain the relevant sub/superscripts when they are varying, but typically these will be fixed throughout a proof. The baseline floor/ceiling choices to have in mind will be $a_v \equiv 0$ and $b_v \equiv \infty$ for all $v \in V$, and we comment soon on other ones the reader might encounter. At the above level of generality, the SOS model satisfies a few essential properties that we recall.

2.3.1. *Monotonicity and the FKG property.* Endow the space of SOS configurations with the natural partial order (induced by pointwise total orders). The above monotonicity statement in boundary conditions, floors/ceilings, and field, is standard; we include a proof for completeness in Appendix A.

Lemma 2.1. *Fix any $\beta > 0$. Suppose that $\lambda' \leq \lambda$, $a_v \leq a'_v$, $b_v \leq b'_v$ and $\phi \leq \phi'$. Then*

$$\mu_{\lambda, V}^{\mathbf{a}, \mathbf{b}, \phi} \leq \mu_{\lambda', V}^{\mathbf{a}', \mathbf{b}', \phi'}.$$

where here and throughout the paper, \leq denotes stochastic domination.

2.3.2. *Domain Markov property.* The next essential property satisfied by the model is the domain Markov property, that describes the law of the SOS measure conditional on its values on some subset of V , as a new SOS measure with some other boundary conditions. Namely, if we fix a subset $W \subset V$, and let ζ be a configuration on W (consistent with the floors and ceilings \mathbf{a}, \mathbf{b}), then

$$\mu_V^\phi(\varphi(V \setminus W) \in \cdot \mid \varphi(W) = \zeta) \stackrel{d}{=} \mu_{V \setminus W}^{(\phi, \zeta)},$$

where the boundary conditions (ϕ, ζ) are the concatenation of η, ζ on $\partial_o(V \setminus W)$.

We end with an observation that if boundary conditions are above/below the floor/ceiling they can be treated as being exactly at the height of the floor/ceiling. Namely, the distribution $\mu_V^{\mathbf{a}, \mathbf{b}, \phi}$ is identical to $\mu_V^{\mathbf{a}, \mathbf{b}, \phi'}$ where for each vertex $w \in \partial_o V$, the boundary condition

$$\phi'_w = \begin{cases} \max_{v \sim w} b_v & \phi_w \geq \max_{v \sim w} b_v \\ \min_{v \sim w} a_v & \phi_w \leq \min_{v \sim w} a_v \\ \phi_w & \text{else} \end{cases}. \quad (2.3)$$

It thus suffices to always consider boundary conditions bounded between $\min_{v \in \partial_i V} a_v$ and $\max_{v \in \partial_i V} b_v$.

2.4. **A few basic estimates.** We present here a few basic results that will be helpful at various points in the paper. The first will be a bound on the maximum that is derived from comparison to the $\lambda = 0$ SOS measure, and crude bounds on the maximum therein. (Much more refined understanding of the maxima was obtained in [9]). For all $k \geq \frac{1}{4} \log |V|$, all β large,

$$\mu_V^{\mathbf{a}, \mathbf{b}, \phi}(\max_{v \in V} \varphi_v \geq \max\{\|\mathbf{a}\|_\infty, \|\phi\|_\infty\} + k) \leq e^{-\beta k}. \quad (2.4)$$

Proof. By monotonicity, it suffices to upper bound the probability that $\max_{v \in V} \varphi_v \geq k/2$ under the measure with $\lambda = 0$, no ceiling, and with boundary conditions that are identically $\max\{\|\mathbf{a}\|_\infty, \|\phi\|_\infty\} + k/2$. By a union bound over the $|V|$ vertices, the probability $Ce^{-4\beta(k/2)}$ of a downward deviation by $k/2$ (this follows from a Peierls argument so long as $\beta \geq 1$, say, per [9, Proposition 3.9]), this measure is coupled to the one with no floor at 0, except with probability $C|V|e^{-2\beta k}$. By a similar union bound, that probability that the no-floor sample has a maximum that goes up from the boundary conditions by an additional $k/2$ is at most $C|V|e^{-4\beta(k/2)}$. Together with the fact that $k \geq \log |V|$, we obtain Eq. (2.4). ■

Various of our estimates in the first few sections of the paper will go through comparison and manipulation of partition functions to obtain convergent cluster expansions corresponding to the SOS model with boundary conditions that are at the ‘‘correct’’ height. The next two bounds are simple a priori bounds on ratios of partition functions with different (constant) boundary conditions.

Lemma 2.2. *For every V and every $k < m$, we have*

$$\frac{Z_V^{\mathbf{a}, \mathbf{b}, k}}{Z_V^{\mathbf{a}', \mathbf{b}', m}} \leq e^{\lambda(m-k)|V|}.$$

so long as $[a_v + m - k, b_v + m - k] \subset [a'_v, b'_v]$ for all v .

Proof. Every configuration contributing to $Z_V^{\mathbf{a}, \mathbf{b}, k}$ can be mapped to one in $Z_V^{\mathbf{a}', \mathbf{b}', m}$ by shifting it by $m - k$. The change in the weights from this map is evidently exactly $e^{\lambda(m-k)|V|}$. ■

The following shows a non-quantitative monotonicity of ratios of partition functions (the top corresponding to a lower measure than the bottom) as one increases the external field.

Lemma 2.3. *For every V , every λ , every $\phi \leq \phi'$, and any floor/ceiling pair $\mathbf{a}, \mathbf{b} \leq \mathbf{a}', \mathbf{b}'$,*

$$\frac{d}{d\lambda} \frac{Z_{\lambda, V}^{\mathbf{a}, \mathbf{b}, \phi}}{Z_{\lambda, V}^{\mathbf{a}', \mathbf{b}', \phi'}} \geq 0.$$

Proof. Since the derivative of $\log(x)$ is positive on \mathbb{R}_+ , it suffices to show the claimed non-negativity of

$$\frac{d}{d\lambda} \log \left(\frac{Z_{\lambda, V}^{\mathbf{a}, \mathbf{b}, \phi}}{Z_{\lambda, V}^{\mathbf{a}', \mathbf{b}', \phi'}} \right) = \mu_{\lambda, V}^{\mathbf{a}', \mathbf{b}', \phi'} \left[\sum_v \varphi_v \right] - \mu_{\lambda, V}^{\mathbf{a}, \mathbf{b}, \phi} \left[\sum_v \varphi_v \right].$$

By the stochastic domination from Lemma 2.1 and the fact that $\sum_v \varphi_v$ is an increasing function, this difference is non-negative as claimed. ■

2.5. The common choices of boundary conditions and floors/ceilings. For much of the paper, we will be focused on a specific class of boundary conditions and floors/ceilings. The baseline choice of floors and ceilings in this paper is $a_v \equiv 0$ and $b_v \equiv \infty$ except when specified otherwise. The special class of boundary conditions are those which are constant $\phi \equiv k$ for some $k \geq 0$. The special class of floors/ceilings will modify the baseline choice along $\partial_i V$ via a *boundary signing* $\eta \in \{+, -, \emptyset, f\}^{\partial_i V}$.

Definition 2.4. If the boundary condition ϕ is constant, $\phi \equiv k$, say, then the *boundary signing* indicates the choices of floors/ceilings where for each $v \in \partial_i V$,

- if $\eta_v = +$ then $a_v = k$ and if $\eta_v = -$, then $b_v = k$
- if $\eta_v = \emptyset$, then $a_v = b_v = k$
- if $\eta_v = f$, then a_v, b_v are the default choices of 0 and ∞ respectively.

Until Section 5, constant ϕ and floors/ceilings given by $a_v \equiv 0, b_v \equiv \infty$ and boundary signings will be the exclusive family of boundary conditions and floors/ceilings we work with and therefore we give it a shortened notation. Define $\mu_{\eta, k, V}$ as the SOS measure with boundary conditions that are identically k , and boundary signing η modifying $a_v \equiv 0, b_v = \infty$ along $\partial_i V$ according to Definition 2.4. Let $Z_{\eta, k, V}$ be the corresponding partition function. In this special context, Lemma 2.2 translates to the following.

Corollary 2.5. *For every V and every $k < m$ and every boundary signing η ,*

$$\frac{Z_{\eta, k, V}}{Z_{\eta, m, V}} \leq e^{\lambda(m-k)|V|}.$$

2.6. The contour representation. An important tool in the analysis of the SOS model is the *contour representation*, essentially a bijection between configurations and a family of non-crossing loops representing level curves of the SOS surface. The contours will be formed out of *dual edges*; for that, let V^* denote the planar dual graph of V generated by considering $(\mathbb{Z}^2)^* = \mathbb{Z}^2 + (\frac{1}{2}, \frac{1}{2})$, then taking $E(V^*)$ to consist of all edges of $E((\mathbb{Z}^2)^*)$ bisecting the edge set $\{e = vv' : v \in V\}$. The vertex set V^* is then naturally induced.

We will use the contour representation of SOS configurations, referring the reader to [9, Section 3] for more background on the below.

Definition 2.6. A geometric contour γ is a connected set of dual-edges in $E(V^*)$ whose elements can be sequenced $e_0, \dots, e_{|\gamma|}$ in such a way that

- $e_i \neq e_j$ if $i \neq j$ except if $i = 0$ and $j = n$;
- e_i and e_{i+1} are incident for every i ;
- if $e_i, e_{i+1}, e_j, e_{j+1}$ all share a single dual vertex, then e_i, e_{i+1} are either the south and west pair of dual edges, or the north and east pair.

The last item above is a canonical splitting rule, where the pairings SW and NE are simply a convention. With this convention, we will call sets $V \subset \mathbb{Z}^2$ simply-connected if there is a geometric contour, whose edge set is exactly $(\partial_e V)^*$, and we will use the notation γ_V to denote this contour.

Definition 2.7. A vertex is exterior to a contour γ if it lies in the infinite connected component of $\mathbb{R}^2 \setminus \gamma$. When there is an underlying domain $V \subset \mathbb{Z}^2$, we use $\text{Ext}(\gamma)$ to denote those vertices of V exterior to γ . A vertex is interior to a contour γ if it lies in a finite connected component of $\mathbb{R}^2 \setminus \gamma$, and $\text{Int}(\gamma)$ denotes the set of such vertices. We say γ nests γ' if $\text{Int}(\gamma') \subset \text{Int}(\gamma)$, and two contours γ, γ' are mutually external if neither nests the other.

We are now able to define up and down contours corresponding to an SOS configuration φ .

Definition 2.8. In an SOS configuration φ , a geometric contour γ is an up- h -contour if $\varphi_x \geq h$ for all $h \in \partial_i \text{Int}(\gamma)$ and $\varphi_x < h$ for all $h \in \partial_o \text{Int}(\gamma)$. Analogously, a geometric contour γ is a down- h -contour if $\varphi_x \leq h$ for all $h \in \partial_i \text{Int}(\gamma)$ and $\varphi_x > h$ for all $h \in \partial_o \text{Int}(\gamma)$. In these cases, use the notation $S(\gamma) = \pm$ for whether it is up or down, and $h(\gamma) = h$.

If we consider SOS configurations with constant boundary $\phi \equiv k$, we can associate to each configuration a contour representation, where the contours it consists of are all signed contours at specific heights. Namely, for a configuration φ , with boundary conditions k , on a domain V , first include the outermost up- $(k+1)$ -contours and down- $(k-1)$ -contours. Interior to each of the outermost up- $(k+1)$ -contours, we can consider the outermost up- $(k+2)$ and down- k -contours, and proceed iteratively in this manner to get a family of signed contours, each having an associated height. Per this construction, if multiple contours are the same geometric contour, just at different heights, the one whose height is smallest if they are up-contours, or largest if they are down-contours, is called the outermost one.

Definition 2.9. A family of contours Γ is said to be admissible if there exists a φ whose contour collection is exactly Γ . (This effectively imposes that contours may not cross, and if they overlap on an edge, their signs must be the same if one nests the other, while their signs must be different if they are mutually external.)

With this, we can express the SOS measure in terms of weights associated to contours.

Definition 2.10. The weight of a contour γ is

$$W(\gamma) = \begin{cases} e^{-\beta|\gamma| - \lambda|\text{Int}(\gamma)|} & \gamma \text{ is an up-contour} \\ e^{-\beta|\gamma| + \lambda|\text{Int}(\gamma)|} & \gamma \text{ is a down-contour} \end{cases}.$$

There is a bijection from height functions that have a constant boundary condition $\phi \equiv h$ outside a set $V \subset \mathbb{Z}^2$ (and no floor/ceiling) and admissible contour collections in V^* . In fact, the information of the heights of the contours is not needed to reconstruct φ , only their signs are needed (the heights can be reconstructed from the boundary inwards). In our contexts, we require contour collections to be compatible with the floor/ceilings. In the default case we are interested in a floor at height zero, and a boundary signing η per Definition 2.4.

Denote by $\mathcal{G}_{\eta, h, V}$ the set of all admissible contours collections whose corresponding configuration is in $\Omega_{\eta, h, V}$. Notice that the boundary signing is imposing a constraint on the contours that no dual edge bisecting $\partial_e V \cap \partial_e \{v\}$ is part of an up contour if $\eta_v \in \{-, \emptyset\}$ (respectively is part of a down contour if $\eta_v \in \{+, \emptyset\}$).

With that, we can rewrite the corresponding partition function in terms of contour weights as:

$$Z_{\eta, h, V} := e^{-\lambda h |V|} \sum_{\Gamma \in \mathcal{G}_{\eta, h, V}} \prod_{\gamma \in \Gamma} W(\gamma). \quad (2.5)$$

The specific choices where $\eta \equiv +$ and $\eta \equiv -$ will be especially pertinent since, e.g., $Z_{+, h(\gamma), \text{Int}(\gamma)}$ describes the possible contour collections found in $\text{Int}(\gamma)$ if γ is an up- h -contour.

3. RENORMALIZED WEIGHTS, ELEMENTARY CONTOURS AND FREE ENERGIES

We study a family of contours which are exponentially unlikely (in their length) no matter what height they are based at. These are small enough that neither the relative benefit of the external field (for down-contours) nor the entropic repulsion effect (for up-contours) are enough to overcome the exponential in β cost in the length of the contour. Following the language of [11, 15], these contours are called *elementary*. We note that the overarching goal of this section, Theorem 3.6, tracks [15] but for contours instead of the *cylinders*—essentially maximal stacks of identical contours—that were considered in that paper. This simplifies some aspects of the proof, while complicating others; in particular, our proof requires stitching together Peierls-type arguments in the standard contour representation of the SOS model, with arguments for using the “renormalized” contour representation defined as follows.

3.1. Renormalized contour weights. The first step we take is to rewrite SOS partition functions in terms of renormalized contour weights according to a specified boundary height h . The role of this renormalization is to move to a model over contour collections with all the contours based at height h , rather than each having their own associated height; the outermost contours of this model will still correspond to the outermost contours of the original model.

Definition 3.1. For a contour γ , the h -renormalized weight is defined as

$$W_h^{\text{rn}}(\gamma) = \begin{cases} e^{-\beta|\gamma| \frac{Z_{+,h+1,\text{Int}(\gamma)}}{Z_{+,h,\text{Int}(\gamma)}}} & \gamma \text{ is an up-contour} \\ e^{-\beta|\gamma| \frac{Z_{-,h-1,\text{Int}(\gamma)}}{Z_{-,h,\text{Int}(\gamma)}}} & \gamma \text{ is a down-contour} \end{cases}.$$

Given these renormalized weights, we can obtain the following equivalence between the partition function on a simply connected domain V and without tracking the heights of the corresponding contours, allowing arbitrary compatible families. Towards that purpose, define $\mathcal{G}_{\eta,h,V}^{\text{rn}}$ as admissible collections of signed contours compatible with η , but with no constraint on the height function they generate respecting a floor at height 0 (unless $h = 0$, in which case no down contours are allowed). With that, we can define

$$Z_{\eta,h,V}^{\text{rn}} = e^{-\lambda h|V|} \sum_{\Gamma \in \mathcal{G}_{\eta,h,V}^{\text{rn}}} \prod_{\gamma \in \Gamma} W_h^{\text{rn}}(\gamma). \quad (3.1)$$

In an element $\Gamma \in \mathcal{G}_{\eta,h,V}^{\text{rn}}$, we say a contour $\gamma \in \Gamma$ is *outermost* if it is not strictly nested by any other contour in Γ . Since contours may overlay fully and in the renormalized collection are all based at the same height (i.e., multiple copies of the same contour can be in Γ), we arbitrarily call one of them the outermost one.

Lemma 3.2. *The standard partition function and the renormalized one are equal, namely for any simply-connected set V ,*

$$Z_{\eta,h,V} = Z_{\eta,h,V}^{\text{rn}}.$$

Moreover, the same equality holds if we specify some outermost contours which must belong to a configuration and only sum over Γ having those as outermost contours.

Proof. If we let Γ_{out} be collections of outermost contours, then by Eq. (2.5) and Definition 3.1,

$$e^{\lambda h|V|} Z_{\eta,h,V} = \sum_{\Gamma_{\text{out}}} \prod_{\gamma \in \Gamma_{\text{out}}} W(\gamma) e^{\lambda h|\text{Int}(\gamma)|} Z_{S(\gamma),h\pm 1,\text{Int}(\gamma)} = \sum_{\Gamma_{\text{out}}} \prod_{\gamma \in \Gamma_{\text{out}}} W_h^{\text{rn}}(\gamma) e^{\lambda h|\text{Int}(\gamma)|} Z_{S(\gamma),h,\text{Int}(\gamma)},$$

where in the middle display, $h + 1$ is used when $S(\gamma) = +$ and $h - 1$ when $S(\gamma) = -$. Performing the same expansion on $e^{\lambda h|\text{Int}(\gamma)|} Z_{S(\gamma),h,\text{Int}(\gamma)}$ and repeating, this gives the claimed equality to $e^{\lambda h|V|} Z_{\eta,h,V}^{\text{rn}}$ concluding the proof. It is evident that the same argument could be performed restricted to sums over Γ_{out} that contain a fixed collection of outermost contours. \blacksquare

The benefit of the above renormalized contour weight representation is that the deletion of a contour from Γ leaves an element of the same $\mathcal{G}_{\eta,h,V}^{\text{rn}}$, whereas in $\mathcal{G}_{\eta,h,V}$ such a deletion could cause conflicts with the

floor at height 0. This enables the use of simple Peierls maps on the renormalized collection so long as the renormalized weights are exponentially small in $|\gamma|$.

3.2. Elementary contours. We next consider *elementary contours*, which are those we can show to have exponentially small renormalized weight regardless of the height h with respect to which they are renormalized. For a subset A in \mathbb{R}^2 , we use $\text{diam}(A)$ to denote $\max\{\|x - y\|_\infty : x, y \in A\}$.

Definition 3.3. A contour γ is called h -elementary if it has

$$\text{diam}(\gamma) \leq \lambda^{-1} \wedge e^{3\beta(h+1)}.$$

Remark 3.4. Note the following two important properties:

- if γ is h -elementary, then it is also $(h + 1)$ -elementary.
- if γ' is nested in γ and γ is h -elementary, then γ' is also h -elementary.

We now introduce the notation $\mathcal{G}_{\eta,h,V}^{\text{el}}$ and the notation $\mathcal{G}_{\eta,h,V}^{\text{rn.el}}$, for compatible (renormalized) collections of contours all of which are h -elementary. With these notations, we can define

$$Z_{\eta,h,V}^{\text{el}} = e^{-\lambda h|V|} \sum_{\Gamma \in \mathcal{G}_{\eta,h,V}^{\text{el}}} \prod_{\gamma \in \Gamma} W(\gamma), \quad \text{and} \quad Z_{\eta,h,V}^{\text{rn.el}} = e^{-\lambda h|V|} \sum_{\Gamma \in \mathcal{G}_{\eta,h,V}^{\text{rn.el}}} \prod_{\gamma \in \Gamma} W_h^{\text{rn}}(\gamma). \quad (3.2)$$

By the same reasoning as in the proof of Lemma 3.2, we obtain the following equivalence.

Lemma 3.5. *The elementary partition function and the renormalized elementary one are equal, namely for a simply-connected set V*

$$Z_{\eta,h,V}^{\text{el}} = Z_{\eta,h,V}^{\text{rn.el}}.$$

Moreover, the same equality holds if we specify some outermost contours which must belong to a configuration, and only sum over Γ having those as outermost contours.

In what follows, we use the notation $\mu_{\eta,h,V}^{\text{el}}$ to denote the distribution over $\mathcal{G}_{\eta,h,V}^{\text{el}}$ which assigns each Γ , the weight $\prod_{\gamma \in \Gamma} W(\gamma)$. Equivalently, this is the SOS measure $\mu_{\eta,h,V}$ conditioned on its contour representation belonging to $\mathcal{G}_{\eta,h,V}^{\text{el}}$, i.e., conditioned on all its outermost contours being h -elementary.

The following theorem is our main result in this section, and establishes that the h -elementary contours always have exponentially small h -renormalized weight. It also gives a first bound, that is sharp up to the constant in the exponent, for the ratio of elementary partition functions at different heights. This is the analogue, in the context of our contours, of [15, Lemma 2.5] and [11, Lemma 2.7].

Theorem 3.6. *There exists β_0 such that for every simply-connected $V \subset \mathbb{Z}^2$ and every $\beta > \beta_0$,*

(1) *For every $0 \leq k < m$ and every boundary signing η ,*

$$e^{-2|V|e^{-(4\beta-\lambda)(k+1)}} \leq e^{-\lambda(m-k)|V|} \frac{Z_{\eta,k,V}^{\text{el}}}{Z_{\eta,m,V}^{\text{el}}} \leq 2e^{-\frac{1}{2}|\mathring{V}|e^{-(4\beta-\lambda)(k+1)}}. \quad (3.3)$$

(2) *For every $h \geq 0$, if γ is h -elementary,*

$$W_h^{\text{rn}}(\gamma) \leq e^{-(\beta-2)|\gamma|}.$$

Proof. We prove this by induction, showing that Item 1 implies Item 2 holds for that same V , and then by showing that if Item 2 holds for V for all connected sets V having $|V| \leq l$, then that implies Item 1 for all sets V' having $|V'| \leq l + 1$. This second step is the core of the argument and entails maps on elementary contour families to estimate the entropic repulsion effect amongst elementary contours, at a given height k .

Base case for Item 1. Consider $V = \{x\}$, whence $\mathring{V} = \emptyset$. In this case, the only contours in $E(V^*)$ are the loops consisting of the four dual-edges surrounding x , i.e., $\gamma_{\{x\}}$ (notice that these contours are always elementary, no matter its height or sign, since $\text{diam}(\gamma_{\{x\}}) = 1$ is less than $\lambda^{-1} \wedge e^{3\beta}$).

If $\eta = \emptyset$ then $Z_{\emptyset,k,V}^{\text{el}} = e^{-\lambda k}$ and the middle quantity in Eq. (3.3) is exactly 1, which satisfies the two bounds since $|\dot{V}| = 0$. If $\eta = +$ then

$$Z_{+,k,V}^{\text{el}} = e^{-\lambda k} \sum_{l \geq 0} e^{-4\beta l - \lambda l}, \quad \text{so that} \quad e^{-\lambda(m-k)} \frac{Z_{+,k,V}^{\text{el}}}{Z_{+,m,V}^{\text{el}}} = \frac{\sum_{l \geq 0} e^{-4\beta l - \lambda l}}{\sum_{l \geq 0} e^{-4\beta l - \lambda l}} = 1,$$

which again satisfies the two bounds in Eq. (3.3). If $\eta = -$, then

$$Z_{-,k,V}^{\text{el}} = e^{-\lambda k} \sum_{l \geq 0} e^{-4\beta l + \lambda l}, \quad \text{so that} \quad e^{-\lambda(m-k)} \frac{Z_{-,k,V}^{\text{el}}}{Z_{-,m,V}^{\text{el}}} = \frac{\sum_{0 \leq l \leq k} e^{-4\beta l + \lambda l}}{\sum_{0 \leq l \leq m} e^{-4\beta l + \lambda l}}.$$

Since $\lambda \leq 1$ and $\beta > \beta_0$, we have $e^{-4\beta + \lambda} < 1$ and therefore the middle quantity in Eq. (3.3) is equal to

$$\frac{1 - e^{(-4\beta + \lambda)(k+1)}}{1 - e^{(-4\beta + \lambda)(m+1)}}.$$

This is at most 1 since $m > k$. As for the lower bound, the denominator is at most 1 (as $k < m$ and $\lambda < 4\beta$), and for the numerator we use $1 - x \geq e^{-2x}$ for $x < \frac{1}{2}$, so it is at least $e^{-2e^{(-4\beta + \lambda)(k+1)}}$ as claimed. Finally if $\eta = f$, then

$$e^{-\lambda(m-k)} \frac{Z_{f,k,V}^{\text{el}}}{Z_{f,m,V}^{\text{el}}} = \frac{\sum_{l \geq -k} e^{-4\beta|l| - \lambda l}}{\sum_{l \geq -m} e^{-4\beta|l| - \lambda l}} = \frac{\sum_{l > 0} e^{-4\beta l - \lambda l} + \sum_{0 \leq l \leq k} e^{-4\beta l + \lambda l}}{\sum_{l > 0} e^{-4\beta l - \lambda l} + \sum_{0 \leq l \leq m} e^{-4\beta l + \lambda l}}.$$

The last fraction evidently lies between the ratio of the first sums and the ratio of the second sums. The first of these ratios is exactly 1 and the second of these ratios is exactly $Z_{-,k,V}^{\text{el}}/Z_{-,m,V}^{\text{el}}$; in particular, both satisfy the upper and lower bounds of Eq. (3.3), and thus so does the above.

Item 1 for $|V| \leq l$ implies *Item 2* for $\gamma : |\text{Int}(\gamma)| = l$. Suppose *Item 1* holds for all sets V having $|V| \leq l$. Fix any h and consider an up-contour γ with $\text{Int}(\gamma) = l$. Recalling Definition 3.1, we have

$$W_h^{\text{rn}}(\gamma) = e^{-\beta|\gamma|} \frac{Z_{+,h+1,\text{Int}(\gamma)}}{Z_{+,h,\text{Int}(\gamma)}} e^{-\beta|\gamma|} \frac{Z_{+,h+1,\text{Int}(\gamma)}^{\text{el}}}{Z_{+,h,\text{Int}(\gamma)}^{\text{el}}},$$

where we used that since γ is h -elementary (and by Remark 3.4 also $h+1$ -elementary), every γ' in $\text{Int}(\gamma)$ must also be both h and $h+1$ -elementary, so both partition functions are equal to their rn.el-versions, and those in turn are equal to the elementary ones since $\text{Int}(\gamma)$ is simply connected per Lemma 3.5. Thus, by the lower bound in Eq. (3.3) from *Item 1* for the domain $\text{Int}(\gamma)$, with $k = h$ and $m = h+1$,

$$W_h^{\text{rn}}(\gamma) \leq e^{-\beta|\gamma| - \lambda|\text{Int}(\gamma)| + 2|\text{Int}(\gamma)|} e^{-(4\beta - \lambda)(h+1)}.$$

Since γ is h -elementary, $|\text{Int}(\gamma)| \leq |\gamma| \text{diam}(\gamma) \leq |\gamma| e^{3\beta(h+1)}$, and thus we get the claimed bound on $W_h^{\text{rn}}(\gamma)$ at $\beta > \beta_0$. For a down-contour γ , by Corollary 2.5 with $k = h-1$ and $m = h$,

$$W_h^{\text{rn}}(\gamma) = e^{-\beta|\gamma|} \frac{Z_{-,h-1,\text{Int}(\gamma)}}{Z_{-,h,\text{Int}(\gamma)}} \leq e^{-\beta|\gamma| + \lambda|\text{Int}(\gamma)|}.$$

Since $|\text{Int}(\gamma)| \leq |\gamma| \text{diam}(\gamma) \leq \lambda^{-1}|\gamma|$ and $|\gamma| \geq 1$, this gives the claimed bound.

Item 2 for all $|V| < l$ implies *Item 1* for $|V| \leq l$. In what follows, suppose *Item 2* holds for all γ having $|\text{Int}(\gamma)| \leq l-1$. Let us begin with a few consequences of the inductive assumption of *Item 2* for $V : |V| \leq l-1$ that will be useful in both the upper and lower bound of *Item 2*.

The first is the easy observation that a bound on the renormalized weights implies a bound on the probability that under the renormalized contour model, a certain family of compatible k -contours is present.

Lemma 3.7. *Consider a fixed family G of compatible k -elementary contours, all of which have $|\text{Int}(\gamma)| \leq l - 1$. Let $Z_{\eta,k,V}^{\text{rn,el}}(G)$ be the analogue of $Z_{\eta,k,V}^{\text{rn,el}}$ with the sum only running over elements of $\mathcal{G}_{\eta,k,V}^{\text{rn,el}}$ containing all of G . Then for any V ,*

$$\frac{Z_{\eta,k,V}^{\text{rn,el}}(G)}{Z_{\eta,k,V}^{\text{rn,el}}} \leq e^{-(\beta-2)\sum_{\gamma \in G} |\gamma|}.$$

If the contours in G are mutually external and we only sum over $\mathcal{G}_{\eta,k,V}^{\text{rn,el}}$ in which they are outermost, the same bound applies to $Z_{\eta,k,V}^{\text{el}}(G)/Z_{\eta,k,V}^{\text{el}}$.

Proof. Consider the map that removes G from the numerator to send it into an element of the denominator. This is an injection, and the relative weight of each pre-image to its image is $\prod_{\gamma \in G} W_k^{\text{rn}}(\gamma)$. By the inductive assumption and the fact that every $\gamma \in G$ has interior smaller than $l - 1$, we get the claim. \blacksquare

Since our contours may overlay, when $|V| = l$, it may admit the contour $\gamma_V = (\partial_e V)^*$, which will have $|\text{Int}(\gamma_V)| = l$ and to which the inductive assumption will not apply. We use a 1-site modification to show therefore that if V has size l , then it is exponentially unlikely to have γ_V as a contour. For $j \in \mathbb{Z}$,

$$\mathcal{E}_{\eta,k,V}^j = \{\Gamma \in \mathcal{G}_{\eta,k,V} : \text{exactly } |j| \text{ copies of } \gamma_V \text{ with sign matching } \text{sgn}(j) \text{ in } \Gamma\}. \quad (3.4)$$

Lemma 3.8. *For every η, k and V such that $|V| \leq l$, we have for every j ,*

$$\mu_{\eta,k,V}^{\text{el}}(\mathcal{E}_{\eta,k,V}^j) \leq \exp(-(\beta - 2)|j||\gamma_V| + 5(1 \vee j)).$$

Proof. In order for the probability not to be zero, it must be that γ_V is k -elementary, and either $\eta_v = \{-, f\}$ for all $v \in \partial_i V$ if j negative, or $\eta_v \in \{+, f\}$ for all $v \in \partial_i V$ if j positive.

Consider first the probability that under $\mu_{\eta,k,V}^{\text{el}}$ there exists at least one contour in Γ that is all of γ_V . In the language of Lemma 3.2, this is equivalent to the fraction $\frac{Z_{\eta,k,V}^{\text{el}}(\{\gamma_V\})}{Z_{\eta,k,V}^{\text{el}}} = \frac{Z_{\eta,k,V}^{\text{rn,el}}(\{\gamma_V\})}{Z_{\eta,k,V}^{\text{rn,el}}}$; however, we cannot yet apply Lemma 3.7 because $|\text{Int}(\gamma_V)| = |V| = l$, not $l - 1$. Therefore, we first consider the modification $M_x \Gamma$ of $\mathcal{G}_{\eta,k,V}^{\text{rn,el}}(G)$ that performs an \oplus operation with a single copy of $\gamma_{\{x\}}$ satisfying $V \setminus \{x\}$ is s.c. (such an x must exist). For each $\Gamma \in \mathcal{G}_{\eta,k,V}^{\text{rn,el}}(\{\gamma_V\})$, let d_x ($|d_x| \leq 2$) be the change in the total number of edges in Γ vs. the $M_x \Gamma$. The resulting configuration evidently belongs to $\mathcal{G}_{\eta,k,V}^{\text{rn,el}}$. Its relative weight satisfies

$$\frac{\prod_{\gamma \in M_x \Gamma} W_k^{\text{rn}}(\gamma)}{\prod_{\gamma \in \Gamma} W_k^{\text{rn}}(\gamma)} \leq e^{\beta d_x + \lambda}.$$

Furthermore, M_x is an injection since the original configuration can be recovered by the \oplus operation with the singleton loop around $\{x\}$. Let

$$\tilde{Z} = \sum_{\Gamma \in \mathcal{G}_{\eta,k,V}^{\text{rn,el}}(\{\gamma_V\})} \prod_{\gamma \in M_x \Gamma} W_k^{\text{rn}}(\gamma) = \sum_{M_x \Gamma \in M_x \mathcal{G}_{\eta,k,V}^{\text{rn,el}}(\{\gamma_V\})} \prod_{\gamma \in M_x \Gamma} W_k^{\text{rn}}(\gamma).$$

Each $M_x \gamma$ appearing in this sum can in turn be mapped to one in which the contour $\tilde{\gamma}$ in $M_x \Gamma$ resulting from the modification to γ_V after $\oplus \gamma_{\{x\}}$, is deleted from the collection. That yields a further weight change $W_k^{\text{rn}}(\tilde{\gamma})$. Since $\tilde{\gamma}$ necessarily has $|\text{Int}(\tilde{\gamma})| \leq l - 1$, this is at most $e^{-(\beta-2)|\tilde{\gamma}|}$. Summing over the possible choices for the contour $\tilde{\gamma}$ (namely summing over possible contours γ' adjacent to x such that $\tilde{\gamma} = \gamma_V \oplus \gamma_{\{x\}} \oplus \gamma'$, the number of such choices being at most 4^r for $|\gamma'| = r$, say), we get

$$\frac{Z_{\eta,k,V}^{\text{el}}(\{\gamma_V\})}{Z_{\eta,k,V}^{\text{el}}} \leq \sum_{r \geq 0} 4^r e^{-(\beta-2)(|\gamma_V| + d_x + r) + \beta d_x + \lambda} \leq e^{-(\beta-2)|\gamma_V| + \lambda + 4}. \quad (3.5)$$

where the 4 comes from the worst-case value of d_x .

We now wish to reduce the event $\mathcal{E}_{\eta,k,V}^j$ to the event of having at least 1 such contour. If j is negative, consider the deletion of $|j| - 1$ such contours from the (standard) contour representation. That gives a change

of $e^{\beta(j-1)|\gamma_V| - \lambda(j-1)|V|}$ to the weight of the configurations, and is an injection from $\mathcal{E}_{\eta,k,V}^j$ to $\mathcal{E}_{\eta,k,V}^{-1}$. Since we assumed γ_V to be k -elementary, $\lambda|V| \leq \lambda \text{diam}(V)|\gamma_V| \leq |\gamma_V|$, and thus we have

$$\mu(\mathcal{E}_{\eta,k,V}^j) \leq e^{-(\beta-1)(j-1)|\gamma_V|} \mu(\mathcal{E}_{\eta,k,V}^{-1}) \leq e^{-(\beta-1)(j-1)|\gamma_V|} e^{-(\beta-2)|\gamma_V| + \lambda + 4}.$$

which easily the claimed bound using $\lambda \leq 1$.

If j is positive, notice that on the event of having $\geq j-1$ many up-contours γ_V , the probability of having at least j is at most the probability of having at least $j-1$, times the conditional probability of having at least one more given those $j-1$. But the conditional measure given at least $j-1$ up-contours γ_V is exactly $\mu_{+,k+j-1,V} = \mu_{+,k+j-1,V}^{\text{el}}$ because all possible contours in V are enclosed in γ_V , which is k -elementary and therefore also $k+j-1$ elementary for $j \geq 1$. For this conditional probability, the bound of Eq. (3.5) applies (with $k+j-1$ instead of k), and iterating that down, this gives

$$\mu_{\eta,k,V}^{\text{el}}(\mathcal{E}_{\eta,k,V}^j) \leq \prod_{i=1}^j e^{-(\beta-2)|\gamma_V| + \lambda + 4},$$

yielding the claim after $\lambda \leq 1$. ■

Together, Lemmas 3.7 and 3.8 imply that if V has $|V| = l$, the probability under $\mu_{\eta,k,V}^{\text{el}}$ of any fixed collection of outermost contours is exponentially small.

Corollary 3.9. *Let V have $|V| \leq l$. The probability of having a specific contour γ as an outermost contour under $\mu_{\eta,k,V}^{\text{el}}$, is at most $e^{-(\beta-3)|\gamma|}$.*

Proof. If $\gamma = \gamma_V$, this is ruled out by Lemma 3.8, a sum over j , and $|\gamma_V| > 4$. Otherwise, the probability of γ being an outermost contour is at most $e^{-(\beta-2)|\gamma|}$ per Lemma 3.7. ■

With the above in hand, we move to proving Item 1; we split up the proofs of the upper and lower bound.

Upper bound in Eq. (3.3). Intuitively, the proof of the upper bound goes by mapping configurations in $Z_{\eta,k,V}^{\text{el}}$ to $Z_{\eta,m,V}^{\text{el}}$ in such a way to inject entropy via singleton spikes $\gamma_{\{x\}}$ that go down by height m , which would have been prohibited under $Z_{\eta,k,V}^{\text{el}}$ due to the floor constraint. We first construct such a map, using the contour representation of the elementary SOS configurations. We then use its alternative representation with renormalized weights together with the inductive hypothesis to argue that under $Z_{\eta,k,V}^{\text{el}}$, most sites are at height k , and thus there is typically $(1 - \varepsilon_\beta)|V|$ many sites to insert these downward spikes.

We begin by working with $\mathcal{G}_{\eta,k,V}^{\text{el}}$, which we recall is the set of contour representations of SOS configurations compatible with the floor at height zero, and all of whose outermost contours are all k -elementary. This in fact implies that all its contours are k -elementary (though the inner contours may not be elementary at the height they are based at). For $\Gamma \in \mathcal{G}_{\eta,k,V}^{\text{el}}$, let

$$V_\Gamma = \{x \in \mathring{V} : \min_{\gamma \in \Gamma} d(x, \text{Int}(\gamma)) > 1\}.$$

Given any $\Gamma \in \mathcal{G}_{\eta,k,V}^{\text{el}}$ and a subset $S \subset V_\Gamma$, define the map T_S as the one that adds to Γ , collections of $k+1$ many copies of the down-contour $\gamma_{\{v\}}$.

Claim 3.10. *For every $k < m$, every $\Gamma \in \mathcal{G}_{\eta,k,V}^{\text{el}}$, and every $S \subset V_\Gamma$, we have $T_S \Gamma \in \mathcal{G}_{\eta,m,V}^{\text{el}}$.*

Proof. This resulting collection $T_S \Gamma$ is in $\mathcal{G}_{\eta,m,V}$ because the addition of these contours (1) is compatible with the remainder of the contour collection since $S \subset V_\Gamma$ and V_Γ ensures that any newly added contours will not be incident to any contours in Γ ; and (2) does not violate the floor constraint since $m \geq k+1$, and in Γ , the height of sites in V_Γ is k . Finally, it is compatible with the signing on the boundary condition since $S \subset V_\Gamma \subset \mathring{V}$. To see that $T_S \Gamma \in \mathcal{G}_{\eta,m,V}^{\text{el}}$, we further claim that in $T_S \Gamma$, every outermost contour is m -elementary. This is because the singleton contours are trivially m -elementary having diameter 1, and all the outermost contours in Γ were k -elementary so they are also m -elementary per Remark 3.4. ■

Each configuration in the set $(T_S \Gamma)_{S \subset V_\Gamma, \Gamma \in \mathcal{G}_{\eta, k, V}^{\text{el}}}$ has a unique pre-image in $\mathcal{G}_{\eta, k, V}^{\text{el}}$, since the preimage can be reconstructed by taking the corresponding SOS configuration, shifting all sites at heights $m - (k + 1)$ up to m , then shifting the whole configuration down by $m - k$. Thus, for any $\theta > 0$,

$$Z_{\eta, m, V}^{\text{el}} \geq e^{-\lambda m |V|} \sum_{\substack{\Gamma \in \mathcal{G}_{\eta, k, V}^{\text{el}} \\ |V_\Gamma| \geq \theta |\dot{V}|}} \sum_{S \subset V_\Gamma} \prod_{\gamma \in T_S \Gamma} W(\gamma).$$

Also, by definition of the map T_S , we have

$$\frac{\prod_{\gamma \in T_S \Gamma} W(\gamma)}{\prod_{\gamma \in \Gamma} W(\gamma)} \geq e^{-(4\beta - \lambda)(k+1)|S|}.$$

(This is an inequality rather than equality because some sites in S could be adjacent.) Thus, for every $\theta > 0$,

$$\begin{aligned} \frac{Z_{\eta, m, V}^{\text{el}}}{Z_{\eta, k, V}^{\text{el}}} &\geq e^{-\lambda m |V|} \sum_{\substack{\Gamma \in \mathcal{G}_{\eta, k, V}^{\text{el}} \\ |V_\Gamma| \geq \theta |\dot{V}|}} \sum_{S \subset V_\Gamma} \frac{\prod_{\gamma \in T_S(\Gamma)} W(\gamma)}{Z_{\eta, k, V}^{\text{el}}} \geq e^{-\lambda(m-k)|V|} \sum_{\substack{\Gamma \in \mathcal{G}_{\eta, k, V}^{\text{el}} \\ |V_\Gamma| \geq \theta |\dot{V}|}} \sum_{S \subset V_\Gamma} e^{-(4\beta - \lambda)(k+1)|S|} \mu_{\eta, k, V}^{\text{el}}(\Gamma) \\ &= e^{-\lambda(m-k)|V|} \sum_{\substack{\Gamma \in \mathcal{G}_{\eta, k, V}^{\text{el}} \\ |V_\Gamma| \geq \theta |\dot{V}|}} \mu_{\eta, k, V}^{\text{el}}(\Gamma) \left(1 + e^{-4\beta - \lambda}(k+1)\right)^{|V_\Gamma|} \\ &\geq \exp\left(-\lambda(m-k)|V| + \frac{1}{2}e^{-(4\beta - \lambda)(k+1)}\theta|\dot{V}|\right) \mu_{\eta, k, V}^{\text{el}}(|V_\Gamma| \geq \theta|\dot{V}|). \end{aligned}$$

Taking $\theta = 3/4$, proving the following display would conclude the proof of the upper bound in Eq. (3.3): there is an absolute constant C such that for $\beta > \beta_0$, and every η , we have

$$\mu_{\eta, k, V}^{\text{el}}(|V_\Gamma| \geq 3|\dot{V}|/4) \geq 1 - C e^{-(\beta - C)|\dot{V}|^{1/3}}. \quad (3.6)$$

as it will in particular be greater than 0.9 for every V . The event $|V_\Gamma| \geq 3|\dot{V}|/4$ is a subset of the union of

- \mathcal{B}_1 : there exists a contour γ confining a vertex in \dot{V} and having $|\gamma| \geq |\dot{V}|^{1/2}$
- \mathcal{B}_2 : the subset of \dot{V} interior to outermost contours γ with $|\gamma| \leq |V|^{1/2}$ is at least $|\dot{V}|/4$.

We first bound the probability of \mathcal{B}_1 . The number of choices for γ having $|\gamma| = r \geq |\dot{V}|^{1/2}$ is at most $|\dot{V}|4^r$, first for picking a site in \dot{V} to root the contour at, then for picking the contour γ . The probability of any fixed such γ being an outermost contour is at most $\exp(-(\beta - 3)r)$ per Corollary 3.9. Thus,

$$\mu_{\eta, k, V}^{\text{el}}(\mathcal{B}_1) \leq \sum_{r \geq |\dot{V}|^{1/2}} |V|4^r e^{-(\beta - 3)r} \leq e^{-(\beta - C)|\dot{V}|^{1/2}}.$$

Next, we consider the probability of \mathcal{B}_2 . For \mathcal{B}_2 to happen, there must be a collection G of outermost elementary contours each of which have $|\gamma| \leq |\dot{V}|^{1/2}$, and for which

$$\sum_{\gamma \in G} |\text{Int}(\gamma) \cap \dot{V}| > \frac{1}{4}|\dot{V}|.$$

In particular, there must exist $\ell = 1, \dots, \frac{1}{2} \log_2 |\dot{V}|$ such that if \mathfrak{U}_ℓ is the set of $\gamma \in G$ having $|\gamma| \in [2^{\ell-1}, 2^\ell)$,

$$\sum_{\gamma \in \mathfrak{U}_\ell} |\text{Int}(\gamma) \cap \dot{V}| \geq \frac{|\dot{V}|}{8\ell^2}.$$

Consider the probability of this for fixed ℓ , and note that the above requires

$$\sum_{\gamma \in \mathfrak{U}_\ell} |\gamma| \geq 4 \sum_{\gamma \in \mathfrak{U}_\ell} 2^{-\ell/2} |\text{Int}(\gamma)| \geq \frac{|\dot{V}|}{2\ell^2} 2^{-\ell/2}. \quad (3.7)$$

Counting over the possible choices for such \mathfrak{U}_ℓ , and applying Lemma 3.7 to bound the probability that \mathfrak{U}_ℓ are outermost contours, the probability that there exists \mathfrak{U}_ℓ having Eq. (3.7) is at most

$$\sum_{r \geq \frac{|\mathring{V}|}{2^{(\ell+1)/2} \ell^2}} 2^r \binom{|\mathring{V}|}{2^{(-\ell-1)} |\mathring{V}|} 4^r e^{-(\beta-2)r} \leq e^{|\mathring{V}| \ell 2^{-(\ell-1)}} e^{-(\beta-3) \frac{1}{2\ell^2} 2^{-\ell/2} |\mathring{V}|}.$$

If β is large, for every ℓ , this is at most $\exp(-\frac{\beta}{3\ell^2 2^{\ell/2}} |\mathring{V}|)$. The coefficient of $|\mathring{V}|$ is decreasing in ℓ and therefore, this is at most its value for $\ell = \frac{1}{2} \log_2 |\mathring{V}|$ which gives $e^{-4\beta |\mathring{V}|^{3/4} / (3(\log_2 |\mathring{V}|)^2)}$. This is at least $\exp(-\beta |\mathring{V}|^{1/2})$ as long as $|\mathring{V}|$ is larger than an absolute constant. The sum of the above probabilities thus yields the bound of Eq. (3.6).

Lower bound in Eq. (3.3). In the other direction, define the identity map T that acts on the subset of configurations in $\mathcal{G}_{\eta, k+1, V}^{\text{el}}$ that have $\{\varphi_V \geq 1\}$ and sends them to $\mathcal{G}_{\eta, k, V}$ by preserving the same contour collection, just viewing them all as being based at 1 height lower. This is equivalent to shifting the height configuration down by 1 everywhere, and thus $\{\varphi_V \geq 1\}$ implies that $T\varphi$ will respect the floor constraint. The core of the argument is to inductively estimate the fraction of $\mathcal{G}_{\eta, k+1, V}^{\text{el}}$ that has $\{\varphi_V \geq 1\}$ (for which this map is valid), as a way to bound the relative weight of $Z_{\eta, k+1, V}^{\text{el}} / Z_{\eta, k, V}^{\text{el}}$.

A potential problem is that the image under this map may not be in $\mathcal{G}_{\eta, k, V}^{\text{el}}$ because contours can be $(k+1)$ -elementary but not k -elementary. To this end, define the good event \mathcal{G} on a contour collection (with a boundary height so that the contours have associated heights) as follows:

$$\mathcal{G} := \left\{ \Gamma : \forall \gamma \in \Gamma, |\gamma| \leq e^{3\beta h_b(\gamma)} \right\}. \quad (3.8)$$

where $h_b(\gamma)$ is its base height, i.e., $h_b(\gamma) = h(\gamma) + 1$ if $S(\gamma) = -$ and $h_b(\gamma) = h(\gamma) - 1$ if $S(\gamma) = +$. Because outermost contours in $\mathcal{G}_{\eta, k+1, V}^{\text{el}}$ have base height $k+1$, and $\text{diam}(\gamma) \leq |\gamma|$, we have by definition of k -elementary, that

$$T(\mathcal{G} \cap \{\varphi_V \geq 1\}) \subset \mathcal{G}_{\eta, k, V}^{\text{el}}.$$

Furthermore, T is an injection because the pre-image can be recovered uniquely by shifting all heights up by 1. As a consequence of this inclusion, and the simple calculation of the weight change $e^{\lambda|V|}$ under application of the map T , for every h ,

$$\frac{Z_{\eta, h, V}^{\text{el}}}{Z_{\eta, h+1, V}^{\text{el}}} \geq e^{\lambda|V|} \mu_{\eta, h+1, V}^{\text{el}}(T^{-1}(\mathcal{G}_{\eta, h, V}^{\text{el}})) \geq e^{\lambda|V|} \mu_{\eta, h+1, V}^{\text{el}}(\mathcal{G}, \varphi_V \geq 1). \quad (3.9)$$

Our goal is now to establish for every h ,

$$\mu_{\eta, h+1, V}^{\text{el}}(\mathcal{G}, \varphi_V \geq 1) \geq \exp\left(-\frac{3}{2}|V|e^{-(4\beta-\lambda)(h+1)}\right), \quad (3.10)$$

since plugging this bound into Eq. (3.9) and iterating over $h = k, \dots, m-1$ yields the claimed lower bound.

Henceforth, fix any h . The core content of Eq. (3.10), is that the probability of $\mathcal{G}^c \cup \{\min_v \varphi_v = 0\}$ is dominated by the event that at some single site $v \in V$ a singleton spike reaches down from height $h+1$ to height 0, as that is what provides the rate $e^{-(4\beta-\lambda)(h+1)}$. Towards this, for a simply-connected subdomain $U \subset V$ and point $y \in U$, we can define the event in $\mathcal{G}_{\zeta, h+1, U}^{\text{el}}$,

$$\tilde{\mathcal{G}}_y = \{\varphi_y \geq 1\} \cap \bigcap_{\gamma: y \in \text{Int}(\gamma)} \{|\gamma| \leq e^{3\beta h_b(\gamma)}\}.$$

Claim 3.11. *Let $U \subsetneq V$ be simply connected. For any $\ell \geq 1$, ζ and every $y \in U$, we have*

$$\mu_{\zeta, \ell, U}^{\text{el}}(\tilde{\mathcal{G}}_y^c) \leq (1 + e^{-\beta})e^{-(4\beta-\lambda)\ell}.$$

Proof. Let us prove the claim inductively over U . The case $|U| = 1$ is governed entirely by the event of $\{\varphi_y = 0\}$ via singleton contours, which has probability at most $e^{-(4\beta-\lambda)\ell}$. Generally, let us decompose

$$\mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c) \leq \mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c, \mathcal{E}_{\zeta,\ell,U}^0) + \sum_{j \neq 0} \mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c, \mathcal{E}_{\zeta,\ell,U}^j). \quad (3.11)$$

where we recall the events $\mathcal{E}_{\zeta,\ell,U}^j$ from Eq. (3.4) and use the shorthand \mathcal{E}^j . We begin with comparing the second term in Eq. (3.11) back to the first, so that the dominant contribution comes on the event \mathcal{E}^0 . This will be helpful again in the future, so we isolate this reduction with a general event A in place of \tilde{G}_y^c .

Let j_0 be the largest number such that $|\gamma_U| \leq e^{3\beta(\ell-j_0+1)}$ (this is the largest number of down γ_U contours we can reveal while still having that conditional on those contours, the interior measure is the elementary one). For $j > -j_0$, the probability of $A \cap \mathcal{E}^j$ is bounded by the probability that there exist $|j|$ contours that are all γ_U , times the probability that under $\mu_{\text{sgn}(j),\ell+j,U}^{\text{el}}(A, \mathcal{E}^0)$. For the $j \leq -j_0$ contributions, their total probability is at most the probability of first having $j_0 - 1$ down contours, times the probability of having a contour of size $|\gamma_U| \geq e^{3\beta(\ell-j_0+1)}$ (as well as $|\gamma_U| \geq 6$ since $|U| \geq 2$), whose probability by Corollary 3.9 is at most $e^{-(\beta-2)e^{3\beta(\ell-j_0+1)}}$. Together, these imply the bound

$$\sum_{j \neq 0} \mu_{\zeta,\ell,U}^{\text{el}}(A, \mathcal{E}^j) \leq \sum_{j \neq 0} e^{-(\beta-2)|j||\gamma_U|} \mu_{\text{sgn}(j),\ell+j,U}^{\text{el}}(A, \mathcal{E}^0) + \max_{j \neq 0} e^{-(\beta-2)(|j|-1)|\gamma_U| + \max\{6, e^{3\beta(\ell-j+1)}\}}.$$

For the first term, we can use the bound $|\gamma_U| \geq 6$. For the second term, we notice that $6(|j| - 1) + \max\{6, e^{3\beta(\ell-j+1)}\}$ is at least 6ℓ for all j . Together, these imply for a general event A and $U : |U| \geq 2$,

$$\sum_{j \neq 0} \mu_{\zeta,\ell,U}^{\text{el}}(A, \mathcal{E}^j) \leq \sum_{j \neq 0} e^{-6(\beta-2)|j|} \mu_{\text{sgn}(j),\ell+j,U}^{\text{el}}(A, \mathcal{E}^0) + e^{-6(\beta-2)\ell}. \quad (3.12)$$

The last term being at most $\frac{1}{4}e^{-\beta}e^{-(4\beta-\lambda)\ell}$, we have reduced Eq. (3.11) to bounding $\mu_{\eta,\ell+j,U}^{\text{el}}(A, \mathcal{E}^0)$.

In what follows, let γ_y^{out} be the outermost contour having $y \in \text{Int}(\gamma_y^{\text{out}})$. We have for general ℓ ,

$$\mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c, \mathcal{E}^0) \leq \mu_{\zeta,\ell,U}^{\text{el}}(\varphi_y = 0, |\gamma_y^{\text{out}}| = 4) + \sum_{6 \leq r \leq e^{3\beta\ell}} \mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c, |\gamma_y^{\text{out}}| = r, \mathcal{E}^0) + \mu_{\zeta,\ell,U}^{\text{el}}(|\gamma_y^{\text{out}}| \geq e^{3\beta\ell}). \quad (3.13)$$

For the first term in Eq. (3.13), observe that if $|\gamma_y^{\text{out}}| = 4$ and $\varphi_y = 0$, the only contours surrounding y are exactly ℓ down contours that are each the singleton contour $\gamma_{\{y\}}$. By explicit calculation, the probability of such configurations is at most $e^{-(4\beta-\lambda)\ell}$. For the third term in Eq. (3.13), we can sum the probability of $\gamma_y^{\text{out}} = \gamma$ over all possible γ having $|\gamma| \geq e^{3\beta\ell}$ using Corollary 3.9, to find that

$$\sum_{r \geq e^{3\beta\ell}} \mu_{\zeta,\ell,U}^{\text{el}}(|\gamma_y^{\text{out}}| = r) \leq \sum_{r \geq e^{3\beta\ell}} 4^r e^{-(\beta-3)r} \leq e^{-(\beta-5)\max\{6, e^{3\beta\ell}\}} \leq e^{-6(\beta-5)\ell}.$$

For the second term in Eq. (3.13), bounding the probability of $|\gamma_y^{\text{out}}| = r$ similarly, and then conditioning on that outermost contour surrounding y , (using that for $r \leq e^{3\beta\ell}$ its interior will necessarily only consist of $h(\gamma_y^{\text{out}})$ -elementary contours),

$$\sum_{r \geq 6} \mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c, |\gamma_y^{\text{out}}| = r, \mathcal{E}_{\zeta,\ell,U}^0) \leq \sum_{r \geq 6} 4^r e^{-(\beta-3)r} \max_{|\gamma_y^{\text{out}}| \geq 6; \text{Int}(\gamma_y^{\text{out}}) \neq U} \max_{\zeta} \mu_{\zeta,\ell \pm 1, \text{Int}(\gamma_y^{\text{out}})}^{\text{el}}(\tilde{G}_y^c).$$

Since $\text{Int}(\gamma_y^{\text{out}})$ is strictly contained in U , we can use the inductive hypothesis on this last quantity, to get

$$\sum_{r \geq 6} \mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c, |\gamma_y^{\text{out}}| = r, \mathcal{E}_{\zeta,\ell,U}^0) \leq e^{-6(\beta-5)}(1 + e^{-\beta})e^{-(4\beta-\lambda)(\ell-1)} \leq \frac{1}{4}e^{-\beta}e^{-(4\beta-\lambda)\ell}.$$

Combining these three bounds, we get for general $\ell \geq 1$, and ζ ,

$$\mu_{\zeta,\ell,U}^{\text{el}}(\tilde{G}_y^c, \mathcal{E}^0) \leq (1 + \frac{1}{2}e^{-\beta})e^{-(4\beta-\lambda)\ell}. \quad (3.14)$$

Plugging this into Eq. (3.12) with ℓ replaced by $\ell + j$ and A replaced by \tilde{G}_y^c gives

$$\begin{aligned} \sum_{j \neq 0} \mu_{\zeta, \ell, U}^{\text{el}}(\tilde{G}_y^c, \mathcal{E}^j) &\leq \sum_{j \neq 0} e^{-6(\beta-2)|j|} (1 + e^{-\beta}) e^{-(4\beta-\lambda)(\ell+j)} + \frac{1}{4} e^{-\beta} e^{-(4\beta-\lambda)\ell} \\ &\leq (1 + e^{-\beta}) e^{-(4\beta-\lambda)\ell} \sum_{j \neq 0} e^{-2(\beta-2)|j|} + \frac{1}{4} e^{-\beta} e^{-(4\beta-\lambda)\ell}, \end{aligned}$$

which is at most $\frac{1}{2} e^{-\beta} e^{-(4\beta-\lambda)\ell}$. Combined with Eq. (3.14) and plugged into Eq. (3.11) we conclude. \blacksquare

In order to go from the one-point estimate to the probability that \tilde{G}_y holds at all $y \in V$, we argue as follows. (A boosting like this in the setting of the SOS model without el could be done by means of FKG, but in the absence of FKG— μ^{el} does not satisfy this and further the event \mathcal{G} is non-monotone—a revealing procedure is required.) For a configuration in $\mathcal{G}_{\eta, h+1, V}^{\text{el}}$, let

$$G_x = \{\varphi : \varphi(\text{Int}(\gamma_x^{\text{out}})) \in \mathcal{G}\} \cap \{\varphi : \varphi(\text{Int}(\gamma_x^{\text{out}})) \geq 1\}.$$

The role of this event is that

$$\mathcal{G} \cap \{\varphi_V \geq 1\} = \bigcap_{x \in V} G_x \quad \text{and} \quad G_x^c \subset \bigcup_{y \in \text{Int}(\gamma_x^{\text{out}})} \tilde{G}_y^c.$$

Our goal is to establish that if we order the vertices of V in such a manner that $V \setminus \bigcup\{y : y < x\}$ is always simply connected (e.g., repeatedly assign x to be an arbitrary boundary vertex of $V \setminus \bigcup\{y : y < x\}$), then

$$\mu_{\eta, h+1, V}^{\text{el}}(G_x \mid (G_y)_{y < x}) \geq 1 - (1 + \varepsilon\beta) e^{-(4\beta-\lambda)(h+1)}. \quad (3.15)$$

from which the claim would follow since

$$\mu_{\eta, h+1, V}^{\text{el}}\left(\bigcap_{x \in V} G_x\right) = \prod_{x \in V} \mu_{\eta, h+1, V}^{\text{el}}(G_x \mid (G_y)_{y < x}) \geq \left(1 - (1 + \varepsilon\beta) e^{-(4\beta-\lambda)(h+1)}\right)^{|V|},$$

where η_i are the induced boundary signings by the exposure of the outermost contours confining $y < x$. This gives the claimed bound of Eq. (3.10) since $1 - x \geq e^{-x}$ for $0 < x < 1/2$.

In order to establish Eq. (3.15), for ease of notation, let

$$\nu_{h+1} = \mu_{\eta_x, h+1, V}^{\text{el}}(\cdot \mid (G_y)_{y < x}, (\mathcal{F}_y)_{y < x}),$$

where, by \mathcal{F}_y we mean the σ -algebra generated by the configuration interior to γ_y^{out} . Notice that ν_{h+1} is exactly the measure $\mu_{\zeta, h+1, U}^{\text{el}}$ where $U = V \setminus \bigcup_{y < x} \text{Int}(\gamma_y^{\text{out}})$ (which is $(\mathcal{F}_y)_{y < x}$ -measurable and is simply-connected by the ordering we imposed over V) and ζ agrees with η on $\partial_i V \setminus \partial_o U$ and takes the opposite sign of γ_y^{out} on every vertex in $\partial_o U$. If U is a singleton, the bound of Eq. (3.15) is immediate as the probability of G_x^c would be $e^{-(4\beta-\lambda)(h+1)}$ if a single-column spike down to height zero is permissible and 0 otherwise.

Suppose now that $|U| \geq 2$. Then we can naturally decompose

$$\nu_{h+1}(G_x^c) \leq \nu_{h+1}(G_x^c, \mathcal{E}^0) + \sum_{j \neq 0} \nu_{h+1}(G_x^c, \mathcal{E}^j), \quad (3.16)$$

and apply Eq. (3.12) to get the bound

$$\nu_{h+1}(G_x^c) \leq \sum_j e^{-(\beta-2)|j|} \nu_{h+1+j}(G_x^c, \mathcal{E}^0) + e^{-6(\beta-2)(h+1)}, \quad (3.17)$$

(where we use the shorthand ν_{h+1+j} for $\mu_{\text{sgn}(j), h+1+j, U}$ when $j \neq 0$). We now examine $\nu_{h+1+j}(G_x^c, \mathcal{E}^0)$. For general ℓ , just as in Eq. (3.13), we can split

$$\nu_\ell(G_x^c, \mathcal{E}^0) \leq \nu_\ell(\varphi_x = 0, |\gamma_x^{\text{out}}| = 4) + \sum_{6 \leq r \leq e^{3\beta\ell}} \nu_\ell(G_x^c, |\gamma_x^{\text{out}}| = r, \mathcal{E}^0) + \nu_\ell(|\gamma_x^{\text{out}}| \geq e^{3\beta\ell}).$$

As argued following Eq. (3.13), the first term will contain the dominant contribution of at most $e^{-(4\beta-\lambda)\ell}$, and the last term will be at most $e^{-6(\beta-5)\ell} \leq \frac{1}{4}e^{-\beta}e^{-(4\beta-\lambda)\ell}$. The middle sum is similarly at most

$$\sum_{r \geq 6} 4^r e^{-(\beta-3)r} \max_{|\gamma_x^{\text{out}}| \geq 6; \text{Int}(\gamma_x^{\text{out}}) \neq U} \max_{\zeta} \mu_{\zeta, \ell \pm 1, \text{Int}(\gamma_x^{\text{out}})}^{\text{el}}(G_x^c).$$

Recalling that G_x^c is a subset of $\bigcup_{y \in \text{Int}(\gamma_x^{\text{out}})} \tilde{G}_y^c$, by a union bound and Claim 3.11,

$$\sum_{6 \leq r \leq e^{3\beta\ell}} \nu_{\ell}(G_x^c, |\gamma_x^{\text{out}}| = r, \mathcal{E}^0) \leq \sum_{r \geq 6} e^{-(\beta-4)r} r^2 (1 + e^{-\beta}) e^{-(4\beta-\lambda)(\ell-1)}.$$

This sum is evidently at most $\frac{1}{4}e^{-\beta}e^{-(4\beta-\lambda)\ell}$. Altogether, we get

$$\nu_{\ell}(G_x^c, \mathcal{E}^0) \leq (1 + \frac{1}{2}e^{-\beta})e^{-(4\beta-\lambda)\ell}.$$

Plugging this bound into Eq. (3.17) with $\ell = h + 1 + j$, we get

$$\nu_{h+1}(G_x^c) \leq \sum_j e^{-(\beta-2)|j|} (1 + \frac{1}{2}e^{-\beta}) e^{-(4\beta-\lambda)(h+1+j)} + e^{-6(\beta-2)(h+1)} \leq (1 + e^{-\beta})e^{-(4\beta-\lambda)(h+1)},$$

yielding the claimed Eq. (3.15) and concluding the proof. \blacksquare

3.3. The elementary cluster expansion. We can use Theorem 3.6 to deduce that the elementary partition functions admit convergent cluster expansions. For a thorough background on cluster expansion, see [19, Chapter 5], though we will only use fairly elementary consequences. We will work with the renormalized elementary partition function $Z_{\eta,k,V}^{\text{rn,el}}$, though we recall that for simply-connected domains, this is the same as $Z_{\eta,k,V}^{\text{el}}$ per Lemma 3.5. For non-simply connected domains, we define $Z_{\eta,k,U}^{\text{rn,el}}$ as in Eq. (3.2) (contours encircling holes of U being admissible).

Lemma 3.12. *For each i , the limit,*

$$f_i^{\text{el}} = \lim_{n \rightarrow \infty} \frac{1}{|\Lambda_n|} \log Z_{\eta,i,\Lambda_n}^{\text{el}},$$

exists. Furthermore, for all (not necessarily simply connected) U , all i and all η ,

$$|\log Z_{\eta,i,U}^{\text{rn,el}} - |U| f_i^{\text{el}}| \leq e^{-\beta/2} |\partial_e U|.$$

Proof. The partition function $Z^{\text{rn,el}}$ fits into the standard cluster expansion framework with hardcore interactions between contours (which serve as the polymers). Furthermore, per Item 2 of Theorem 3.6, they satisfy the requisite exponential decay property on their weights for the Kotecky–Preiss condition [29] to hold at large β . The lemma is then a direct corollary of convergence of the cluster expansion. We have used the bound $e^{-\beta/2}$ on the sum of weights of all possible elementary contours confining a vertex v in their interior, which holds for β larger than some absolute constant. \blacksquare

Using the bounds we obtained in Item 1 from Theorem 3.6, we can provide uniform lower bounds on $f_i^{\text{el}} - f_j^{\text{el}}$ in certain windows of the parameter λ . Namely, we get an infinite sequence of disjoint intervals I_i in which the boundary condition i elementary partition function dominates. Define

$$I_i = [\lambda_-^{(i)}, \lambda_+^{(i)}] \quad \text{for} \quad \lambda_-^{(i)} := e^{-4\beta i - 2\beta}, \quad \text{and} \quad \lambda_+^{(i)} := e^{-4\beta i - \beta}. \quad (3.18)$$

Lemma 3.13. *If $\lambda \in I_i$, then for every $j \neq i$ we have*

$$f_i^{\text{el}} - f_j^{\text{el}} \geq e^{-4\beta(i \wedge (j+1)) - 3\beta}.$$

Proof. Applying Item 1 of Theorem 3.6 with $\lambda \in I_i$, we find that for every $k < m$,

$$\lambda(m-k) - 2e^{-(4\beta-\lambda)(k+1)} \leq \frac{1}{|V|} \log \frac{Z_{\eta,k,V}^{\text{el}}}{Z_{\eta,m,V}^{\text{el}}} \leq \lambda(m-k) - \frac{1}{2} \frac{|\hat{V}|}{|V|} e^{-(4\beta-\lambda)(k+1)} + \frac{\log 2}{|V|}. \quad (3.19)$$

Taking $k = i$ and $m = j > i$, we see that

$$\frac{1}{|V|} \log \frac{Z_{\eta,i,V}^{\text{el}}}{Z_{\eta,j,V}^{\text{el}}} \geq \lambda(j-i)(1 - 2e^{-(4\beta-\lambda)(i+1)+4\beta i+2\beta}) \geq \lambda(j-i)(1 - 2e^{-4\beta+e^{-4\beta i}(i+1)+2\beta}).$$

Since $e^{-4\beta i}(i+1) \leq 1$, we see that for β large, this satisfies

$$\frac{1}{|V|} \log \frac{Z_{\eta,i,V}^{\text{el}}}{Z_{\eta,j,V}^{\text{el}}} \geq \lambda(j-i)(1 - \varepsilon_\beta),$$

and ε_β can be taken to be $e^{-\beta}$. Since this bound is uniform over V , we can take V to be sequences of boxes going to \mathbb{Z}^2 to get the lower bound in the limit.

If $j < i$, the upper bound of Eq. (3.19) with $k = j$ and $m = i$ implies

$$\frac{1}{|V|} \log \frac{Z_{\eta,j,V}^{\text{el}}}{Z_{\eta,i,V}^{\text{el}}} \leq e^{-4\beta i-\beta}(i-j) - \frac{1}{2}e^{-4\beta(j+1)} + O\left(\frac{|\partial V|}{|V|}\right).$$

Since $j < i$, we can write $e^{-4\beta i-\beta} = e^{-4\beta(j+1)}e^{-4\beta(i-j-1)-\beta}$ and use the fact that $e^{-4\beta(i-j-1)-\beta}(i-j) - \frac{1}{2}$ is at most $-\frac{1}{2} + e^{-\beta} \leq -e^{-\beta}$ uniformly over $i-j \geq 1$, yielding

$$\frac{1}{|V|} \log \frac{Z_{\eta,j,V}^{\text{el}}}{Z_{\eta,i,V}^{\text{el}}} \leq -e^{-4\beta(j+1)-\beta} + O\left(\frac{|\partial V|}{|V|}\right).$$

Taking V to be sequences of boxes going to \mathbb{Z}^2 , and negating the above, we get the claimed bound. \blacksquare

4. AN INFINITE SEQUENCE OF CRITICAL EXTERNAL FIELD VALES

The aim in this section is to establish an infinite sequence of discontinuous phase transitions as one varies the external field λ , with critical points $(\lambda_c^{(k)})_{k \geq 0}$ ordered $\lambda_c^{(0)} > \lambda_c^{(1)} > \dots$ and accumulating at zero, such that for all $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$ the SOS measure with boundary conditions k is thermodynamically stable. In particular, if one starts with boundary conditions k on any domain, then the interface will remain rigid about height k in the bulk of that domain whenever $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$.

Moreover, and crucially for the purposes of this paper, we probe the precise behavior as λ approaches the critical points, demonstrating that both of the heights k and $k-1$ will be stable as long as the domain is smaller than $|\lambda - \lambda_c^{(k-1)}|^{-1}$. It is in this and the following section that the true benefits of the contour representation and our more geometric arguments arise, allowing us to use much shorter arguments to get sharp results analogous to [11] for all $\lambda \geq 0$ rather than only for $\lambda \geq c_\beta > 0$ fixed and bounded away from zero.

4.1. Rigidity of height h when $\lambda \in I_h$. For each h , the full window $[\lambda_c^{(h)}, \lambda_c^{(h-1)}]$ will contain the interval I_h from Eq. (3.18), so we begin by establishing the rigidity in the smaller interval away from the critical points. A general contour collection in $\mathcal{G}_{\eta,h,V}^{\text{rn}}$ consists both of h -elementary contours, and ones that are not h -elementary. The above estimates give us the ability to prove that even the contours that are not h -elementary have exponentially small h -renormalized weight when $\lambda \in I_h$.

Lemma 4.1. *If $\lambda \in I_h$, then for every (not necessarily elementary) γ , we have*

$$W_h^{\text{rn}}(\gamma) \leq e^{-(\beta-4)|\gamma|}.$$

The proof of Lemma 4.1 goes by treating nested collections of non h -elementary contours together, so that the partition functions on their complements are exactly the elementary ones, for which we know that being at height h is the most preferred when $\lambda \in I_h$, per Lemma 3.12. The key step in the argument is to bound the total weight contribution over all choices of non-elementaries that can be nested in γ , noticing that their being non-elementary is enough to beat out the entropy over the possible locations they can be placed (also using that if they are far from the boundary of their nesting contour, there is a cost due to the height- h

elementary free energy being the dominant one). This argument was performed in [15, Lemmas 2.13–2.14] for their cylinders rather than our contours, using an encoding of *clusters of non-elementary cylinders* into weighted trees. With only minor modifications, this argument can be imported into our setting of contours, and we therefore defer the details to Appendix B.

Given Lemma 4.1, we can deduce convergence of the cluster expansion for $Z_{h,V}$ (see the next subsection) as well as the following rigidity of the (full) SOS measure with h -boundary conditions when $\lambda \in I_h$.

Lemma 4.2. *Suppose $\lambda \in I_h$. For every V , every η , and every $x \in V$,*

$$\mu_{\eta,h,V}(|\varphi_x - h| \geq r) \leq e^{-4(\beta-5)r}.$$

Proof. It suffices to bound the probability that there is a set of all contours nesting x have total size at least $4r$. The bound on this goes via a standard Peierls estimate, with the set of all such configurations being mapped, via deletion of their outermost contour, to the set of configurations whose set of all contours nesting x has total size at least $4(r-1)$ (since $|\varphi_x - h|$ would change by 1 under this operation). If in the contour representation of φ , γ_x^{out} is the outermost contour confining x , we can enumerate over all such choices of the contour having size $r \geq 4$, there being at most 4^r many such choices, and recall from Lemma 3.2 that the probability of a specific contour γ being γ_x^{out} is exactly $W_h^{\text{rn}}(\gamma)$. Together with Lemma 4.1, we get

$$\mu_{\eta,h,V} \left(\sum_{\gamma: x \in \text{Int}(\gamma)} |\gamma| \geq 4r \right) \leq \sum_{r \geq 4} 4^r e^{-(\beta-4)r} \mu_{\eta,h,V} \left(\sum_{\gamma: x \in \text{Int}(\gamma)} |\gamma| \geq 4(r-1) \right),$$

which, when iterated implies the claimed bound for large β . ■

Truncated free energies and cluster expansion. Having established exponential tails not only on elementary γ , but also when $\lambda \in I_h$, non-elementary γ based at height h , we can deduce that the full partition function $Z_{\eta,h,V}$ admits a convergent cluster expansion when $\lambda \in I_h$. In order to compare this with other partition functions (which may not admit convergent cluster expansions), we introduce “truncated” partition functions which modify the renormalized weights so that they admit convergent cluster expansions at all reference heights and all λ . These will coincide with the usual partition function whenever the usual partition function has exponential tails on its renormalized weights.

Definition 4.3. For an contour γ , define its h -truncated weight

$$W_h^{\text{tr}}(\gamma) = \min \{ W_h^{\text{rn}}(\gamma), e^{-(\beta-5)|\gamma|} \}. \quad (4.1)$$

As a consequence of Item 2 of Theorem 3.6, we have for every elementary γ , $W_h^{\text{tr}}(\gamma) = W_h^{\text{rn}}(\gamma)$. Also, by Lemma 4.1, if $\lambda \in I_h$, then every non-elementary contour γ also has $W_h^{\text{tr}}(\gamma) = W_h^{\text{rn}}(\gamma)$. We can now use the truncated weights to construct truncated partition functions.

Definition 4.4. For general η, h, V , define the truncated partition function

$$Z_{\eta,h,V}^{\text{tr}} = e^{-\lambda h|V|} \sum_{\Gamma \in \mathcal{G}_{\eta,h,V}^{\text{rn}}} \prod_{\gamma \in \Gamma} W_h^{\text{tr}}(\gamma).$$

By definition of W_h^{tr} and $Z_{\eta,h,V}^{\text{tr}}$, we always have

$$Z_{\eta,h,V}^{\text{tr}} \leq Z_{\eta,h,V}^{\text{rn}}, \quad (4.2)$$

By the earlier observation that $W_h^{\text{tr}}(\gamma) = W_h^{\text{rn}}(\gamma)$ for all γ when $\lambda \in I_h$, we have

$$Z_{\eta,h,V}^{\text{tr}} = Z_{\eta,h,V}^{\text{rn}} \quad \text{for all } \lambda \in I_h. \quad (4.3)$$

By Lemma 3.2, if V is simply connected both the right-hand sides above can be replaced by $Z_{\eta,h,V}$.

Below, we collect some standard consequences of convergent cluster expansions that we get for the truncated partition functions. Towards this, define the following finite-volume free energies:

$$f_{\eta,h,V}^{\text{tr}} = \frac{1}{|V|} \log Z_{\eta,h,V}^{\text{tr}}, \quad \text{and} \quad f_{\eta,h,V} = \frac{1}{|V|} \log Z_{\eta,h,V}. \quad (4.4)$$

As argued in Lemma 3.12, the uniform $e^{-(\beta-5)|\gamma|}$ bound on truncated contour weights implies that the truncated partition functions admit convergent cluster expansions, yielding the following.

Lemma 4.5. *For each h , the following infinite-volume free energy*

$$f_h^{\text{tr}} := \lim_{n \rightarrow \infty} f_{\eta, h, \Lambda_n}^{\text{tr}},$$

exists. Furthermore, for all (not necessarily simply connected) U , all h , and all η ,

$$|\log Z_{\eta, h, U}^{\text{tr}} - |U| f_h^{\text{tr}}| \leq e^{-\beta/2} |\partial_e U|.$$

At the same time, due to Eq. (4.3) and Lemma 3.2, when $\lambda \in I_h$,

$$f_h := \lim_{n \rightarrow \infty} \frac{1}{|\Lambda_n|} \log Z_{\eta, h, \Lambda_n},$$

exists and is equal to f_h^{tr} because the partition functions are equal. Since

$$\left| \frac{1}{|V|} \log \frac{Z_{\eta, j, V}}{Z_{\zeta, h, V}} \right| \leq (\beta + \lambda) \cdot |h - j| \cdot \frac{|\partial_e V|}{|V|},$$

(by crudely forcing the heights along the boundary to be at h), we find that when $\lambda \in I_h$,

$$f_j = \lim_{n \rightarrow \infty} f_{\eta, j, \Lambda_n} \text{ exists, and } f_j = f_h = f_h^{\text{tr}} \text{ for all } j \neq h. \quad (4.5)$$

We now show using the truncated free energies, that there exist an infinite sequence of critical values of λ dictating transitions between which height maximizes the infinite-volume truncated free energy.

Proposition 4.6. *There exists $(\lambda_c^{(i)})_{i=0}^\infty > 0$ such that for every i , it has $I_i \subset [\lambda_c^{(i)}, \lambda_c^{(i-1)}]$ and*

- (1) *If $\lambda_c^{(i)} < \lambda < \lambda_c^{(i-1)}$, then $f_i^{\text{tr}} > f_j^{\text{tr}}$ for all $j \neq i$.*
- (2) *If $\lambda = \lambda_c^{(i)}$, then $f_{i+1}^{\text{tr}} = f_i^{\text{tr}} > f_j^{\text{tr}}$ for $j \notin \{i, i+1\}$.*

Proof. We first use Eq. (4.5), then Eq. (4.2) to deduce that for every $\lambda \in I_i$, we have for all j ,

$$f_i^{\text{tr}} = f_i = f_j \geq f_j^{\text{tr}}. \quad (4.6)$$

We now lower bound the λ -derivative of $f_i^{\text{tr}} - f_j^{\text{tr}}$ to deduce that there is a single λ at which the former becomes strictly larger than the latter. We start by considering the derivative for finite volumes. (The boundary signing can be arbitrary here, so let us take it as free and drop it from the notation.) Explicit differentiation gives

$$\left| \frac{d}{d\lambda} f_{i, V}^{\text{tr}} + i \right| \leq \frac{1}{|V|} \mu_{i, V}^{\text{tr}} \left[\sum_{\gamma \in \Gamma} |\text{Int}(\gamma)| \right],$$

(where when the derivative does not exist due to the min in $W^{\text{tr}}(\gamma)$, this bound holds for both the right and left derivatives). By standard cluster expansion reasoning, the total volume confined by $\gamma \in \Gamma$ has an exponential tail beyond $e^{-\beta|V|}$, and therefore the expectation on the right above is at most $e^{-\beta/2|V|}$, uniformly over λ . This implies that for any simply connected V , for all i, j

$$\left| \frac{d}{d\lambda} (f_{i, V}^{\text{tr}} - f_{j, V}^{\text{tr}}) - (j - i) \right| \leq e^{-\beta/2}. \quad (4.7)$$

Taking $V = \Lambda_n$ and $n \rightarrow \infty$, this implies that $f_i^{\text{tr}}, f_j^{\text{tr}}$ are (Lipschitz-)continuous in λ , and for all $\lambda > 0$, we have for all i, j

$$\frac{d}{d\lambda} (f_i^{\text{tr}} - f_j^{\text{tr}}) \geq (1 - \varepsilon_\beta)(j - i).$$

(Again, when there are issues of non-differentiability, this lower bound holds both for the right and left derivatives.) When this lower bound is combined with Eq. (4.6), the desired sequence of critical values $\lambda_c^{(i)}$ and the claimed strict inequalities on truncated free energies follows by continuity in λ . \blacksquare

Eq. (4.7) also gives the following relation between $f_k^{\text{tr}} - f_{k\pm 1}^{\text{tr}}$ and $d^\pm(\lambda)$ from Eq. (1.2).

Corollary 4.7. *We have for every k that for $\lambda < \lambda_c^{(k-1)}$,*

$$(1 - \varepsilon_\beta)(\lambda_c^{(k-1)} - \lambda) \leq f_k^{\text{tr}} - f_{k-1}^{\text{tr}} \leq (1 + \varepsilon_\beta)(\lambda_c^{(k-1)} - \lambda),$$

and similarly, for $\lambda > \lambda_c^{(k)}$, we have $(1 - \varepsilon_\beta)(\lambda - \lambda_c^{(k)}) \leq f_k^{\text{tr}} - f_{k+1}^{\text{tr}} \leq (1 + \varepsilon_\beta)(\lambda - \lambda_c^{(k)})$.

In what follows, to simultaneously consider the $i = 0$ case, by convention set $\lambda_c^{(-1)} = \infty$.

4.2. Extending control of non-elementary contours to $\lambda \in [\lambda_c^{(h)}, \lambda_c^{(h-1)}]$. Our aim is now to prove that the rigidity at height h (e.g., exponential tails on $W_h^{\text{tr}}(\gamma)$ even for non-elementary γ) holds throughout the full window $[\lambda_c^{(h)}, \lambda_c^{(h-1)}]$ where Proposition 4.6 showed the h -free energy dominates.

From this point forth, our arguments deviate much more significantly, even at the high level, from the strategy of [11] as their arguments became more constrained by the usage of cylinders instead of contours, and begin to necessitate λ bounded away from zero by a β -dependent constant, or equivalently necessitate the boundary level bounded by a β -dependent constant. By contrast, our arguments are significantly shorter than those in [11] and allow us to interpolate all the way down to $\lambda = 0$.

Proposition 4.8. *Let $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$. For every (including non-elementary) contour γ ,*

$$W_k^{\text{tr}}(\gamma) = W_k^{\text{rn}}(\gamma). \quad (4.8)$$

In particular, if $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$, then

$$Z_{\eta,k,V}^{\text{rn}} = Z_{\eta,k,V}^{\text{tr}}. \quad (4.9)$$

and if V is simply connected, $Z_{\eta,k,V} = Z_{\eta,k,V}^{\text{tr}}$.

Proof. We prove the proposition by induction. Eq. (4.8) is already proved for elementary contours in Theorem 3.6 so it suffices to show it for γ that is not k -elementary. Suppose Eq. (4.8) has been shown for every contour γ' having $|\text{Int}(\gamma')| \leq m$ and show it for γ such that $|\text{Int}(\gamma)| = m + 1$. For ease of notation, let $V \cup \{x\} = \text{Int}(\gamma)$ where x is a prescribed point on the boundary $\partial_i(V \cup \{x\})$ such that V is simply connected. Recalling the definition Eq. (4.1), if γ is an up-contour, we wish to show

$$\frac{Z_{+,k+1,V \cup x}}{Z_{+,k,V \cup x}} \leq e^{|\gamma|}. \quad (4.10)$$

If we let $\tilde{\mp}$ be the boundary signing on $\partial_i V$ that is $+$ on all vertices in $\partial_i(V \cup x)$, and free on $\partial_i V \setminus \partial_i(V \cup x)$, the first bound we make in this direction is

$$\frac{Z_{+,k+1,V \cup x}}{Z_{+,k,V \cup x}} \leq e^{\lambda k} \frac{Z_{+,k+1,V \cup x}}{Z_{\tilde{\mp},k,V}} = e^{\lambda k} \frac{Z_{\tilde{\mp},k+1,V}}{Z_{\tilde{\mp},k,V}} \frac{Z_{+,k+1,V \cup x}}{Z_{\tilde{\mp},k+1,V}},$$

using in the first inequality the fact that $Z_{+,k,V \cup x} \geq e^{-\lambda k} Z_{\tilde{\mp},k,V}$. The second ratio on the right is exactly

$$(\mu_{+,k+1,V \cup x}(\varphi_x = k + 1))^{-1} = (\mu_{+,k+1,V \cup x}(\varphi_x \leq k + 1))^{-1}.$$

By monotonicity, the probability of $\{\varphi \leq k + 1\}$ is only increasing as we increase λ , and so the right-hand side above increases as we decrease λ . Therefore, we upper bound the right-hand side by decreasing λ to $\lambda' \in I_{k+1}$. At that point, the rigidity Lemma 4.2 implies the above display is at most $1 + \varepsilon_\beta$.

Plugging this bound in, we are left with

$$\frac{Z_{+,k+1,V \cup x}}{Z_{+,k,V \cup x}} \leq (1 + \varepsilon_\beta) e^{\lambda k} \frac{Z_{\tilde{\mp},k+1,V}}{Z_{\tilde{\mp},k,V}}. \quad (4.11)$$

For the ratio on the smaller domain V , we first use Lemma 2.3 (saying that the ratio is decreasing in λ) to increase the ratio by decreasing λ to $\lambda_c^{(k)}$, at which, by the inductive hypothesis, both partition functions

are exactly their truncated versions. By definition of $\lambda_c^{(k)}$ and Proposition 4.6, we have $f_{k+1}^{\text{tr}} = f_k^{\text{tr}}$ at $\lambda_c^{(k)}$, whence by Lemma 4.5,

$$\frac{Z_{\lambda_c^{(k)}, \tilde{\gamma}, k+1, V}^{\text{tr}}}{Z_{\lambda_c^{(k)}, \tilde{\gamma}, k, V}^{\text{tr}}} \leq \frac{Z_{\lambda_c^{(k)}, \tilde{\gamma}, k+1, V}^{\text{tr}}}{Z_{\lambda_c^{(k)}, \tilde{\gamma}, k, V}^{\text{tr}}} \leq e^{2e^{-\beta/2}|\partial_e V|}.$$

Plugging this bound into Eq. (4.11), we are left with

$$\frac{Z_{+, k+1, V}}{Z_{\tilde{\gamma}, k, V}} \leq (1 + \varepsilon_\beta) e^{\lambda k} e^{2e^{-\beta/2}|\partial_e V|}.$$

Since $\lambda \leq e^{-4\beta k + 2\beta} \wedge 1$, we have $(1 + \varepsilon_\beta) e^{\lambda k} \leq e^2$ say. Since we are assuming the confining contour γ bounding the region $V \cup x$ is not k -elementary, it has $|\gamma| \geq e^{3\beta}$, so for β large,

$$2e^{-\beta/2}|\partial_e V| + 2 \leq |\partial_e(V \cup x)|,$$

which together with the above implies the claimed bound of Eq. (4.10) since $|\partial_e(V \cup x)| = |\gamma|$.

If γ is instead a negative contour, the argument proceeds analogously. Here,

$$(\mu_{-, k-1, V \cup x}(\varphi_x \geq k-1))^{-1},$$

will be increasing as we increase λ , and taking λ to $\lambda' \in I_{k-1}$ is an increase of λ . Similarly, the ratio $Z_{-, k-1, V}^{\text{tr}}/Z_{-, k, V}^{\text{tr}}$ will be increasing in λ , so we can increase λ to $\lambda_c^{(k-1)}$ as needed. ■

We can also deduce the following, showing exponential tails on all contours based at heights $k \pm 1$ if they are small as compared to $d(\lambda_n)$ from Eq. (1.2) so the free energy difference does not overcome the cost in the boundary. Later, this will be crucial to understanding bottlenecks and the formation of critical droplets when λ is microscopically close to a critical value.

Corollary 4.9. *Suppose $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$. If γ is a down contour and $\text{diam}(\gamma) \leq (f_k^{\text{tr}} - f_{k+1}^{\text{tr}})^{-1}$, then $W_{k+1}^{\text{rn}}(\gamma) = W_{k+1}^{\text{tr}}(\gamma)$, and if γ is an up contour with $\text{diam}(\gamma) \leq (f_k^{\text{tr}} - f_{k-1}^{\text{tr}})^{-1}$ then $W_{k-1}^{\text{rn}}(\gamma) = W_{k-1}^{\text{tr}}(\gamma)$.*

Proof. Consider γ a down contour having $\text{diam}(\gamma) \leq (f_k^{\text{tr}} - f_{k+1}^{\text{tr}})^{-1}$. For these, we have

$$\frac{Z_{-, k, \text{Int}(\gamma)}}{Z_{-, k+1, \text{Int}(\gamma)}} \leq \frac{Z_{-, k, \text{Int}(\gamma)}^{\text{tr}}}{Z_{-, k+1, \text{Int}(\gamma)}^{\text{tr}}},$$

where we used Eq. (4.9) on the numerator, and Eq. (4.2). By Lemma 4.5 applied to both terms on the right,

$$W_{k+1}^{\text{rn}}(\gamma) \leq \exp(-\beta|\gamma| + (f_k^{\text{tr}} - f_{k+1}^{\text{tr}})^{-1}|\text{Int}(\gamma)| + 2e^{-\beta/2}|\gamma|).$$

Using the assumption that $\text{diam}(\gamma) \leq (f_k^{\text{tr}} - f_{k+1}^{\text{tr}})^{-1}$, and $|\text{Int}(\gamma)| \leq \text{diam}(\gamma) \cdot |\gamma|$, we get that this is at most $\exp(-(\beta-2)|\gamma|)$ as claimed. The reasoning for a up contour from $k-1$ is analogous. ■

We will also use the following more technical corollary deducing for the full, as opposed to truncated, partition functions that up to an error of $\varepsilon_\beta|\partial_e U|$, the full partition function on U with boundary k is dominant when $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$ (cf. the crude error size $\beta|\partial_e U|$ one gets by forcing the configuration values along the boundary).

Corollary 4.10. *Consider an arbitrary connected domain U (possibly not simply connected) and any boundary signings η and ζ . If $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$, then for all j ,*

$$\frac{Z_{\eta, j, U}}{Z_{\zeta, k, U}} \leq e^{e^{-\beta/3}|\partial_e U||j-k|}.$$

Proof. Suppose first that $j = k + 1$. By Lemma 2.3, we can decrease λ to $\lambda_c^{(k)}$, only increasing the ratio between the partition functions. Since Lemma 3.2 only applied to simply connected domains, we cannot move to Z^{tr} immediately. Instead, if we call H_i the holes of U on the one hand, at $\lambda_c^{(k)}$,

$$Z_{\eta,k+1,U} \leq \frac{Z_{\eta,k+1,U \cup \bigcup_i H_i}}{\prod_i Z_{\eta,k+1,H_i}} = \frac{Z_{\eta,k+1,U \cup \bigcup_i H_i}^{\text{tr}}}{\prod_i Z_{\eta,k+1,H_i}^{\text{tr}}},$$

where the inequality is by inclusion, and the equality uses Proposition 4.8 and simply-connectedness of $U \cup \bigcup_i H_i$, and of H_i . By Lemma 4.5, therefore,

$$Z_{\eta,k+1,U} \leq \exp \left(f_{k+1}^{\text{tr}} |U| + e^{-\beta/2} (|\partial_e(U \cup \bigcup_i H_i)| + \sum_i |\partial_e H_i|) \right). \quad (4.12)$$

On the other hand, if we define $\bar{Z}_{\zeta,k,U}$ as the restricted partition function where no contours confining any hole H_i in their interior are allowed, then

$$Z_{\zeta,k,U} \geq \bar{Z}_{\zeta,k,U} = \bar{Z}_{\zeta,k,U}^{\text{rn}} = \bar{Z}_{\zeta,k,U}^{\text{tr}}.$$

From the convergence of the cluster expansion for $Z_{\zeta,k,U}^{\text{tr}}$, it is not hard to deduce that

$$\frac{\bar{Z}_{\zeta,k,U}^{\text{tr}}}{Z_{\zeta,k,U}^{\text{tr}}} \geq \exp \left(-e^{-\beta/2} \sum_i |\partial_e H_i| \right).$$

Combining the above two, and applying Lemma 4.5 to $Z_{\zeta,k,U}^{\text{tr}}$, we get

$$Z_{\zeta,k,U} \geq \exp \left(f_k^{\text{tr}} |U| - e^{-\beta/2} (|\partial_e U| + \sum_i |\partial_e H_i|) \right).$$

Combined with Eq. (4.12) and the fact from Proposition 4.6 that the limiting truncated free energies f_{k+1}^{tr} and f_k^{tr} are equal at $\lambda_c^{(k)}$, we get

$$\frac{Z_{\eta,k+1,U}}{Z_{\zeta,k,U}} \leq \exp \left(e^{-\beta/2} (|\partial_e U| + 2 \sum_i |\partial_e H_i| + |\partial_e(U \cup \bigcup_i H_i)|) \right).$$

That sum of the various boundary terms counts each edge in $\partial_e U$ at most three times, leaving that the above is at most $e^{3e^{-\beta/2} |\partial_e U|} \leq e^{e^{-\beta/3} |\partial_e U|}$ at large β . If $j = k + \ell$ more generally, we first express

$$\frac{Z_{\eta,j,U}}{Z_{\zeta,k,U}} = \frac{Z_{\eta,k+\ell,U}}{Z_{\eta,k+\ell-1,U}} \dots \frac{Z_{\eta,k+1,U}}{Z_{\zeta,k,U}}.$$

We can then evaluate the last term at $\lambda_c^{(k+1)}$, the second-to-last one at $\lambda_c^{(k+2)}$, etc, each of these only increasing the ratio of partition functions per Lemma 2.3, and then being calculated as above. The other direction, where $j < k$ is analogous, with the increase of λ to $\lambda_c^{(k-1)}$ increasing the partition function ratio. ■

4.3. A geometric rigidity statement about height k when $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$. As a consequence of Proposition 4.8, we can deduce rigidity of the SOS distribution with height k boundary conditions just like Lemma 4.2 throughout the entire regime $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$: for all η and all $x \in V$

$$\mu_{\eta,k,V}(|\varphi_x - h| \geq r) \leq e^{-4(\beta-6)r}, \quad \text{for all } \lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]. \quad (4.13)$$

We go beyond that, showing a more geometric rigidity statement. Namely, if $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$, we show that the set of height k sites qualitatively behave like a highly supercritical percolation.

For this, we will need a few geometric definitions.

Definition 4.11. We say a set of sites $L \subset \mathbb{Z}^2$ forms a *loop* if the vertices of L can be ordered $v_0, \dots, v_{|L|}$ such that $v_0 = v_{|L|}$, for every other $i \neq j$, $v_i \neq v_j$, and $v_i \sim v_{i+1}$ for all i . They form a loop surrounding a set A if A is in a finite connected component of $\mathbb{R}^2 \setminus L$.

Given a configuration φ on a domain U we say $\mathcal{L} \subset U$ is a height- k loop if it is a loop and $\varphi_x = k$ for all $x \in \mathcal{L}$. For a fixed set A , the *outermost* height- k loop in U surrounding A is measurable with respect to the sites exterior to it. In particular, this outermost loop can be revealed by starting from $\partial_i U$ and revealing its $*$ -connected components (connected components with respect to $*$ -adjacency, meaning adjacent or diagonally adjacent to) of non-height- k sites.

Lemma 4.12. *Consider a domain V containing Λ_{m+r} , with boundary conditions k and suppose $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$. For every $\delta > 0$, the outermost loop \mathcal{L} of height- k sites in Λ_{m+r} surrounding the origin satisfies*

$$|\mathcal{L}| \leq 4(1 + \delta)(m + r), \quad \text{and} \quad \Lambda_m \subset \text{Int}(\mathcal{L}),$$

except with $\mu_{\eta,k,V}$ -probability $me^{-(\beta-C)r} + e^{-\delta(\beta-C)(m+r)}$.

Proof. Consider the outermost contours under $\mu_{\eta,k,V}$ whose interiors intersect $\partial_i \Lambda_{m+r}$, as well as any mutually external contours which are incident to those, and so forth. This leaves a collection of contours, call it G the union of whose interiors confines the $*$ -connected components of non-height- k sites of $\partial_i U$.

The probability that $|\mathcal{L}| > 4(1 + \delta)(m + r)$ is bounded by the probability that $\sum_{\gamma \in G} |\gamma| > \delta(m + r)$. This is governed by a Peierls map, since all contours in G are mutually external, each contribute a weight of $W_k^{\text{rn}}(\gamma)$ per Lemma 3.2, which is at most $e^{-(\beta-5)|\gamma|}$ using Proposition 4.8. The number of choices for G having $\sum_{\gamma \in G} |\gamma| = l$ is at most $2^{\partial_i \Lambda_{m+r}} C^l$, while the probability of any fixed G is at most $\prod_{\gamma \in G} W_k^{\text{rn}}(\gamma) \leq e^{-(\beta-5)l}$, so the sum over $l \geq \delta(m + r)$ comes out to at most $e^{-\delta(\beta-C)(m+r)}$ for some universal C .

The probability that $\Lambda_m \not\subset \text{Int}(\mathcal{L})$ is bounded by the probability that some incident sequence of contours of G intersect both $\partial_i \Lambda_{m+r}$ and $\partial_o \Lambda_m$, which is seen to be at most $me^{-(\beta-C)r}$ by a union bound over the $4(m + r)$ sites where the sequence of incident contours could start, and then a Peierls map like the above to bound the probability of such a sequence of incident contours having size at least r . \blacksquare

5. SPATIAL MIXING PROPERTIES IN FINITE VOLUMES

In the previous section, we derived a sequence of critical external field values $(\lambda_c^{(k)})_k$ dictating transitions in which truncated free energy is largest, and showed that when $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$, we have rigidity for the height k boundary SOS model. Our aim is now to deduce spatial mixing properties when the boundary conditions start from a different height. Namely, we show that if we start with some arbitrary boundary condition, the distance between the boundary and the nearest loop of height- k sites has an exponential tail. Inside of the outermost such k -loop the surface will be rigid, and this will in particular be used to couple SOS measures with different boundary conditions away from their boundary.

This section is sensitive to the isoperimetry properties of the underlying domain and we henceforth restrict attention to boxes Λ_m . Also, we now need to allow consideration of non-constant boundary conditions, so we let $\phi \in \mathbb{Z}_+^{\partial_o V}$ denote general boundary conditions on V , unless otherwise specified the floors and ceilings will be at 0 and ∞ respectively everywhere, and we will use the shorthand notation μ_{ϕ, Λ_m} for this measure on Λ_m . The main theorem in this section is the following.

Theorem 5.1. *Let $\beta > \beta_0$, $k_m < a_m \leq \frac{m}{\log m}$ and $\lambda_m \in (\lambda_c^{(k_m)}, \lambda_c^{(k_m-1)})$. Consider the concentric boxes $\Lambda_{m/2} \subset \Lambda_m$. There is an absolute constant $C > 0$ such that for all $m \geq C\beta a_m / \mathfrak{d}(\lambda)$ (where $\mathfrak{d}(\lambda)$ is defined as in Eq. (1.2)) and all boundary conditions ϕ, ϕ' having $\|\phi\|_\infty \vee \|\phi'\|_\infty < a_m$, we have*

$$\left\| \mu_{\phi, \Lambda_m}(\varphi(\Lambda_{m/2}) \in \cdot) - \mu_{\phi', \Lambda_m}(\varphi(\Lambda_{m/2}) \in \cdot) \right\|_{\text{TV}} \leq e^{-\beta m / C}.$$

When λ is kept fixed independent of the domain size, we desire a stronger spatial mixing estimate, and towards that purpose need to allow for unbounded boundary conditions.

Corollary 5.2. *If for some $\varepsilon > 0$, $\lambda > \varepsilon$, then the maximum in Theorem 5.1 can be replaced by a supremum over all possible boundary conditions ϕ, ϕ' , so long as $m \geq C(\beta, \varepsilon)$ for $C(\beta, \varepsilon) > 0$ that is finite for all $\lambda > 0$ (it blows up as $\varepsilon \downarrow 0$).*

We will also prove a similar weak spatial mixing statement to Theorem 5.1 for annuli of side-length n and thickness m , with boundary conditions differing on the inner boundary: see Theorem 5.11.

In this section, we mix arguments based on the height function representation of the SOS model, and loops of height k , $\{\geq k\}$ and $\{\leq k\}$, with those based on the contour representation. This enables us to combine more percolation-theoretic arguments based on coarse-graining, together with the cluster expansion-based estimates of the previous sections, and is key to our results holding arbitrarily close to $\lambda = 0$.

5.1. k -contours in boxes with height $k \pm 1$. We first aim to establish that started from height $k \pm 1$, there will be a k -contour surrounding the concentric box of half the side-length in any box of side-length larger than $d^\mp(\lambda)$, which we recall are defined as

$$d^+(\lambda) = \min_k \{(\lambda_c^{(k)} - \lambda) : \lambda_c^{(k)} \geq \lambda\} \quad \text{and} \quad d^-(\lambda) = \min_k \{(\lambda - \lambda_c^{(k)}) : \lambda_c^{(k)} \leq \lambda\}.$$

(Recall also that $d(\lambda) = \min\{d^+(\lambda), d^-(\lambda)\}$.) Notice that if the boundary conditions on Λ_m are $k + 1$ and there is a down k -contour γ surrounding $\Lambda_{m'}$, then necessarily $\partial_i \text{Int}(\gamma)$ forms a loop (per Definition 4.11) of height- $\{\leq k\}$ sites. Similarly for $k - 1$ boundary conditions, an up k -contour, and height- $\{\geq k\}$ loops.

The following lemma will be essential to our spatial mixing arguments, and will get boosted into exponential tails to get to height k from any height, rather than just $k \pm 1$ using monotonicity and coarse-graining.

Lemma 5.3. *For any $\varepsilon > 0$, there is $\beta_0(\varepsilon)$, $\delta_0(\varepsilon)$, and absolute constant C_0 , such that the following holds. Suppose V is a simply-connected subset of Λ_m having $|\partial_e V| \leq 4(1 + \delta_0)m$, and $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$, with*

$$m \geq C_0 \beta / d^-(\lambda);$$

Then, for $\eta \in \{-, f\}$,

$$\mu_{\eta, k+1, V}(\exists \gamma \in \Gamma : h(\gamma) = k; \text{Int}(\gamma) \supset \Lambda_{(1-\varepsilon)m}) \geq 1 - e^{-\beta m / C}.$$

The same holds under $\mu_{\eta, k-1, V}$ for $\eta \in \{+, f\}$ if $m \geq C_0 \beta / d^+(\lambda)$ instead.

Proof. We prove the lemma with $k + 1$ -boundary conditions, the $k - 1$ case being symmetrical. For readability, we drop the boundary signing subscript η from the notation.

We consider the geometric properties of the family of large (at least $(\log m)^2$ length) k -contours, to show that with high probability one of them must confine the box $\Lambda_{(1-\varepsilon)m}$. Throughout this proof, we call a contour γ *macroscopic* if it is outermost and has $|\gamma| \geq (\log m)^2$. For a (standard) contour collection Γ , let Γ_k^{mac} be the collection of outermost macroscopic k -contours, and let Γ_k^{mic} be the set of outermost k -contours that are not macroscopic. We use the following basic isoperimetric bound from [10, Lemma 2.6].

Fact 5.4. *For every $\varepsilon > 0$, there exists $\delta(\varepsilon) > 0$ such that if $(\gamma_i)_i$ are any collection of mutually external geometric contours in Λ_m satisfying*

$$\sum_i |\gamma_i| \leq 4(1 + \delta)m, \quad \text{and} \quad \sum_i |\text{Int}(\gamma_i)| \geq (1 - \delta)m^2. \quad (5.1)$$

Then, if γ_1 is the one with the largest interior, $\text{Int}(\gamma_1)$ contains $\Lambda_{(1-\varepsilon)m}$.

Given Fact 5.4, it suffices to show that in a typical sample from $\mu_{\eta, k+1, V}$, the collection Γ_k^{mac} satisfies the two conditions of Eq. (5.1) for $\delta = \delta(\varepsilon)$. Towards this, we will prove Lemma 5.3 with $\delta_0 = \frac{1}{2}\delta(\varepsilon)$ and cover the complement of the event in Eq. (5.1) by the following bad events:

$$\mathcal{B}_1 := \left\{ \sum_{\gamma \in \Gamma_k^{\text{mac}}} |\gamma| > 4(1 + \delta)m \right\}, \quad (5.2)$$

$$\mathcal{B}_2 := \bigcup_{1 \leq i \leq 2 \log_2 \log m} \mathcal{B}_{2,i}, \quad \text{where} \quad \mathcal{B}_{2,i} = \left\{ |\{\gamma \in \Gamma_k^{\text{mic}} : |\gamma| \in [2^i, 2^{i+1}]\}| \geq \frac{1}{\sqrt{\beta}} e^{-(1.1)^i} m^2 \right\}, \quad (5.3)$$

$$\mathcal{B}_3 := \mathcal{B}_1^c \cap \mathcal{B}_2^c \cap \left\{ \sum_{\gamma \in \Gamma_k^{\text{mac}}} |\text{Int}(\gamma)| \leq (1 - \delta)m^2 \right\}. \quad (5.4)$$

Observe that on the intersection of $(\mathcal{B}_{2,i}^c)_{1 \leq i \leq 2 \log_2 \log m}$, we have

$$\sum_{\gamma \in \Gamma_k^{\text{mic}}} |\text{Int}(\gamma)| \leq \sum_i 2^{2(i+1)} \frac{1}{\sqrt{\beta}} e^{-(1.1)^i} m^2 \leq \frac{C}{\sqrt{\beta}} m^2,$$

which for β large, is at most $\delta m^2/2$ area confined by contours in Γ_k^{mic} . We show that each of $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ have small probability in Eqs. (5.2) to (5.4), which when summed yield Lemma 5.3.

Bound on \mathcal{B}_1 : too much length in macroscopic loops. We first show that for β large enough, $\delta_0 \leq \delta/2$,

$$\mu_{k+1,V}(\mathcal{B}_1) \leq e^{-(\beta-C)\delta m}. \quad (5.5)$$

Consider a fixed collection of mutually external macroscopic k -contours Γ_k^{mac} in \mathcal{B}_1 . The weight of configurations with this specific realization of Γ_k^{mac} can be written as

$$Z_{k+1,V}(\Gamma_k^{\text{mac}}) := \overline{Z}_{k+1,\text{Ext}(\Gamma_k^{\text{mac}})} \prod_{\gamma \in \Gamma_k^{\text{mac}}} e^{-\beta|\gamma|} Z_{-,k,\text{Int}(\gamma)},$$

where \overline{Z} indicates here that the sum is restricted to having no outermost macroscopic k -contours, and no contours confining any of the $\gamma \in \Gamma_k^{\text{mac}}$ (ensuring compatibility with Γ_k^{mac} indeed being the collection of outermost macroscopic k contours), and the boundary signing is that induced by $-$ on $\partial_i V$ and $+$ on $\partial_0 \text{Int}(\gamma)$ for $\gamma \in \Gamma_k^{\text{mac}}$. By inclusion, the first term is at most $Z_{k+1,\text{Ext}(\Gamma_k^{\text{mac}})}$. At the same time, the denominator in $\mu_{k+1,V}(\Gamma_k^{\text{mac}}) = Z_{k+1,V}(\Gamma_k^{\text{mac}})/Z_{k+1,V}$ evidently satisfies

$$Z_{k+1,V} \geq e^{-\beta|\partial_e V|} Z_{-,k,V} \geq e^{-\beta|\partial_e V|} Z_{k,\text{Ext}(\Gamma_k^{\text{mac}})} \prod_{\gamma \in \Gamma_k^{\text{mac}}} Z_{-,k,\text{Int}(\gamma)}. \quad (5.6)$$

Dividing through, and using $|\partial_e V| \leq 4(1 + \delta_0)m$ we get

$$\mu_{k+1,V}(\Gamma_k^{\text{mac}}) \leq e^{\beta(4(1+\delta_0)m - \sum_{\gamma \in \Gamma_k^{\text{mac}}} |\gamma|)} \frac{Z_{k+1,\text{Ext}(\Gamma_k^{\text{mac}})}}{Z_{k,\text{Ext}(\Gamma_k^{\text{mac}})}}.$$

By Corollary 4.10 applied to $U = \text{Ext}(\Gamma_k^{\text{mac}})$, whence $|\partial_e \text{Ext}(\Gamma_k^{\text{mac}})| \leq 4(1 + \delta_0)m + \sum_{\gamma \in \Gamma_k^{\text{mac}}} |\gamma|$,

$$\mu_{k+1,V}(\Gamma_k^{\text{mac}}) \leq e^{\beta(4(1+\delta_0)m - \sum_{\gamma} |\gamma|) + e^{-\beta/3}(4(1+\delta_0)m + \sum_{\gamma \in \Gamma_k^{\text{mac}}} |\gamma|)}.$$

At this point, we can bound $\mu_{k+1,V}(\mathcal{B}_1)$ by summing over the possible choices of $\Gamma_k^{\text{mac}} \in \mathcal{B}_1$. Let K denote the total number of contours in Γ_k^{mac} , let M denote $\sum_{\gamma \in \Gamma_k^{\text{mac}}} |\gamma|$, and note that $K \leq M/(\log m)^2$ since these are all macroscopic contours. Then bounding $\binom{m^2}{K} \leq e^{2K \log m}$, we have

$$\mu_{k+1,V}(\mathcal{B}_1) \leq \sum_{M \geq 4(1+\delta)m} \sum_{K \leq M/(\log m)^2} e^{2K \log m} 4^M e^{-(\beta-1)(M-4(1+\delta_0)m)}.$$

Since $K \leq M/(\log m)^2$, we have $e^{2K \log m} \leq e^M$, (for $m \geq 10$ say) and this is at most

$$\sum_{M \geq 4(1+\delta)m} e^M e^{-(\beta-1)(M-4(1+\delta_0)m)} \leq e^{4(1+\delta_0)m} \sum_{L \geq 4\delta_0 m} e^{-(\beta-C)L},$$

where we used that $\delta \geq 2\delta_0$. This implies the desired Eq. (5.5).

Bound on \mathcal{B}_2 : too much area in microscopic loops. Our next aim is to show that for every $i \leq 2 \log_2 \log m$,

$$\mu_{k+1,V}(\mathcal{B}_{2,i}) \leq e^{-(\beta-C)m^{3/2}}. \quad (5.7)$$

which when summed over $i \leq 2 \log_2 \log m$ bounds the probability of \mathcal{B}_2 by $e^{-(\beta-C)m^{4/3}}$, say.

For ease of notation, let $\Gamma_{k,i}^{\text{mic}} = \{\gamma \in \Gamma_k^{\text{mic}} : |\gamma| \in [2^i, 2^{i+1}]\}$ so that $\mathcal{B}_{2,i} = \{|\Gamma_{k,i}^{\text{mic}}| \geq \beta^{-1/2} e^{-(1.1)^i} m^2\}$. Consider a fixed collection of outermost non-elementary k -contours $\Gamma_k = (\Gamma_k^{\text{mic}}, \Gamma_k^{\text{mac}})$ belonging to $\mathcal{B}_{2,i}$. The total weight of configurations with a specific realization of $\Gamma_{k,i}^{\text{mic}}$, denoted $Z_{k+1,V}(\Gamma_{k,i})$ is

$$Z_{k+1,V}(\Gamma_{k,i}^{\text{mic}}) = \bar{Z}_{k+1, \text{Ext}(\Gamma_{k,i}^{\text{mic}})} \prod_{\gamma \in \Gamma_{k,i}^{\text{mic}}} e^{-\beta|\gamma|} Z_{-,k, \text{Int}(\gamma)}.$$

where \bar{Z} indicates that it has no non-elementary outermost down contours of size between $[2^i, 2^{i+1}]$ and also no non-elementary contours nesting any $\gamma \in \Gamma_{k,i}^{\text{mic}}$. By inclusion, the \bar{Z} term is at most $Z_{k+1, \text{Ext}(\Gamma_{k,i}^{\text{mic}})}$. Dividing out by $Z_{k+1,V}$ for which we have the lower bound analogous to Eq. (5.6),

$$\frac{Z_{k+1,V}(\Gamma_{k,i}^{\text{mic}})}{Z_{k+1,V}} \leq e^{4\beta(1+\delta_0)m - \beta \sum_{\gamma \in \Gamma_{k,i}^{\text{mic}}} |\gamma|} \frac{Z_{k+1, \text{Ext}(\Gamma_{k,i}^{\text{mic}})}}{Z_{k, \text{Ext}(\Gamma_{k,i}^{\text{mic}})}}.$$

By Corollary 4.10, this is at most

$$e^{\beta(4(1+\delta_0)m - \sum_{\gamma \in \Gamma_{k,i}^{\text{mic}}} |\gamma|) + e^{-\beta/3} |\partial_e \text{Ext}(\Gamma_{k,i})|}.$$

Now observe that on the event $\mathcal{B}_{2,i}$,

$$\sum_{\gamma \in \Gamma_{k,i}^{\text{mic}}} |\gamma| \geq 2^i \beta^{-1/2} e^{-(1.1)^i} m^2 = \beta^{-1/2} e^{i \log 2 - (1.1)^i} m^2.$$

For all $i \leq 2 \log_2 \log m$, this is at least $m^{3/2}$, say, so long as $m \geq C\beta^{1/4}$. Now to sum over the possible choices of $\Gamma_{k,i}^{\text{mic}}$, let $\sum_{\gamma \in \Gamma_{k,i}^{\text{mic}}} |\gamma| = M$, let K count the number of contours in $\Gamma_{k,i}^{\text{mic}}$, and note the bound $|\partial_e \text{Ext}(\Gamma_{k,i})| \leq 4(1+\delta_0)m + M$. We then obtain

$$\mu_{k+1,V}(\mathcal{B}_{2,i}) \leq \sum_{M \geq \beta^{-1/2} e^{i \log 2 - (1.1)^i} m^2} \sum_{K \leq 2^{-i} M} \binom{m^2}{K} C^M e^{-(\beta-1)(M-4(1+\delta_0)m)}.$$

Letting $2^{-i} M = \rho m^2$, and using the fact that $M \geq m^{3/2}$, this is at most

$$\sum_{\beta^{-1/2} e^{-(1.1)^i} \leq \rho \leq 1} (\rho m^2) e^{\rho m^2 (1 + \log \frac{2}{\rho})} e^{-(\beta-C)M}.$$

Since $\rho \geq \beta^{-1/2} e^{-(1.1)^i}$ we can upper bound $\log(2/\rho)$ and use $2^i \beta \rho m^2 \geq 4\rho m^2 (1.1)^i \log \beta$ to see that the second exponential dominates the first, whence for all $m \geq C\beta^{1/4}$, we get

$$\mu_{k+1,V}(\mathcal{B}_{2,i}) \leq e^{-(\beta-C)e^{-(1.1)^i} m^2}.$$

Using that $e^{-(1.1)^i} m^2 \geq m^{3/2}$ for all $i \leq 2 \log_2 \log m$ then yields the claimed Eq. (5.7).

Bound on \mathcal{B}_3 : too little area in macroscopic loops. For the last bound, let

$$\Delta_{\text{tr}} = f_k^{\text{tr}} - f_{k+1}^{\text{tr}}. \quad (5.8)$$

which is positive and at least $(1 - \varepsilon_\beta) d^-(\lambda)$ per Corollary 4.7 for $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)}]$. The goal of this last part is to establish the following bound:

$$\mu_{k+1,V}(\mathcal{B}_3) \leq e^{-\beta m/C}. \quad (5.9)$$

On \mathcal{B}_2^c , the total boundary length confined by outer contours of size at least Δ_{tr}^{-1} is at most

$$\sum_{i \geq -\log_2 \Delta_{\text{tr}}} 2^{i+1} \beta^{-1/2} e^{-(1.1)^i} m^2.$$

We can absorb the 2^{i+1} prefactor by changing this to $\beta^{-1/2}e^{-(1.05)^i}$, and summing that out, we get that $\sum_{i \geq -\log_2 \Delta_{\text{tr}}} \sum_{\gamma \in \Gamma_{k,i}^{\text{mic}}} |\gamma| \leq \beta^{-1/2} m^2 e^{-\Delta_{\text{tr}}^{-\alpha}}$, for some universal constant $\alpha > 0$.

Now consider any realization of $G_k = ((\Gamma_{k,i}^{\text{mic}})_{i \geq -\log_2 \Delta_{\text{tr}}}, \Gamma_k^{\text{mac}})$ belonging to \mathcal{B}_3 . The partition function associated to these contour families is

$$Z_{k+1,V}(G_k) = \overline{Z}_{k+1, \text{Ext}(G_k)} \prod_{\gamma \in G_k} e^{-\beta|\gamma|} Z_{-,k, \text{Int}(\gamma)},$$

where \overline{Z} here indicates that there are no outermost k -contours of size larger than Δ_{tr}^{-1} , and none of the ones in G_k are surrounded by another contour. As such, for every contour γ appearing in the sum corresponding to \overline{Z} , it has $\text{diam}(\gamma) \leq |\gamma|/2 \leq \Delta_{\text{tr}}^{-1}$. By Corollary 4.9, this implies that for all contours appearing in \overline{Z} , they have $W_{k+1}^{\text{rn}}(\gamma) = W_{k+1}^{\text{tr}}(\gamma)$. By inclusion, then,

$$Z_{k+1,V}(G_k) \leq Z_{k+1, \text{Ext}(G_k)}^{\text{tr}} \prod_{\gamma \in G_k} e^{-\beta|\gamma|} Z_{-,k, \text{Int}(\gamma)}.$$

Dividing by $Z_{k+1,V}$, on which we have the lower bound analogous to Eq. (5.6), we get

$$\mu_{k+1,V}(G_k) \leq e^{4\beta(1+\delta_0)m - \beta \sum_{\gamma \in G_k} |\gamma|} \frac{Z_{k+1, \text{Ext}(G_k)}^{\text{tr}}}{Z_{k, \text{Ext}(G_k)}}.$$

By Proposition 4.8, $Z_{k, \text{Ext}(G_k)} = Z_{k, \text{Ext}(G_k)}^{\text{tr}}$ since $\lambda \in [\lambda_c^{(k)}, \lambda_c^{(k-1)}]$, so by Lemma 4.5,

$$\mu_{k+1,V}(G_k) \leq e^{\beta(4(1+\delta_0)m - \sum_{\gamma \in G_k} |\gamma|)} e^{-\Delta_{\text{tr}} |\text{Ext}(G_k)| + 2e^{-\beta/2} |\partial_e \text{Ext}(G_k)|}.$$

On the event \mathcal{B}_3 , we must have $|\text{Ext}(G_k)| \geq \delta m^2/2$ as under \mathcal{B}_2^c at most $\delta m^2/2$ area is confined by $\bigcup_i \text{Int}(\Gamma_{k,i}^{\text{mic}})$ and on \mathcal{B}_3 at most $(1-\delta)m^2$ is confined in Γ_k^{mac} . Thus,

$$\mu_{k+1,V}(G_k) \leq e^{4(\beta+1)(1+\delta_0)m - (\beta-1) \sum_{\gamma \in G_k} |\gamma|} e^{-\Delta_{\text{tr}} \delta m^2/2} \leq e^{(4\beta+C)m - \Delta_{\text{tr}} \delta m^2/2}.$$

Now consider the entropy over the choices of G_k that belong to \mathcal{B}_3 . Since $\mathcal{B}_3 \subset \mathcal{B}_1^c$, for the macroscopic contours, the total number of choices of Γ_k^{mac} in \mathcal{B}_1^c is evidently at most

$$\binom{m^2}{4(1+\delta_0)m/(\log m)^2} 4^{4(1+\delta_0)m} \leq \exp(10m).$$

For the microscopic contours, since $\mathcal{B}_3 \subset \mathcal{B}_2^c$, for each i , there are at most $\beta^{-1/2}e^{-(1.1)^i} m^2$ many contours in $\Gamma_{k,i}^{\text{mic}}$, each of size at most 2^{i+1} ; thus, the number of choices is at most

$$\prod_{i=-\log \Delta_{\text{tr}}}^{2 \log_2 \log n} e^{\beta^{-1/2}e^{-(1.1)^i} m^2 (1+(1.1)^i + \log \beta)} 4^{\beta^{-1/2} 2^{i+1} e^{-(1.1)^i} m^2} = \exp\left(\beta^{-1/3} m^2 \sum_{i=-\log \Delta_{\text{tr}}}^{2 \log_2 \log n} 2^i e^{-(1.1)^i}\right)$$

Since the sum is dominated by the smallest $i = -\log \Delta_{\text{tr}}$, the count above is at most $\exp(\beta^{-1/3} e^{-\Delta_{\text{tr}}^{-\alpha}} m^2)$ for some universal constant $\alpha > 0$. Therefore, summing up over the choices of $G_k \in \mathcal{B}_3$, we get

$$\mu_{k+1,V}(\mathcal{B}_3) \leq \exp\left(-\left(\frac{1}{2} \Delta_{\text{tr}} \delta - \beta^{-1/3} e^{-\Delta_{\text{tr}}^{-\alpha}}\right) m^2 + (4\beta + C)m\right).$$

We have the naive upper bounds of $\Delta_{\text{tr}} \leq \max\{\lambda, (1+\varepsilon_\beta)e^{-4\beta}\}$ on Δ_{tr} , we are really only interested in the behavior when $\Delta_{\text{tr}} \leq 1$, say. Since α is universal (independent of β, λ), it is easy to check that for large β (depending only on α) we have $\beta^{-1/3} e^{-\Delta_{\text{tr}}^{-\alpha}} \leq \Delta_{\text{tr}}/4$. As such, we deduce that

$$\mu_{k+1,V}(\mathcal{B}_3) \leq \exp\left(-\frac{1}{4} \Delta_{\text{tr}} \delta m^2 + (4\beta + C)m\right).$$

This gives Eq. (5.9) so long as $m \geq 20\beta\delta^{-1}/d^-(\lambda)$. ■

5.2. The outermost k -loop in general domains with $k \pm 1$ boundary conditions. We next use Lemma 5.3 to show that in general domains (not necessarily boxes) with $k \pm 1$ boundary conditions, the distance of the outermost loop of height- k sites to the boundary will have an exponential tail. For any domain V , define the shrinkage of V by r as

$$S_r(V) = \{v \in V : d(v, \partial_i V) \geq r\}.$$

For $V : S_r(V) \supset \Lambda_{m/2}$, let $\mathcal{L}_{\leq k}$ be the outermost loop of $\{\leq k\}$ sites surrounding the origin.

Lemma 5.5. *Let V be any simply connected set having $\partial \Lambda_{m/2} \subset V \subset \partial \Lambda_m$. For every $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$, if $r \geq C_0 \beta / d^-(\lambda)$ and $m \geq 2r$, then for $\eta \in \{-, f\}$*

$$\mu_{\eta, k+1, V}(\mathcal{L}_{\leq k} \subset V \setminus S_r(V)) \geq 1 - |\partial_i S_r(V)| e^{-\beta r / C}.$$

The same holds under $\mu_{\eta, k-1, V}$ for $\eta \in \{+, f\}$ if $r \geq C_0 \beta / d^+(\lambda)$.

Proof. We prove one direction, the other direction being symmetrical. For every x at distance r from ∂V , let $B_x = B_r(x)$ be the ball of radius r about x , and let $A_x = B_r(x) \setminus B_{r/2}(x)$ be the annulus between B_x and its concentric box of half the radius. By construction, B_x, A_x are interior to V . Let E_x be the event that there is a loop of $\{\leq k\}$ -sites in the annulus A_x . If E_x occurs for all $x \in \partial_i S_r(V)$, then $\mathcal{L}_{\leq k} \subset V \setminus S_r(V)$.

Therefore, the complement of that is a subset of the event $\bigcup_{x \in \partial_i S_r(V)} E_x^c$, and it suffices to upper bound the probability of this union. We use a union bound for this and consider any fixed x . For fixed $x \in \partial_i S_r(V)$, of which there are at most m^2 , consider the probability of E_x^c . Let A'_x be the annulus $B_r(x) \setminus B_{3r/4}(x)$ and first consider the event $E_{x,0}$ that there is an outermost loop of $\{\leq k+1\}$ -height sites \mathcal{C}_x in A'_x having total length $|\mathcal{C}_x| \leq 4(1 + \delta_0)r$. We use the bound

$$\mu_{k+1, V}(E_x^c) \leq \mu_{k+1, V}(E_{x,0}^c) + \mu_{k+1, V}(E_x | E_{x,0}). \quad (5.10)$$

For the first probability in Eq. (5.10), note that the event $E_{x,0}^c$ is an increasing event, so its probability is only larger by decreasing λ to $\lambda_c^{(k)}$. Upon doing so, the probability of $E_{x,0}^c$ is upper bounded by Lemma 4.12 whereby it is at most $e^{-\delta_0 \beta r / 4}$ for r larger than an absolute constant.

For the second probability in Eq. (5.10), we can condition on the outermost loop of $\{\leq k+1\}$ sites, \mathcal{C}_x , and note that it being measurable with respect to its exterior, we get

$$\mu_{k+1, V}(E_x^c | E_{x,0}) \leq \max_{\mathcal{C}_x} \max_{\phi \leq k+1} \mu_{\phi, \text{Int}(\mathcal{C}_x)}(E_x^c),$$

where the maximum over \mathcal{C}_x runs over those confining $B_{3r/4}$ in their interior and having length at most $4(1 + \delta_0)r$. Since E_x^c is an increasing event, we can take the maximal boundary conditions $k+1$ on $\text{Int}(\mathcal{C}_x)$, at which point Lemma 5.3 with $\varepsilon = 1/2$ (so that δ_0 above is $\delta_0(\varepsilon = 1/2)$) provides an upper bound of $e^{-\beta r / C}$ so long as $r \geq C_0 \beta / d^-(\lambda)$ and β large. Combining the above two estimates and taking a union bound over the $|\partial_i S_r(V)|$ many possible choices of x gives the claimed bound with a different absolute C . ■

5.3. The case of general boundary heights. Our aim is now to show that if the boundary conditions are at any height between 0 and $m / \log m$, the interface will, with high probability, contain a loop of height- k sites surrounding the origin when $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$. The argument essentially goes by repeatedly using monotonicity to apply Lemma 5.5 to each level from the boundary condition height to k one at a time.

Lemma 5.6. *Suppose $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$, $k+1 \leq j \leq \frac{m}{\log m}$, and $m \geq 2C_0 j \beta / d^-(\lambda)$. For $\eta \in \{-, f\}$,*

$$\mu_{\eta, j, \Lambda_m}(\mathcal{L}_{\leq k} \subset \Lambda_m \setminus \Lambda_{3m/4}) \geq 1 - e^{-\beta m / C}.$$

The analogue holds for $\eta \in \{+, f\}$, $0 \leq j \leq k-1$, and $\mathcal{L}_{\geq k}$ replacing $\mathcal{L}_{\leq k}$ if $m \geq 2C_0 k \beta / d^+(\lambda)$.

Proof. We show the bound supposing that $j \geq k+1$, the other direction being analogous. Let \mathcal{C}_i denote $\text{Int}(\mathcal{L}_{\leq i})$ for all i . Let r_i be the minimal r such that $\mathcal{L}_{\leq i-1}$ is a subset of $\mathcal{C}_i \setminus S_r(\mathcal{C}_i)$, in other words the maximal distance from $\mathcal{L}_{\leq i}$ that $\mathcal{L}_{\leq i-1}$ gets. In order for $\mathcal{L}_{\leq k}$ to not be a subset of $\Lambda_m \setminus \Lambda_{3m/4}$, it must be the case that $\sum_{i=k}^{j-1} r_i$ is at least $m/4$. We fix any possible sequence $(r_i)_i$ with $r_i \geq 0$ for all i and such that

$\sum_{i=k}^{j-1} r_i = m/4$. Since $j \leq m$, there are at most 2^{2m} many such choices. Since $m \geq 2C_0 j \beta / d^-(\lambda)$, there must be some subsequence of these (r_i) , call them $(r_{i_l})_l$, such that they are each larger than $C_0 \beta / d^-(\lambda)$ and whose sum exceeds $m/8$.

We now upper bound the probability that for all l , we have $\mathcal{L}_{\leq i_l}$ intersects $S_{r_{i_l}}(\mathcal{C}_{i_l+1})$. We do so iteratively from the boundary inwards. Namely, for fixed such $(r_{i_l})_l$, the probability of interest is at most

$$\prod_l \mu_{\eta, j, \Lambda_m}(\mathcal{L}_{\leq i_l} \cap S_{r_{i_l}}(\mathcal{C}_{i_l+1}) \neq \emptyset \mid \varphi(\Lambda_m \setminus \mathcal{C}_{i_l+1})).$$

(Here we used the fact that \mathcal{C}_{i_l+1} is measurable with respect to the height values on its exterior.) The conditional measure is then exactly $\mu_{-, i_l+1, \mathcal{C}_{i_l+1}}$. By monotonicity, since the event in question is an increasing event, the probability is only larger if we lift the floor up to $i_l - k$ whence the measure will be an additive translation of $\mu_{-k+1, \mathcal{C}_{i_l+1}}$ with $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$. At that point, the probability we are interested in is exactly that bounded by Lemma 5.5, whence it is at most $|\partial_i S_r(\mathcal{C}_{i_l+1})| e^{-r_{i_l}/C}$. Performing the sum over all possible choices of r_{i_l} , and plugging this in, we get

$$\mu_{j, \Lambda_m}(\mathcal{C}_k \cap \Lambda_{3m/4} \neq \emptyset) \leq 4^m (m^2)^j e^{-\beta m/C}.$$

So long as $j \leq m/(\log m)$, this will be at most $e^{-\beta m/C}$ for a different C . ■

Remark 5.7. At this point, one can see Theorem 1.4 by applying (the second part of) Lemma 5.6 with $j = 0$, revealing the outermost loop $\mathcal{L}_{\geq k}$, then using the rigidity at height k interior to $\mathcal{L}_{\geq k}$ per the estimates of Section 4.3 to see that interior to $\mathcal{L}_{\geq k}$, most sites are exactly at height k when $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$.

In the case where $\lambda > 0$ is kept independent of the box size, the external field is strong enough to allow us to consider all possible boundary conditions (unbounded) and still obtain spatial mixing. For this purpose, the following lemma shows that the surface drops to some finite height quickly even started arbitrarily high.

Lemma 5.8. *For every $\lambda > 0$, if $m \geq C\lambda^{-1}$, the outermost loop of height $\{\leq C\lambda^{-2}\}$ in Λ_m confines $\Lambda_{(1-\varepsilon)m}$ in its interior except with probability e^{-cm} .*

Proof. It is fairly straightforward to check that as long as a box has side-length L at least $L_0 := \lceil 8/\lambda \rceil$ the values of φ will have finite moments uniformly over all (unbounded) boundary conditions: see [11, Proposition 3.2, item (i)] for a short self-contained argument. In particular, in a domain V , for any x at distance at least L_0 away from ∂V , for every s there exists a constant $C(\beta, s) > 0$ such that for all $\lambda > 0$,

$$\sup_{\phi} \mu_{\phi, V}[\varphi_x^s] \leq \lambda^{-s} C(\beta, s).$$

Now for simplicity, consider the domain Λ_m , and tile it by overlapping boxes of side-length $2L_0 \times 2L_0$ that overlap on half of their area (i.e., centered at the vertices of $\Lambda_m \cap L_0 \mathbb{Z}^2$).

Consider any such box B_x . Uniformly over the boundary conditions on that box, by Markov's inequality, and a union bound, the probability that there is no $\{\leq L_0 \lambda^{-1}\}$ loop around the concentric $L_0 \times L_0$ box is at most $C(\beta, s) L_0^{-s+1}$. Call such a box bad, and call a box good if there is such a loop. The set of bad blocks form a 1-dependent percolation in the graph induced by the boxes, with two boxes called adjacent if they intersect. By classical reasoning (see e.g., [35]), this will be stochastically dominated by an independent percolation with a parameter that can be made arbitrarily small by taking L_0 large, from which we deduce that the probability of a path of bad blocks from $\partial \Lambda_m$ to $\partial \Lambda_{(1-\varepsilon)m}$ is exponentially unlikely in εm . ■

5.4. From $\{\geq k\}$ and $\{\leq k\}$ loops to spatial mixing. We now wish to use Lemma 5.6 to deduce weak spatial mixing amongst the set of boundary conditions ranging from 0 to some upper bound a_m (and in the case where λ is bounded away from zero, any arbitrary boundary conditions). The following lemma boosts the probability of two configurations individually having k -loops to the joint probability of having a common k -loop (a loop of sites whose height is k in both configurations).

Lemma 5.9. *Suppose $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$. For any two boundary conditions ϕ, ϕ' and any coupling of $\varphi \sim \mu_{\phi, \Lambda_m}$ and $\varphi' \sim \mu_{\phi', \Lambda_m}$, the probability that (φ, φ') don't have a common k -loop surrounding $\Lambda_{(1-2\varepsilon)m}$ is at most the sum of the marginal probabilities of not having a k -loop surrounding $\Lambda_{(1-\varepsilon)m}$, plus $e^{-\beta\varepsilon m/C}$.*

Proof. Let \mathcal{E}_k and \mathcal{E}'_k be the events that there is a loop of height- k sites surrounding $\Lambda_{(1-\varepsilon)m}$ in φ and φ' respectively. Let \mathcal{S}_k be the event that there is a common k -loop in the annulus $\Lambda_{(1-\varepsilon)m} \setminus \Lambda_{(1-2\varepsilon)m}$. The desired probability is at most

$$\mu_{\phi, \Lambda_m}(\mathcal{E}_k^c) + \mu_{\phi', \Lambda_m}((\mathcal{E}'_k)^c) + \mathbb{P}(\mathcal{E}_k, \mathcal{E}'_k, \mathcal{S}_k^c).$$

It suffices to bound the third probability. The event \mathcal{S}_k^c implies the existence of a path \mathcal{P} from $\partial_i \Lambda_{(1-\varepsilon)m}$ to $\partial_o \Lambda_{(1-2\varepsilon)m}$ such that for all $x \in \mathcal{P}$ either x is interior to a contour inside the outermost k -loop of φ , or interior to a contour inside the outermost k -loop of φ' . In particular, if we consider the collection of all outermost contours inside the outermost k -loop of φ , and color them RED and do the same for φ' and color them BLUE, there must be a collection of RED and BLUE contours whose union is a connected subset \mathcal{P} of Λ_m^* connecting $\partial_i \Lambda_{(1-\varepsilon)m}$ to $\partial_o \Lambda_{(1-2\varepsilon)m}$. Let $M \geq \varepsilon m$ denote the total number of edges in \mathcal{P} , there being at most $4m3^M$ many choices for it, and another 3^M many choices of for choosing whether each edge in \mathcal{P} was colored RED, BLUE, or both. There must then be a subset of \mathcal{P} , call it $\tilde{\mathcal{P}}$ which has at least $M/2$ many edges, and all of whose edges were colored RED, or all of whose edges were colored BLUE. This gives us the following bound:

$$\mathbb{P}(\mathcal{E}_k, \mathcal{E}'_k, \mathcal{S}_k^c) \leq \sum_{M \geq \varepsilon m} 4m3^{2M} \max_{\mathcal{P}: |\mathcal{P}|=M} 2^M \max_{\tilde{\mathcal{P}} \subset \mathcal{P}: |\tilde{\mathcal{P}}| \geq M/2} (\mu_{\phi, \Lambda_m}(\mathcal{E}_k, E_{\tilde{\mathcal{P}}}) + \mu_{\phi', \Lambda_m}(\mathcal{E}'_k, E'_{\tilde{\mathcal{P}}})) .$$

where we are using $E_{\tilde{\mathcal{P}}}, E'_{\tilde{\mathcal{P}}}$ to denote the event that the contours in $\tilde{\mathcal{P}}$ are all outermost contours interior to \mathcal{L}_k in φ or φ' respectively.

Fix M , a colored collection of contours \mathcal{P} , a subset $\tilde{\mathcal{P}}$ forming mutually external contours, and w.l.o.g. consider the first of the two probabilities on the right-hand side above. We can condition on the outermost k -loop in μ_{ϕ, Λ_m} , call its interior \mathcal{C}_k , this being measurable with respect to its exterior and inducing boundary conditions k on its interior. The bound on the probability of $E_{\tilde{\mathcal{P}}}$ then becomes a standard consequence of the rigidity at height k in \mathcal{C}_k . Namely, by the Peierls map that deletes the contours making up $\tilde{\mathcal{P}}$ from the renormalized contour collection for $\varphi(\mathcal{C}_k)$, and using Proposition 4.8 to bound the weight change, uniformly over \mathcal{C}_k , we have

$$\mu_{\emptyset, k, \mathcal{C}_k}(E_{\tilde{\mathcal{P}}}) \leq e^{-(\beta-5)|\tilde{\mathcal{P}}|} \leq e^{-(\beta-5)M/2} .$$

Plugging this bound in to the earlier sum, together with its equivalent for φ' , implies $\mathbb{P}(\mathcal{E}_k, \mathcal{E}'_k, \mathcal{S}_k^c)$ is at most $\exp(-\beta\varepsilon m/C)$ as claimed. \blacksquare

Corollary 5.10. *If $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$ and $k \vee \|\phi\|_\infty < a_m \leq \frac{m}{\log m}$, as long as $m \geq 2C_0 a_m \beta / d(\lambda)$,*

$$\mu_{\phi, \Lambda_m}(\mathcal{L}_k \subset \Lambda_m \setminus \Lambda_{2m/3}) \geq 1 - e^{-\beta m/C} ,$$

If $\lambda > \varepsilon$ for some $\varepsilon > 0$ fixed, then we can drop the upper bound on $\|\phi\|_\infty$, at the expense of also assuming m is at least some constant $C(\varepsilon) > 0$.

Proof. Consider φ, φ' being drawn from the SOS measures with minimal and maximal boundary conditions $\phi \equiv a_m$ and $\phi' \equiv 0$, and couple them using the grand monotone coupling. For each of φ, φ' , the individual probability of not having a $\{\leq k\}$ or $\{\geq k\}$ -loop surrounding $\Lambda_{3m/4}$ is bounded by $e^{-\beta m/C}$ per Lemma 5.6 (applicable to both since $d(\lambda) = d^-(\lambda) \wedge d^+(\lambda)$). Interior to the outermost such loop, the rigidity at height k as argued in Lemma 4.12 implies that \mathcal{L}_k will in fact contain $\Lambda_{2m/3}$ in its interior, except with probability $e^{-\beta m/C}$. Applying Lemma 5.9, they will therefore have a common k -loop except with probability $e^{-\beta m/C}$ for some other C . Since we are using the grand monotone coupling, if φ, φ' share a common k -loop, then so do the coupled draws from μ_{ρ, Λ_m} for all boundary conditions $\rho : \|\rho\|_\infty \leq a_m$, as desired.

When $\lambda > 0$ independently of m , we use the grand coupling of all possible boundary conditions, noting that if a constant boundary condition at ϕ^+ is coupled to the one with boundary conditions at 0, then all boundary conditions in between also share a k -loop. For those two, first use Lemma 5.8 to get a $\{\leq C\lambda^{-2}\}$ -height loop surrounding $\Lambda_{4m/5}$ uniformly over ϕ^+ , then apply the above reasoning interior to that. ■

With the above, we are able to deduce the claimed spatial mixing bound on boxes.

Proof of Theorem 5.1. For any two boundary conditions ϕ, ϕ' having $\|\phi\|_\infty \vee \|\phi'\|_\infty < a_m$ inducing configurations φ, φ' , we construct a coupling such that they agree on $\Lambda_{m/2}$ except with exponentially small probability. First, expose (under an arbitrary coupling) from the outside in, their outermost common k -loop surrounding $\Lambda_{m/2}$ if it exists. By Corollary 5.10 together with Lemma 5.9, such a shared k -loop exists except with probability $e^{-\beta m/C}$. On the event of existence of such a shared k -loop, this being measurable with respect to the randomness exterior to it, we can use the identity coupling of the two configurations on their interior so that they agree with probability 1 on $\Lambda_{m/2}$. ■

Proof of Corollary 5.2. The proof is the same as the above proof of Theorem 5.1 except with the application of Corollary 5.10 using its second part where we allow arbitrary, unbounded, boundary conditions. ■

5.5. Spatial mixing in an annulus. We conclude the section by using the above to also show a type of weak spatial mixing result in annuli, showing that if two boundary conditions differ on the inner boundary of an annulus, they can be coupled with high probability in an annulus of half the thickness close to the outer boundary. For ease of notation, let $A_{\bar{m},m}$ be the annulus $\Lambda_{\bar{m}} \setminus \Lambda_{\bar{m}-m}$.

Theorem 5.11. *Suppose that $\frac{\bar{m}}{2} \geq m \geq 4C_0\beta a_m/d(\lambda)$, and $\log \bar{m} \leq a_m \leq \frac{m}{\log m}$. Then for any pair of boundary conditions ϕ, ϕ' on $\partial_i \Lambda_{\bar{m}-m}$, having $\|\phi\|_\infty, \|\phi'\|_\infty < a_m$, we have*

$$\|\mu_{(0,\phi),A_{\bar{m},m}}(\varphi(A_{\bar{m},m/2}) \in \cdot) - \mu_{(0,\phi'),A_{\bar{m},m}}(\varphi(A_{\bar{m},m/2}) \in \cdot)\|_{\text{TV}} \leq 4\bar{m}(e^{-\beta m/C} + e^{-\beta a_m/C}).$$

where $(0, \phi)$ boundary conditions are 0 on $\partial_o \Lambda_{\bar{m}}$ and ϕ on $\partial_i \Lambda_{\bar{m}-m}$.

Proof. Let φ, φ' be independent draws from the two relevant distributions. We wish to show that with high probability, there will be a common loop of height k sites in the annulus $A_{\bar{m},3m/4} \setminus A_{\bar{m},m}$. If we expose the *innermost* shared k -loop in the annulus, it being measurable with respect to its interior, we can use the identity coupling exterior to the common k -loop to bound the total variation distance by the probability of non-existence of such a common k -loop in $A_{\bar{m},3m/4} \setminus A_{\bar{m},m}$.

The existence of such a common height- k loop will follow if around every vertex in $\Lambda_{\bar{m}-3m/4}$, there is a common k -loop in the ball $B_{v,m/4}$ surrounding that vertex. For any $v \in \partial_i \Lambda_{\bar{m}-3m/4}$, consider the probability of non-existence of a common k -loop in $B_{v,m/4}$ surrounding v .

First, expose the two configurations φ and φ' on the complement of $B_{v,m/4}$. Except with probability $e^{-\beta a_m/C}$, the maximal height of the configurations revealed on $\partial_o B_{v,m/4}$ will be $2a_m$. In order to see this, observe that the configuration is stochastically below the one that has all its boundary conditions at height a_m on the annulus at $\lambda = 0$, for which the tail bound on the maximum height from Eq. (2.4) applies since $a_m \geq \log \bar{m}$. Given those boundary conditions, we can apply Corollary 5.10 and Lemma 5.9 to deduce that the two configurations will share a k -loop surrounding v in $B_{v,m/4}$ except with probability $e^{-\beta m/C}$. A union bound over the at most $4\bar{m}$ vertices in $\partial \Lambda_{\bar{m}-3m/4}$ concludes the proof. ■

6. MIXING TIME UPPER BOUNDS

We now turn to the dynamical part of the paper, starting with the proofs of our mixing time upper bounds. Our focus is on the mixing time of the Glauber dynamics on the $n \times n$ box Λ_n . In order to have a finite state space, so that the mixing time is finite, we introduce a ceiling, and as is convention for such models, we will take the ceiling height to also be n . The floor will always be fixed to 0, but the ceiling will be variable in this section, and therefore it will be useful to use the notation Λ_n^ℓ for the domain Λ_n with ceiling at height ℓ .

The main content of this section is establishing the upper bound on the mixing time when $d(\lambda)$ is going to zero with n . The proof goes in two stages:

- (1) Reducing the mixing time on Λ_n with ceiling n to the mixing time on Λ_n with $\log n$ ceiling.
- (2) Reducing the mixing time on Λ_n with $\log n$ ceiling to the maximum of the mixing times on boxes of size n_0 and an annulus of width n_0 , with arbitrary boundary conditions, for $n_0 = \tilde{\Theta}(d(\lambda))$.

Finally, we apply a standard bound on the mixing time on boxes/annuli with cut-width n_0 , by an exponential in n_0 (possibly times a $\log n$ factor). Together, these arguments take up Sections 6.2 to 6.4.

In Section 6.5, we justify the $O(1)$ inverse gap bound when d is fixed independent of n ; given the weak spatial mixing with arbitrary boundary conditions for fixed λ in Corollary 5.2 this is essentially reproducing the argument of [12]. In Section 6.6, we explain the minor modifications needed for these upper bounds to apply to the torus, instead of zero boundary conditions.

6.1. Mixing time preliminaries. Let us recall some basics from the study of mixing times for Markov chains with finite state spaces; we refer the reader to [34] for more on the topic.

Consider a state space Ω , with a transition kernel $P(x, y)$ describing the rate at which a continuous-time Markov chain jumps from x to y , reversible with respect to a stationary distribution μ . Let $(X_t^{x_0})_{t \geq 0}$ be the corresponding Markov chain initialized from state x_0 , and let $L = (I - P)$ be its infinitesimal generator.

Let gap denote the gap in the spectrum of $-L$, i.e., the size of its smallest non-zero eigenvalue, and let t_{MIX} be its mixing time, i.e.,

$$t_{\text{MIX}}(\varepsilon) = \inf\{t > 0 : \max_{x_0 \in \Omega} \|\mathbb{P}(X_t^{x_0} \in \cdot) - \mu\|_{\text{TV}} < \varepsilon\}.$$

By convention, $t_{\text{MIX}} = t_{\text{MIX}}(1/4)$. The ε -mixing time satisfies a useful sub-multiplicativity property that $t_{\text{MIX}}(\varepsilon) \leq t_{\text{MIX}} \log_2 \frac{2}{\varepsilon}$. The inverse of the spectral gap, sometimes called the relaxation time, is closely tied to the mixing time. Namely,

$$(\text{gap}^{-1} - 1) \log\left(\frac{1}{2\varepsilon}\right) \leq t_{\text{MIX}}(\varepsilon) \leq \text{gap}^{-1} \log\left(\frac{1}{\varepsilon \mu_{\min}}\right). \quad (6.1)$$

The SOS Glauber dynamics for an SOS distribution μ of the form of Eq. (2.1) has transition kernel $P(\varphi, \varphi')$ which is non-zero only if φ, φ' differ at exactly one vertex, say v , by exactly one height at that vertex. In that case, if $\varphi^{(v, \iota)}$ denotes the configuration φ with the value at v changed to $\varphi_v + \iota$,

$$P(\varphi, \varphi') = \frac{\mu(\varphi' \mid (\varphi_w)_{w \neq v})}{\sum_{\iota \in \{-1, 0, 1\}} \mu(\varphi^{(v, \iota)} \mid (\varphi_w)_{w \neq v})}$$

This is easily checked to be reversible with respect to μ .

The monotonicity of Lemma 2.1 implies monotonicity for the SOS Glauber dynamics, whereby if $x_0 \geq y_0$, then $X_t^{x_0} \geq X_t^{y_0}$ for all $t \geq 0$. In particular, there is a *grand coupling* of all the SOS dynamics chains from all the possible initializations, such that with probability 1, $X_t^{x_0} \geq X_t^{y_0}$ for all $t \geq 0$ and all $x_0 \geq y_0$.

6.2. Reducing the mixing time to $\log n$ ceiling. Let us introduce the notation

$$T_{\Lambda_n^\ell} := t_{\text{MIX}}(\Lambda_n^\ell)$$

for the mixing time of SOS dynamics at parameters β, λ on Λ_n with floor at height 0, ceiling at height ℓ , and height 0 boundary conditions. Our reduction of the ceiling from n to $\log n$ uses the general censoring inequality of [41]. A similar argument at $\lambda = 0$ can be found in [9, Section 6.3]. The censoring inequality says that if started from a maximal initial state in a monotone dynamics, the total-variation distance of the configuration at time t to stationarity can only increase if pre-specified updates are ignored at pre-specified times; furthermore this can only stochastically increase the law of the configuration at time t . While this usually refers to ignoring updates at certain sites for the Glauber dynamics, we apply it to censor moves that would take the SOS dynamics below or above certain heights, effectively imposing different floors/ceilings on the Glauber dynamics. This application of censoring was also used in [9, Theorem 2.2].

Definition 6.1. A censoring scheme for an SOS Glauber dynamics chain X_t prescribes a sequence of times $t_0 < t_1 < \dots$, subsets V_i of Λ_n , and heights $a_i < b_i$ such that between times $[t_{i-1}, t_i]$, the only updates that are permitted are those that move heights of vertices in V_i and only if they move them between a_i and b_i .

We therefore have in our context that if X_t^ℓ (resp. X_t^0) is the SOS Glauber dynamics on Λ_n^ℓ initialized from the maximal (resp., minimal) height of ℓ (resp., 0) everywhere, and \bar{X}_t^ℓ (resp., \bar{X}_t^0) is a censoring of it per Definition 6.1, then for all $t \geq 0$,

$$\bar{X}_t^\ell \geq X_t^\ell, \quad \text{and} \quad \bar{X}_t^0 \leq X_t^0. \quad (6.2)$$

Our goal in this subsection is to use censoring to show the following reduction.

Lemma 6.2. *There is an absolute constant C such that for all $\beta > \beta_0$, for any $\lambda \geq 0$, we have*

$$T_{\Lambda_n^n} \leq CnT_{\Lambda_n^{\log n}}.$$

Proof. For a general initial state x_0 , use the shorthand $X_t^{x_0}$ to denote the Glauber dynamics chain X_t initialized from x_0 , and use X_t^μ to denote the one initialized from the stationary distribution μ_{0, Λ_n^n} . Under the grand coupling, by a union bound and monotonicity,

$$\max_{x_0} \|\mathbb{P}(X_t^{x_0} \in \cdot) - \mu_{0, \Lambda_n^n}\|_{\text{TV}} \leq \sum_{v \in \Lambda_n} \mathbb{P}(X_t^{x_0}(v) \neq X_t^\mu(v)) \leq \sum_{v \in \Lambda_n} \mathbb{P}(X_t^n(v) \neq X_t^0(v)). \quad (6.3)$$

Writing $\mathbb{P}(X_t^n(v) \neq X_t^0(v)) = \mathbb{P}(X_t^n(v) - X_t^0(v) \geq 1)$ and using Markov's inequality, and then Eq. (6.2),

$$\max_{x_0} \|\mathbb{P}(X_t^{x_0} \in \cdot) - \mu_{0, \Lambda_n^n}\|_{\text{TV}} \leq \sum_v \mathbb{E}[X_t^n(v)] - \mathbb{E}[X_t^0(v)] \leq \sum_v \mathbb{E}[\bar{X}_t^n(v)] - \mathbb{E}[\bar{X}_t^0(v)]. \quad (6.4)$$

for any censoring scheme of the form of Definition 6.1. The censoring scheme we use is defined as follows. Define a sequence of epochs, each of length $\Delta t := C_1 T_{\Lambda_n^{\log n}} \log n$ for a large absolute constant C_1 , via

$$t_0 = 0, \quad \text{and} \quad t_i = t_{i-1} + \Delta t.$$

for $1 \leq i \leq N := \frac{2n}{\log n}$ (note that $t_N \leq CnT_{\Lambda_n^{\log n}}$). Define corresponding floor and ceiling pairs

$$a_i = n - \frac{i}{2} \log n, \quad \text{and} \quad b_i = a_i + \log n.$$

Let \bar{X}_t^n be Glauber dynamics initialized from all- n , with the following censoring:

- For each i , for times in $[t_{i-1}, t_i]$, updates are only allowed in $\Lambda_n \times [a_i, b_i]$.

Let \bar{X}_t^0 be the Markov chain initialized from all-0, with the following censoring scheme:

- For all $t \geq 0$, updates are only allowed in $\Lambda_n \times [a_N, b_N] = \Lambda_n \times [0, \log n]$.

By the definition of $T_{\Lambda_n^n} = t_{\text{MIX}}(\Lambda_n^n)$ it suffices to show that at $t = t_N$, each difference in expected values on the right-hand side of Eq. (6.4) is within n^{-4} of the stationary expected value $\mu_{0, \Lambda_n^n}[\varphi_v]$. We start with the minimal chain since that is easier; since all heights are bounded above by n , it suffices to show that the TV-distance between $\mathbb{P}(\bar{X}_t^0(v) \in \cdot)$ and $\mu_{0, \Lambda_n^n}(\varphi_v \in \cdot)$ are within n^{-5} , and the bound on the difference of expectations will follow. Towards that, notice that since $\bar{X}_t^0 \leq \log n$ everywhere, \bar{X}_t^0 is exactly an (uncensored) SOS Glauber dynamics on $\Lambda_n^{\log n}$ initialized from all-0. As such, after time $t \geq \Delta t$, by submultiplicativity of TV-distance to equilibrium, we have that

$$\|\mathbb{P}(\bar{X}_t^0(v) \in \cdot) - \mu_{0, \Lambda_n^{\log n}}\|_{\text{TV}} \leq n^{-6}.$$

At the same time, since the μ_{0, Λ_n^n} -probability that $\{\max_v \varphi_v \leq \log n\}$ is $1 - n^{-6}$ per Eq. (2.4), we get that

$$\|\mathbb{P}(\bar{X}_t^0(v) \in \cdot) - \mu_{0, \Lambda_n^n}\|_{\text{TV}} \leq \|\mathbb{P}(\bar{X}_t^0(v) \in \cdot) - \mu_{0, \Lambda_n^{\log n}}\|_{\text{TV}} + \|\mu_{0, \Lambda_n^{\log n}} - \mu_{0, \Lambda_n^n}\|_{\text{TV}} \leq O(n^{-6}). \quad (6.5)$$

We now turn to controlling the chain \bar{X}_t^n . We will show that for every $i < N$, the following event holds with high probability.

$$E_i := \bigcap_{v \in \Lambda_n} \{\bar{X}_{t_i}^n(v) \leq b_{i+1}\}.$$

If this event holds for all i , then between times $[t_i, t_{i+1}]$, the chain \bar{X}_t^n is exactly an SOS chain with floor a_{i+1} and ceiling b_{i+1} (and boundary condition 0 which is equivalent to boundary condition a_{i+1} as noted in Eq. (2.3)) run for a time Δt . Namely if Y_s^i is a standard Glauber dynamics with floor and boundary

condition a_i , ceiling b_i , initialized from $\bar{X}_{t_{i-1}}^n$, then on the event E_{i-1} , we have $\bar{X}_{t_{i-1}+s}^n \stackrel{d}{=} Y_s^{(i)}$ for $s \leq \Delta t$. Note that each $Y_s^{(i)}$ is a vertical translate of Glauber dynamics on $\Lambda_n^{\log n}$ with 0 boundary conditions, and thus has mixing time $T_{\Lambda_n^{\log n}}$. With this, we can write

$$\|\mathbb{P}(\bar{X}_t^n \in \cdot) - \mu_{0, \Lambda_n^n}\|_{\text{TV}} \leq \sum_{i \leq N} \mathbb{P}(E_{i-1}, E_i^c) + \max_{\eta: \|\eta\|_\infty \leq \log n} \|\mathbb{P}(Y_{\Delta t}^{(N)} \in \cdot \mid Y_0^{(N)} = \eta) - \mu_{0, \Lambda_n^n}\|_{\text{TV}}. \quad (6.6)$$

The total variation distance is bounded just like the bound for \bar{X}_t^0 . For the probabilities in the sum, since on E_{i-1} , we have $\bar{X}_{t_{i-1}+s}^n \stackrel{d}{=} Y_s^{(i)}$ for all $s \leq \Delta t$, it suffices to bound the probability (maximized over its possible initializations) that $Y_{\Delta t}^{(i)}$ has maximum height $b_{i+1} = a_i + \frac{1}{2} \log n$.

By definition of Δt and the sub-multiplicativity of TV-distance to stationarity, $Y_{\Delta t}^{(i)}$ will be within distance n^{-6} of its equilibrium distribution, which we said is a vertical shift of a_i above a sample from $\mu_{0, \Lambda_n^{\log n}}$. A sample from that stationary distribution, will have maximum height at most $\frac{1}{2} \log n$ except with probability n^{-6} per Eq. (2.4). Together this implies an $O(n^{-6})$ bound on the probability that $Y_{\Delta t}^{(i)}$ will have maximum height larger than b_{i+1} , and thus on $\mathbb{P}(E_{i-1}, E_i^c)$. Plugging in to Eq. (6.6), using $N = o(n)$, and combining with Eq. (6.5) gives an n^{-5} bound on $\mathbb{E}[\bar{X}_{t_N}^n(v)] - \mathbb{E}[\bar{X}_{t_N}^0(v)]$ implying that $T_{\Lambda_n^n} \leq t_N$ by Eq. (6.4). ■

6.3. Reducing the mixing time to domains of cut-width n_0 . Now that we have reduced the ceiling down to $\log n$, we wish to bound the mixing time on $\Lambda_n^{\log n}$ with zero boundary conditions, by the mixing time on smaller domains of side-length n_0 , the minimal scale at which spatial mixing kicks in (depending on $d(\lambda)$).

Recall the definition of the annulus $A_{n,m} = \Lambda_n \setminus \Lambda_{n-m}$, and let $A_{n,m}^\ell$ denote the annulus $A_{n,m}$ with ceiling at height ℓ . Let

$$T_{A_{n,m}^\ell} := t_{\text{MIX}}(A_{n,m}^\ell),$$

i.e., the mixing time on $A_{n,m}^\ell$ with 0 boundary conditions. We also need to consider mixing times over more general boundary conditions, so we use a superscript ϕ to indicate a boundary condition different from 0.

Lemma 6.3. *There exists a universal C_0 such that for all $n \geq m \geq \max\{C_0 \beta \log n / d(\lambda), \log n \log \log n\}$,*

$$T_{\Lambda_n^{\log n}} \leq C \log n \cdot \max\left\{\max_{\phi} T_{\Lambda_m^{\log n}}^{\phi}, \max_{\phi} T_{A_{n,m}^{\log n}}^{\phi}\right\}.$$

where the maxima are over boundary conditions having $\|\phi\|_\infty \leq \log n$.

Proof. Similar to Eq. (6.3), under the grand coupling,

$$\max_{x_0} \mathbb{P}(X_t^{x_0} \neq X_t^\mu) \leq \sum_{v \in \Lambda_n} \mathbb{P}(X_t^{\log n}(v) \neq X_t^0(v)),$$

where the maximal initialization is now identically $\log n$ since that is the ceiling height. We aim to localize the dynamics by introducing Markov chains that only make updates in local neighborhoods around v . For every $v \in \Lambda_{n-m}$, let $B_{v,m}$ be the ball of radius m around v , and let $X_{t,v}^{\log n}$ and $X_{t,v}^0$ be the Glauber dynamics chains that only make updates inside $B_{v,m}$ (still with the same floors and ceilings). In particular, these will be SOS chains on $B_{v,m}$ with boundary conditions $\log n$, and 0 respectively. For $v \in A_{n,m}$, use the annulus $A_{n,2m}$ as its corresponding block, and (overloading notation slightly), let $X_{t,v}^{\log n}, X_{t,v}^0$ be the SOS Glauber dynamics chains that only make updates inside $A_{n,2m}$; these will be SOS chains on $A_{n,2m}$ with boundary conditions that are 0 on the outer boundary, and $\log n$ or 0 respectively on the inner boundary.

Fix any $v \in \Lambda_n$, and notice that by monotonicity, under the grand coupling,

$$\mathbb{P}(X_t^{\log n}(v) \neq X_t^0(v)) \leq \mathbb{P}(X_{t,v}^{\log n}(v) \neq X_{t,v}^0(v)).$$

This latter probability is bounded by

$$\|\mathbb{P}(X_{t,v}^{\log n}(v) \in \cdot) - \mu_{\log n, B_v}\|_{\text{TV}} + \|\mathbb{P}(X_{t,v}^0(v) \in \cdot) - \mu_{0, B_v}\|_{\text{TV}} + \|\mu_{\log n, B_v}(\varphi_v \in \cdot) - \mu_{0, B_v}(\varphi_v \in \cdot)\|_{\text{TV}}, \quad (6.7)$$

where we used B_v as a generic stand-in for the block of v , be it the ball $B_{v,m}$ or the annulus $A_{n,2m}$. If $v \in \Lambda_{n-m}$ so that $B_v = B_{v,m}$, as long as t is at least $C_1 \log n$ (for a large absolute constant C_1) times $T_{\Lambda_m^{\log n}}^\phi$ with boundary conditions $\phi \equiv \log n$, the first of these is at most n^{-10} , and as long as it is at least $C_1 \log n$ times the mixing time on $T_{\Lambda_m^{\log n}}^\phi$ with $\phi \equiv 0$ the second of these is at most n^{-10} . Similarly when $v \in A_{n,m}$ as long as t is at least $C_1 \log n$ times $T_{A_{n,m}}^{(0,\phi)}$ with $\phi \equiv \log n$ or with $\phi \equiv 0$.

Finally, the third term in Eq. (6.7) is governed by the spatial mixing estimates we established in the previous section. Namely, if $v \in \Lambda_{n-m}$ so that $B_v = B_{v,m}$, then by Theorem 5.1 with $a_m = \log n$, it is at most $e^{-\beta m/C}$ which is at most n^{-10} , for large β . Similarly, if $v \in A_{n,m}$ so that $B_v = A_{n,2m}$, then by Theorem 5.11 with the choices $\bar{m} = n$ and $a_m = \log n$, this is at most n^{-10} as well.

Altogether, we get that so long as t is at least C_1 times $\log n$ times $\max\{\max_\phi T_{\Lambda_m^{\log n}}^\phi, \max_\phi T_{A_{n,m}}^\phi\}$ for some absolute constant C_1 , we have for all $v \in \Lambda_n$,

$$\max_{v \in \Lambda_n} \mathbb{P}(X_t^{\log n}(v) \neq X_t^0(v)) \leq O(n^{-10}),$$

which when summed over $v \in \Lambda_n$, yields the claimed bound on the mixing time. \blacksquare

6.4. Concluding the upper bound. We are now in position to conclude the upper bounds on the inverse spectral gap in Theorem 1.1 when λ approaches a critical $\lambda_c^{(k)}$ and/or 0.

Proof of Item 2 of Theorem 1.1: upper bound. By Lemma 6.2 and Lemma 6.3, with

$$m = \max\{C_0 \beta \log n / d(\lambda), \log n \log \log n\},$$

(which is less than n as required so long as $d(\lambda) \geq \frac{C_0 \beta \log n}{n}$ as assumed in Item 2 of Theorem 1.1) we obtain for some absolute constant C , the bound

$$T_{\Lambda_n^n} \leq Cn \log n \cdot \left(\max_\phi T_{\Lambda_m^{\log n}}^\phi \vee \max_\phi T_{A_{n,m}}^\phi \right). \quad (6.8)$$

The following is an upper bound on the mixing times on the right-hand side of the above that are exponential in $m \log n$, and essentially conclude the proof.

Lemma 6.4. *There is an absolute constant C , such that for all $\beta > 0$, all $\lambda \geq 0$,*

$$\max_{\phi: \|\phi\|_\infty \leq \ell} T_{\Lambda_m^\ell}^\phi \leq e^{C\beta m \ell + \lambda}, \quad \text{and} \quad \max_{\phi: \|\phi\|_\infty \leq \ell} T_{A_{n,m}^\ell}^\phi \leq n e^{C\beta m \ell + \lambda}.$$

Lemma 6.4 is a by-now standard estimate using the canonical path method ([34, Corollary 13.21]) to bound the mixing time by the cut-width of the underlying graph, a technique first developed in [28] and applied in the context of the Ising model [37]. We could not find a statement at the requisite level of generality for the SOS dynamics with general field, floors, and ceilings, (without a field, a similar bound is in [9, Proposition 2.3]) so we have included a proof for completeness in Appendix C. Plugging the bounds of Lemma 6.4 into Eq. (6.8),

$$T_{\Lambda_n^n} \leq \exp\left(C_\beta (\log n)^2 \max\left\{\frac{1}{d(\lambda)}, \log \log n\right\}\right),$$

and the analogous bound on the inverse spectral gap per Eq. (6.1). Notice that if we simply assume $\lambda \geq 0$, by Lemma 6.2 and Lemma 6.4, we also always have an upper bound of $\exp(C_\beta n \log n)$. \blacksquare

Remark 6.5. The upper bound of Lemma 6.4 can be improved to a bound of $e^{C\beta m}$ (without the dependence on ℓ in the exponent when λ is kept away from zero, say $\lambda > \varepsilon$ uniformly, since the cost to force a set of sites to all take value 1 is only exponential in the number of sites: see the argument in [12, Theorem 4.1].

6.5. Sharp upper bounds using strong spatial mixing. When $d(\lambda)$ is fixed positive, independent of n , the above reasoning gives an upper bound of $\exp(n^{o(1)})$. In this section, we give a better upper bound using the exponential-rate spatial mixing properties of the model, using strong spatial mixing which controls decay of correlations in the presence of nearby boundary conditions.

Proof of Item 1 of Theorem 1.1. Let us first define the notion of strong spatial mixing we will use.

Definition 6.6. The SOS model has SSM with constant C above scale n_0 if for all $n \geq n_0$, for all $y \in \partial\Lambda_n$, and any subset $A \subset \Lambda_n$

$$\max_{\phi, \phi': \phi_x = \phi'_x \forall x \neq y} \|\mu_{\phi, \Lambda_n}(\varphi(A) \in \cdot) - \mu_{\phi', \Lambda_n}(\varphi(A) \in \cdot)\|_{\text{TV}} \leq e^{-d(A, y)/C}.$$

It is a classical result of [40] that weak spatial mixing implies strong spatial mixing (above a certain scale) for finite-range 2D spin systems with finite state space per spin. As noted in [12, Section 2.1], in our context with infinitely-many possible heights available to each site, the strong spatial mixing still follows from weak spatial mixing so long as $\lambda > 0$ uniformly in n . This is because as long as $n_0 > \lceil 8\beta/\lambda \rceil$, there is a constant probability that a $n_0 \times n_0$ box entirely takes height 1, say, uniformly over its boundary conditions ([12, Eq. (2.1)]), whence the rest of the boosting of weak spatial mixing to strong spatial mixing can be carried out, resulting in the following consequence of Corollary 5.2.

Lemma 6.7. *For every $\varepsilon > 0$, there exists a $C_0(\varepsilon, \beta) > 0$ such that for all λ having $d(\lambda) > \varepsilon$, the SOS model has SSM with constant C_0 above scale C_0 .*

To go from Lemma 6.7 to an $O(1)$ inverse spectral gap is quite standard. We sketch the steps below with relevant references to demonstrate that so long as $\lambda > 0$ uniformly in n , there is no issue with this step: we refer to [12, Section 2.2] for further details. Let $L_0 = C_1 C_0$ where C_0 is given by Lemma 6.7 and C_1 is a constant depending on C_0 to be chosen.

- (1) Cover the box by overlapping ℓ^∞ balls (called blocks) of side-length L_0 ;
- (2) The corresponding block dynamics that assigns blocks independent Poisson clocks and performs a heat-bath update on a block when its clock rings is contractive as long as L_0 is sufficiently large as a function of C_0 (this follows from path coupling with the Hamming distance, and the amenability of \mathbb{Z}^2); in particular, the block dynamics has an $O(1)$ inverse spectral gap [12, Eq. (2.3)];
- (3) When $\lambda > 0$, each individual block has an $O(1)$ inverse spectral gap (though blowing up as $\lambda \downarrow 0$), since its base side-lengths are $L_0 = O(1)$: see [12, Proposition 2.2].

The $O(1)$ inverse spectral gap for the SOS Glauber dynamics follows from these by the classical block dynamics bound on the inverse spectral gap [38, Proposition 3.4]. ■

Remark 6.8. The dependence of n_0 for which Definition 6.6 holds depends exponentially on the m_0 at which the weak spatial mixing Theorem 5.1 kicks in, so using the above approach when the weak spatial mixing only holds above a diverging scale, e.g., above $1/d(\lambda)$, would give very sub-optimal mixing time upper bounds, namely double exponential in $1/d(\lambda)$, rather than simply exponential in $1/d(\lambda)$.

6.6. The case of the torus. We describe the modifications to get the upper bounds for the torus.

Proof of Item 2 of Theorem 1.3: upper bound. In the bounds of Section 6.2, the key role played by the boundary conditions was in enabling the a priori bound ensuring that the maximum height of the interface at equilibrium is at most $\frac{1}{2} \log n$ except with probability n^{-5} , say. Such a bound evidently does not hold at $\lambda = 0$ as there will be no force keeping the interface pinned close to height 0, but as long as we assume $\lambda \geq n^{-1+o(1)}$ say, the following lemma replaces that estimate.

Lemma 6.9. *Suppose $\lambda \geq \frac{e^{8\beta} \log n}{n}$. For $\beta > \beta_0$, the SOS model on $\Lambda_n^{\log n}$ with periodic boundary conditions, which we denote using the shorthand $\phi = p$, has*

$$\mu_{p, \Lambda_n^{\log n}}(\max_v \varphi_v \geq \frac{1}{2} \log n) \leq n^{-5}.$$

Proof. Fix a vertex v , the transitivity of the torus implying we can simply fix the origin vertex o , and consider $\mu_{p, \Lambda_n^{\log n}}(\varphi_o \geq \frac{1}{2} \log n)$. Due to the ceiling at height $\log n$, we have $\mu_{p, \Lambda_n^{\log n}}$ is stochastically dominated by $\mu_{\log n, \Lambda_n^{\log n}}$, i.e., with maximal boundary conditions, so it suffices to bound $\mu_{\log n, \Lambda_n^{\log n}}(\varphi_o \geq \frac{1}{2} \log n)$.

Let k be the k for which $\lambda \in (\lambda_c^{(k)}, \lambda_c^{(k-1)})$ and using monotonicity, lower λ down to I_{k+1} , whence $d(\lambda) \geq \frac{e^{4\beta} \log n}{n}$. Moreover, by Eq. (3.18) and the fact that $\lambda \leq \frac{n}{\log n}$ we know $k+1$ to be at most $(C/\beta) \log n$. By Corollary 5.10 with $a_m = \log n$ (applicable since $d(\lambda) \geq \frac{n}{e^{4\beta} \log n}$), there will be a height k loop surrounding $\Lambda_{n/2}$ except with probability e^{-n} . Interior to that height k -loop, the rigidity results at height k from Eq. (4.13) imply that the probability that φ_0 exceeds $k+r$ is at most $e^{-4(\beta-C)r}$. Thus, for β large, for every vertex v in the torus,

$$\mu_{p, \Lambda_n^{\log n}}(\varphi_v \geq \frac{1}{2} \log n) \leq n^{-7},$$

whence a union bound over the n^2 many such vertices implies the bound on their maximum. \blacksquare

With this upper bound in hand, the remainder of the proof of Eq. (6.6) goes through unchanged when $d(\lambda) \geq \frac{c\beta \log n}{n}$. The second reduction of Lemma 6.3 works as stated on the torus, though it makes more sense in the torus to only use blocks of the form $B_v = B_{v,m}$ (ℓ^∞ -balls of radius m about each vertex v since the graph is vertex transitive and our use of the annulus was to deal with behavior near the boundary). On the other hand, when $d(\lambda) \leq \frac{c\beta \log n}{n}$, the reduction of Lemma 6.2 together with the canonical paths bound of Lemma 6.4 gives an upper bound of $\exp(C\beta n \log n)$ on the torus. \blacksquare

Proof of Item 1 of Theorem 1.3. Regarding the $O(1)$ inverse gap estimates when $d(\lambda) > \varepsilon$ uniformly in n , the proof sketched in Section 6.5 is easily checked to go through without any modifications. \blacksquare

7. MIXING TIME LOWER BOUNDS

Our aim in this section is to establish the complementary lower bounds on the mixing times on Λ_n^n with zero and periodic boundary conditions. Our lower bounds are based on *bottlenecks* in the state space. Let us recall the basic Cheeger inequality lower bound for Markov chains. The Cheeger constant of a Markov chain with state space Ω and stationary distribution μ is

$$\Phi_\star = \min_{A \subset \Omega} \frac{Q(A, A^c)}{\mu(A)\mu(A^c)}, \quad \text{where} \quad Q(A, A^c) = \sum_{\omega \in A} \mu(\omega) \sum_{\omega' \in A^c} P(\omega, \omega'). \quad (7.1)$$

Cheeger's inequality then states that the inverse spectral gap has

$$\text{gap}^{-1} \geq \frac{1}{2} \Phi_\star^{-1}. \quad (7.2)$$

One can easily see that $Q(A, A^c) \leq \mu(\partial A)$ where ∂A indicates all configurations in A having positive transition rates to A^c . It therefore suffices to construct a set A having small conditional probability $\mu(\partial A)/\mu(A) = \mu(\partial A | A)$ to lower bound the inverse spectral gap.

7.1. Height-zero boundary conditions. We show a bottleneck between configurations that are predominantly at height $k-1$ vs. predominantly at height k , when λ is close to $\lambda_c^{(k-1)}$, but in $(\lambda_c^{(k)}, \lambda_c^{(k-1)})$. This forms a bottleneck because any contour collection whose diameter is too small (depending on $d^+(\lambda)$) would prefer to remain at height $k-1$ than rise up to k , even if the equilibrium measure prefers height k .

Proof of Item 2 of Theorem 1.1: lower bound. For a configuration φ , let $\mathcal{V}_{\geq k} \subset \Lambda_n$ denote its collection of all $\{\geq k\}$ sites. Let

$$A_r = \{\text{all connected components of } \mathcal{V}_{\geq k} \text{ have size at most } r\}.$$

Letting $o = (0, 0)$ denote the origin, our choice of bottleneck set will be

$$A = A_r, \quad \text{for} \quad r = \min\left\{\frac{n}{8}, \frac{1}{8d^+(\lambda)}\right\}.$$

We firstly claim that for an absolute constant C_0 , as long as $n \geq C_0 \beta k / d^+(\lambda)$,

$$\mu_{0, \Lambda_n^n}(A^c) \geq 1/2. \quad (7.3)$$

Indeed, using Lemma 5.6, we get a loop of height- k sites surrounding $\Lambda_{n/2}$, say, except with probability $e^{-\beta n/C}$, and by definition such a loop must have size at least $\frac{n}{2} \geq r$.

In order for a configuration to belong to ∂A , there must be a site $v \in \Lambda_n$, such that if its height is changed the configuration leaves A ; this requires that v be adjacent to some connected component of $\mathcal{V}_{\geq k}$ having size at least $r/2$. In particular,

$$\partial A \subset \bigcup_v A \cap A_{r/2, v}^c,$$

where $A_{r/2, v}^c$ is the event that there is a connected component of $\mathcal{V}_{\geq k}$ of size at least $r/2$ incident to v . By a union bound and Eq. (7.3),

$$\Phi_\star \leq 2n^2 \max_v \mu_{0, \Lambda_n^n}(A_{r/2, v}^c \mid A).$$

On the event A , there must be a loop (possibly using boundary sites which all have height 0) of height- $\{\leq k-1\}$ sites interior to B_v , the ℓ^∞ ball of radius r about v in Λ_n (otherwise there would be a $\{\geq k\}$ path from v to $\partial_o B_v \setminus \partial_o \Lambda_n$, which would necessarily have length at least r , violating A). Exposing the outermost such loop, calling its strict interior $\mathcal{C}_{k-1} \subset B_v$, this bounds the desired probability by

$$\mu_{0, \Lambda_n^n}(A_{r/2, v}^c \mid A) \leq \max_{\mathcal{C}_{k-1} \subset B_v} \max_{\phi: \|\phi\|_\infty \leq k-1} \mu_{\phi, \mathcal{C}_{k-1}}(A_{r/2, v}^c \mid A).$$

The events $A, A_{r/2, v}$ are both decreasing on this state space, and so by the FKG inequality, the probability in question is at most $\mu_{\phi, \mathcal{C}_{k-1}}(A_{r/2, v}^c)$ (without the conditioning on A). By monotonicity, this probability is maximized by the maximal boundary conditions ϕ which would be $k-1$, so we get

$$\Phi_\star \leq 2n^2 \max_v \max_{\mathcal{C}_{k-1} \subset B_v} \mu_{k-1, \mathcal{C}_{k-1}}(A_{r/2, v}^c).$$

We control this last probability via a Peierls map. The event $A_{r/2, v}^c$ requires there to be an outermost contour of diameter at least $r/2$ (but at most $2r$ since $\mathcal{C}_{k-1} \subset B_v$) confining v in its interior. Considering the renormalized contour representation and performing the Peierls operation that deletes the outermost contour γ_v confining v in its interior, the weight change under application of this map is $W_{k-1}^{rn}(\gamma_v)$ which by Corollary 4.9 is at most $e^{-(\beta-5)|\gamma_v|}$ because $\text{diam}(\gamma_v) \leq 2r \leq 1/(4d^+(\lambda))$, which in turn is at most $(f_k^{\text{tr}} - f_{k-1}^{\text{tr}})^{-1}$ per Corollary 4.7. Summing over the possible choices of ℓ , we arrive at

$$\Phi_\star \leq 2n^2 \sum_{\ell \geq r} 4^\ell e^{-(\beta-5)\ell} \leq n^2 e^{-(\beta-C)r}.$$

This then implies the claimed bound on the inverse gap per Eq. (7.2). ■

7.2. The case of the torus. We describe how to get the analogous lower bound on the torus.

Proof of Item 2 of Theorem 1.3: lower bound. Let k be such that $d(\lambda)$ is attained by $\lambda_c^{(k-1)}$, whether by $d(\lambda) = d^+(\lambda)$, in which case $\lambda = \lambda_c^{(k-1)} - d(\lambda)$, or if $d(\lambda) = d^-(\lambda)$ in which case $\lambda = \lambda_c^{(k-1)} + d(\lambda)$. Consider the following two events:

$$\begin{aligned} A_r &= \{\text{all connected components of } \mathcal{V}_{\geq k} \text{ have size at most } r\}, \\ A'_r &= \{\text{all connected components of } \mathcal{V}_{\leq k-1} \text{ have size at most } r\}. \end{aligned}$$

Similar to the proof with zero boundary conditions, fix $r = \min\{\frac{n}{8}, \frac{1}{d(\lambda)}\}$. Notice that A_r and A'_r are disjoint events because on A_r , the boundary in Λ_{2r} of the $\{\geq k\}$ components intersecting $\partial_i \Lambda_{2r}$, say, forms a $\{\leq k-1\}$ path of diameter at least r . As such, for every λ , one of these two events must have probability less than $1/2$. Thus, the bound of Eq. (7.3) is ensured by the choice of $A = A_r$ vs. A'_r .

From that point on, the argument is essentially identical to the argument in the previous subsection if $A = A_r$ with the balls B_v now all being identical translates of one another. We had used that $\text{diam}(\gamma_v) \leq 2r \leq 1/(4d(\lambda))$ in order to reason that $W_{k-1}^{rn}(\gamma_v)$ had an exponentially small weight. This holds still if $d(\lambda) = d^+(\lambda)$, while if $d(\lambda) = d^-(\lambda)$ then $\lambda \in (\lambda_c^{(k-1)}, \lambda_c^{(k)})$ whence the bound on the diameter isn't even needed for the exponentially small weight per Proposition 4.8. The reasoning if $A = A'_r$ is symmetrical. ■

APPENDIX A. MONOTONICITY WITH GENERIC FIELDS AND FLOORS/CEILINGS

Recall the general form of the SOS measure with arbitrary floors $\mathbf{a} = (a_v)_v$ and ceilings $\mathbf{b} = (b_v)_v$, boundary conditions ϕ , and external field λ from Eq. (2.1).

Proof of Lemma 2.1. It in fact suffices to check this for V being a single vertex, say the origin, because given that, one can run a column Glauber dynamics (fully resampling the height at a vertex), whose $t \rightarrow \infty$ limit serves as the coupling attaining the stochastic domination expressed above. At a single vertex, we check the stochastic domination between the two distributions. The distribution can be expressed as

$$p(\varphi) := e^{-\lambda\varphi_v} \left(\prod_{w \sim v} e^{-\beta|\varphi_v - \phi_w|} \right) \mathbf{1}_{\varphi_v \geq a_v} \mathbf{1}_{\varphi_v \leq b_v},$$

and $p'(\varphi)$ is defined analogously with the primed parameters. The ratio $p'(\varphi)/p(\varphi)$ can be broken up into the individual ratios in the above product. We claim that each of these are individually increasing in φ_v . Since the product of increasing functions will also be increasing, that gives us that the ratio of the mass functions is increasing, from which the stochastic domination follows.

The ratio $e^{-(\lambda' - \lambda)\varphi_v}$ is increasing in φ_v because $\lambda' - \lambda \leq 0$. The ratios of the indicator functions are clearly increasing in φ_v since

$$\frac{\mathbf{1}_{\varphi_v \geq a'_v}}{\mathbf{1}_{\varphi_v \geq a_v}} = \begin{cases} 0 & \varphi_v \in [a_v, a'_v] \\ 1 & \varphi_v \in [a'_v, \infty) \end{cases}, \quad \text{and} \quad \frac{\mathbf{1}_{\varphi_v \leq b'_v}}{\mathbf{1}_{\varphi_v \leq b_v}} = \begin{cases} 1 & \varphi_v \in (-\infty, b_v] \\ \infty & \varphi_v \in [b_v, b'_v) \end{cases},$$

(both measures are supported on $[a_v, b'_v]$ so we don't care about the cases where it is 0 divided by 0). Finally,

$$\frac{e^{-|\varphi_v - \phi'_w|}}{e^{-|\varphi_v - \phi_w|}} = e^{|\varphi_v - \phi_w| - |\varphi_v - \phi'_w|} = \exp \left(\begin{cases} \phi_w - \phi'_w & \varphi_v \leq \phi_w \\ (\varphi_v - \phi_w) - (\phi'_w - \varphi_v) & \varphi_v \in (\phi_w, \phi'_w) \\ \phi'_w - \phi_w & \varphi_v \geq \phi'_w \end{cases} \right).$$

The middle term can be re-expressed as $2(\varphi_v - \frac{1}{2}(\phi_w + \phi'_w))$. It is easy to see with this writing that for $\varphi_v \in (\phi_w, \phi'_w)$, this is in absolute value less than $\phi'_w - \phi_w$, so it is also increasing in φ_v . ■

APPENDIX B. EXPONENTIAL TAILS ON NON-ELEMENTARY RENORMALIZED WEIGHTS

In this section we follow the strategy of [15] adapted to contours to provide the proof of Lemma 4.1.

B.1. Clusters of non-elementary contours. We begin by canonically splitting any contour collection in $\mathcal{G}_{\eta, h, V}$ up into clusters of contours γ that are non-elementary at $h(\gamma)$, and in between regions where the contributions are given by corresponding renormalized elementary partition functions. This follows the exposition leading up to Proposition 2.6 in [11] (also [15, Eq. (2.33)]).

Definition B.1. [Cluster of non-elementary contours] A set of contours \mathcal{C} is called a *non-elementary cluster* if one contour in \mathcal{C} nests all the others, along each nesting path ($\gamma_1 \subset \gamma_2 \subset \dots$) only the heights exterior to the outermost contour and interior to the innermost contour are h , and every $\gamma \in \mathcal{C}$ is $h(\gamma)$ -non-elementary.

Any collection of contours (γ_i) each of which are $h(\gamma_i)$ -non-elementary can be uniquely decomposed into a set of clusters of non-elementary contours.

In words, given a set of γ each of which is $h(\gamma)$ -non-elementary, a cluster of non-elementary contours is an excursion away from height h . For a cluster of non-elementary contours \mathcal{C} , and each $\gamma \in \mathcal{C}$, let

$$\text{Ann}(\gamma; \mathcal{C}) = \text{Int}(\gamma) \setminus \bigcup_{\gamma' \in \mathcal{C}: \gamma' \subsetneq \gamma} \text{Int}(\gamma'),$$

if $\exists \gamma' \in \mathcal{C} : \gamma' \subsetneq \gamma$, and let $\text{Ann}(\gamma; \mathcal{C}) = \emptyset$ otherwise. We can then define

$$\text{Ann}(\mathcal{C}) = \bigcup_{\gamma} \text{Ann}(\gamma; \mathcal{C}).$$

Definition B.2. If \mathcal{C} is a cluster of non-elementary contours, its h -renormalized weight is given by

$$W_h^{\text{rn}}(\mathcal{C}) = e^{-\beta \sum_{\gamma \in \mathcal{C}} |\gamma|} \frac{\prod_{\gamma \in \mathcal{C}} Z_{S(\gamma), h(\gamma), \text{Ann}(\gamma; \mathcal{C})}^{\text{el}}}{Z_{h, \text{Ann}(\mathcal{C})}^{\text{rn.el}}}, \quad (\text{B.1})$$

where we recall $S(\gamma) = +$ if γ is an up contour and $S(\gamma) = -$ if γ is a down contour.

Remark B.3. Notice that on the numerator of Eq. (B.1), it does not matter that we are writing Z^{el} rather than $Z^{\text{rn.el}}$ for the not-necessarily simply connected domain $\text{Ann}(\gamma; \mathcal{C})$ since all its holes are necessarily $h(\gamma)$ -non elementary and therefore cannot be confined by $h(\gamma)$ -elementary contours.

The following shows that these renormalized weights together with the renormalized weights of the elementary contours give a different way of writing the renormalized partition function (c.f., Lemma 3.2). Let $\mathcal{C}_{\eta, h, V}$ be the set of all admissible collections of clusters of h -non-elementary contours. When we write the pair $(C, \Gamma) \in (\mathcal{C}_{\eta, h, V}, \mathcal{G}_{\eta, h, V}^{\text{rn.el}})$, we further impose that in each \mathcal{C} , the contours of Γ are compatible with the outermost and innermost contours of \mathcal{C} (i.e., the boundary of $\text{Ann}(\mathcal{C})$).

Lemma B.4. For every simply-connected V , every h , and every boundary signing η ,

$$Z_{\eta, h, V}^{\text{rn}} = e^{-\lambda h |V|} \sum_{(C, \Gamma) \in (\mathcal{C}_{\eta, h, V}, \mathcal{G}_{\eta, h, V}^{\text{el}})} \prod_{\mathcal{C} \in C} W_h^{\text{rn}}(\mathcal{C}) \prod_{\gamma \in \Gamma} W_h^{\text{rn}}(\gamma).$$

Proof. We can express the full SOS partition function as

$$\begin{aligned} Z_{\eta, h, V} &= \sum_{C \in \mathcal{C}_{\eta, h, V}} Z_{h, \text{Ext}(C)}^{\text{rn.el}} \prod_{\mathcal{C} \in C} e^{-\beta \sum_{\gamma \in \mathcal{C}} |\gamma|} \prod_{\gamma \in \mathcal{C}} Z_{S(\gamma), h(\gamma), \text{Ann}(\gamma; \mathcal{C})}^{\text{el}} \\ &= \sum_{C \in \mathcal{C}_{\eta, h, V}} Z_{h, \text{Ext}(C)}^{\text{rn.el}} \prod_{\mathcal{C} \in C} W_h^{\text{rn}}(\mathcal{C}) Z_{h, \text{Ann}(C)}^{\text{rn.el}}. \end{aligned}$$

Expanding out each $Z^{\text{rn.el}}$ term per Eq. (3.2) yields the claim. \blacksquare

Using our understanding of the elementary partition functions in the windows I_h from Eq. (3.18), we get the following upper bound on the renormalized weights of clusters of non-elementary contours.

Lemma B.5. If $\lambda \in I_h$, for every cluster \mathcal{C} of non-elementary contours,

$$W_h^{\text{rn}}(\mathcal{C}) \leq \exp \left(-(\beta - 1) \sum_{\gamma \in \mathcal{C}} |\gamma| - \sum_{\gamma \in \mathcal{C}} e^{-4\beta(h \wedge (h(\gamma)+1)) - 3\beta} |\text{Ann}(\gamma; \mathcal{C})| \right).$$

Proof. Notice first that

$$Z_{h, \text{Ann}(\mathcal{C})}^{\text{rn.el}} \geq \prod_{\gamma \in \mathcal{C}} Z_{S(\gamma), h, \text{Ann}(\gamma; \mathcal{C})}^{\text{rn.el}}.$$

since the latter is only more restrictive in terms of which contours are permitted in $\text{Ann}(\mathcal{C})$. Therefore,

$$W_h^{\text{rn}}(\mathcal{C}) \leq e^{-\beta \sum_{\gamma \in \mathcal{C}} |\gamma|} \prod_{\gamma \in \mathcal{C}} \frac{Z_{S(\gamma), h(\gamma), \text{Ann}(\gamma; \mathcal{C})}^{\text{el}}}{Z_{S(\gamma), h, \text{Ann}(\gamma; \mathcal{C})}^{\text{rn.el}}}.$$

Using Lemma 3.12 and the lower bounds on the elementary free energies from Lemma 3.13 when $\lambda \in I_h$ on each of the terms in the product, we conclude (noting that $\sum_{\gamma \in \mathcal{C}} |\partial_e \text{Ann}(\gamma; \mathcal{C})| \leq 2 \sum_{\gamma \in \mathcal{C}} |\gamma|$). ■

B.2. Weights of non-elementary contours. For a cluster of non-elementary contours \mathcal{C} , let γ^{out} be its outermost contour (there being only one by the definition of such clusters). We begin with the following bound on the total contribution from all clusters of non-elementary contours nesting a point $x \in V$, and having $|\gamma^{\text{out}}| \geq r$.

Lemma B.6. *If $\lambda \in I_h$, then for every r ,*

$$\sum_{\substack{\mathcal{C}: x \in \text{Int}(\gamma^{\text{out}}) \\ |\gamma^{\text{out}}| \geq r}} W_h^{\text{rn}}(\mathcal{C}) \leq \exp(-(\beta - 4)r).$$

Proof. We follow the general approach of [15, Lemmas 2.13–2.14] to perform the sum by associating to each cluster of non-elementary contours a corresponding tree-like object. We associate to \mathcal{C} a *witness* as follows:

- Take the set of (dual) edges associated to the collection of contours rooted at their respective heights, and color all such edges RED.
- For each contour γ , and each contour γ' nested in γ , add a straight line in the e_1 direction of BLUE faces at $h(\gamma)$ connecting γ to γ' . Then for each γ' nested in γ , erase all BLUE faces in their interiors, and continue the procedure for those γ' .

This gives a treelike structure on contours, where a contour γ is the parent of γ' if they are connected by BLUE faces, and either γ nests γ' or γ is closer (in this ordering) to their common nesting contour. For each γ , let $F_B(\gamma)$ be the set of BLUE faces between γ and its children (these being distinct). We can upper bound the renormalized weight $W_h^{\text{rn}}(\mathcal{C})$ associated to such a \mathcal{C} by the following:

$$W_h^{\text{rn}}(\mathcal{C}) \leq \prod_{\gamma} e^{-(\beta-1)|\gamma| - |F_B(\gamma)| e^{-4\beta(h \wedge h(\gamma)) - 3\beta}}.$$

In order to see this, recall Lemma B.5 and notice that for each γ that is not a leaf in the tree construction, either the BLUE faces constituting $F_B(\gamma)$ are part of $\text{Ann}(\gamma; \mathcal{C})$, or they are part of $\text{Ann}(\gamma'; \mathcal{C})$ where γ' is the innermost nesting contour of γ . that is to say that they are either at height $h(\gamma)$ or at height $\max\{0, h(\gamma) \pm 1\}$.

At the same time, any child of γ must be either $h(\gamma)$ -non elementary or $h(\gamma')$ -non-elementary where γ' is innermost nesting contour. so in particular, it must have size at least $L_0 := e^{3\beta \max\{1, h(\gamma^{\text{out}})\}}$.

We now perform the count one generation of the tree corresponding to \mathcal{C} , by enumerating over γ^{out} of size at least r , then for each point along γ^{out} , deciding whether to start a blue path or not, enumerating over the length of the blue path k , then picking up a weight corresponding to all the possible choices of subtree to place there. Assuming (inductively) that the total weight of all subtrees whose outermost contour has length at least L_0 has a bound of $\exp(-(\beta - 4)L_0)$, we get

$$\sum_{\substack{\mathcal{C}: x \in \text{Int}(\gamma^{\text{out}}) \\ |\gamma^{\text{out}}| \geq r}} W_h^{\text{rn}}(\mathcal{C}) \leq \sum_{\ell \geq r} 4^\ell e^{-(\beta-1)\ell} (1 + e^{-(\beta-4)L_0} \sum_{k=1}^{\infty} e^{-k e^{-4\beta(h \wedge h(\gamma^{\text{out})) - 3\beta}})^\ell).$$

The series in k sums up to at most $e^{4\beta h(\gamma^{\text{out}}) + 3\beta}$ so this is at most

$$\sum_{\ell \geq r} e^{-(\beta-3)\ell} \exp(\ell e^{-(\beta-4)L_0} e^{3\beta \max\{1, h(\gamma^{\text{out}})\}} e^{4\beta h(\gamma^{\text{out}}) + 3\beta}).$$

For β large, the second exponential is easily seen to be at most $\exp(\varepsilon_\beta \ell)$, whence the sum is easily seen to be at most $\exp(-(\beta - 4)r)$ as claimed. ■

Proof of Lemma 4.1. If γ is elementary, the bound follows from Theorem 3.6. Suppose now that γ is an up non-elementary contour. Then,

$$W_h^{\text{rn}}(\gamma) = e^{-\beta|\gamma|} \frac{Z_{\eta, h+1, \text{Int}(\gamma)}}{Z_{\eta, h, \text{Int}(\gamma)}} = \mu_{+, h, \text{Int}(\gamma)} \left(\bigcap_{x \in \partial_i \text{Int}(\gamma)} \{\varphi(x) \geq h+1\} \right).$$

In order for $\varphi_x \geq h+1$ for all $x \in \partial_i \text{Int}(\gamma)$, then there must be a positive contour γ' in $\Gamma(\varphi)$ that lines up exactly with γ ; If γ is not h -elementary, then by Lemma B.4, this probability is bounded by the Peierls map that deletes its entire cluster \mathcal{C} (whose weight is $W_h^{\text{rn}}(\mathcal{C})$), whence Lemma B.6 performs the summation over the choice of the cluster and concludes the proof. The case where γ is a down non-elementary contour is analogous. ■

APPENDIX C. CANONICAL PATHS BOUNDS FOR SOS MODELS

We show that the inverse gap on domains of cut-width m and ceiling ℓ is at most exponential in $m\ell$.

Proof of Lemma 6.4. Recall (e.g., [34, Corollary 13.21]) that for a reversible Markov chain with a finite state space Ω , transition matrix P , and stationary distribution μ , its spectral gap is lower bounded by the congestion

$$\min_{\Gamma = (\gamma_{ab})_{a, b \in \Omega}} \max_{x, y \in \Omega: P(x, y) > 0} \frac{1}{\mu(x)P(x, y)} \sum_{w, z: (x, y) \in \gamma_{w, z}} \mu(w)\mu(z)|\gamma_{w, z}|, \quad (\text{C.1})$$

where each γ_{ab} is a sequence of positive rate transitions starting at a and ending at b . We first do the argument for Λ_m , then describe what changes for the annulus $A_{n, m}$.

Enumerate the vertices of Λ_m in lexicographic order, denoted v_1, v_2, \dots . For two configurations a, b , the path γ_{ab} is described by the updates where each of these vertices is sequentially processed, and by processed we mean that if $a_v \neq b_v$, then there is a minimal sequence of transitions to take a_v to b_v . Evidently, $|\gamma_{a, b}| \leq m^2\ell$ for all a, b . Further, for fixed (x, y) , there is an injection between pairs of configurations (w, z) such that $(x, y) \in \gamma_{w, z}$ and Ω , defined as follows. If (x, y) entails a transition at site v_i , let \tilde{w} agree with w on $(v_j)_{j>i}$ and let \tilde{w} agree with z on $(v_j)_{j<i}$; finally let it take the value of x on v_i . Similarly, let \tilde{z} agree with z on $(v_j)_{j>i}$, agree with w on $(v_j)_{j<i}$, and take the value of y at v_i . Given i , we can evidently recover (w, z) from (\tilde{w}, \tilde{z}) . Moreover, the configuration x is exactly \tilde{w} , so given x and i which we can read off from (x, y) , this gives an injection from $(w, z) : (x, y) \in \gamma_{wz}$ to \tilde{z} , and the congestion can be rewritten as

$$\frac{m^2\ell}{P(x, y)} \sum_{\tilde{z} \in \Omega} \frac{\mu(w)\mu(z)}{\mu(\tilde{w})\mu(\tilde{z})} \mu(\tilde{z}).$$

For any valid transition, $P(x, y)$ is clearly lower bounded by $e^{-4\beta\ell - \lambda}$. In the ratio of weights between w, z and \tilde{w}, \tilde{z} the external field quantities cancel out exactly, as do all gradients except those along the cut between $(v_j)_{j<i}$ and $(v_j)_{j>i}$. In the lexicographic ordering on Λ_m , this cut is at most m for every i , so the largest that ratio can be is $e^{\beta m \ell}$. The sum over \tilde{z} then gives a 1 since μ is a probability measure, altogether yielding the desired.

We can absorb the m^2 and ℓ , and the polynomial in $m^2\ell$ factor that comes from translating spectral gap into mixing time, by adjusting the constant in the exponent, getting the desired.

In order to get the claimed bound for the annulus, follow the same reasoning, but instead of using the lexicographic ordering, we use one that attains a cut-width of at most $2m$. This ordering will be obtained by starting from the $m \times m$ block in one corner, performing the lexicographic order there, then processing around the annulus one row/column (whichever is thinner, depending on the orientation) one at a time. The extra factor of n in the bound here comes from the translation of spectral gap to mixing time. ■

Acknowledgements. R.G. acknowledges the support of NSF DMS-2246780 and E.L. acknowledges the support of NSF DMS-2054833.

REFERENCES

- [1] K. S. Alexander, F. Dunlop, and S. Miracle-Solé. Layering and wetting transitions for an SOS interface. *J. Statist. Phys.*, 142(3):524–576, 2011.
- [2] R. Brandenberger and C. E. Wayne. Decay of correlations in surface models. *J. Statist. Phys.*, 27(3):425–440, 1982.
- [3] J. Brémont, A. El Mellouki, and J. Fröhlich. Random surfaces in statistical mechanics: roughening, rounding, wetting. . . . *J. Statist. Phys.*, 42(5-6):743–798, 1986.
- [4] W. K. Burton, N. Cabrera, and F. C. Frank. The growth of crystals and the equilibrium structure of their surfaces. *Philos. Trans. Roy. Soc. London Ser. A*, 243:299–358, 1951.
- [5] P. Caputo and S. Ganguly. Uniqueness, mixing, and optimal tails for Brownian line ensembles with geometric area tilt. 2023. Preprint, [arXiv:2305.18280](https://arxiv.org/abs/2305.18280).
- [6] P. Caputo, D. Ioffe, and V. Wachtel. Confinement of Brownian polymers under geometric area tilts. *Electron. J. Probab.*, 24:Paper No. 37, 21, 2019.
- [7] P. Caputo, D. Ioffe, and V. Wachtel. Tightness and line ensembles for Brownian polymers under geometric area tilts. In *Statistical mechanics of classical and disordered systems*, volume 293 of *Springer Proc. Math. Stat.*, pages 241–266. Springer, Cham, 2019.
- [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] F. Cesi and F. Martinelli. On the layering transition of an SOS surface interacting with a wall. I. Equilibrium results. *J. Statist. Phys.*, 82(3):823–913, 1996.
- [12] F. Cesi and F. Martinelli. On the layering transition of an SOS surface interacting with a wall. II. The Glauber dynamics. *Comm. Math. Phys.*, 177(1):173–201, 1996.
- [13] J. T. Chalker. The pinning of an interface by a planar defect. *J. Phys. A*, 15(9):L481–L485, 1982.
- [14] A. Dembo, E. Lubetzky, and O. Zeitouni. On the limiting law of line ensembles of Brownian polymers with geometric area tilts. *Annales de l’I.H.P. Probabilités et statistiques*. To appear.
- [15] E. I. Dinaburg and A. E. Mazel. Layering transition in SOS model with external magnetic field. *J. Statist. Phys.*, 74(3):533–563, 1994.
- [16] R. Dobrushin, R. Kotecký, and S. Shlosman. *Wulff construction*, volume 104 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1992.
- [17] N. Feldheim and S. Yang. Typical height of the $(2+1)$ -D Solid-on-Solid surface with pinning above a wall in the delocalized phase. *Stochastic Processes and their Applications*, 165:168–182, 2023.
- [18] P. L. Ferrari and H. Spohn. Constrained Brownian motion: fluctuations away from circular and parabolic barriers. *Ann. Probab.*, 33(4):1302–1325, 2005.
- [19] S. Friedli and Y. Velenik. *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*. Cambridge University Press, 2017.
- [20] J. Fröhlich and T. Spencer. Kosterlitz-Thouless transition in the two-dimensional plane rotator and Coulomb gas. *Phys. Rev. Lett.*, 46(15):1006–1009, 1981.
- [21] J. Fröhlich and T. Spencer. The Kosterlitz-Thouless transition in two-dimensional abelian spin systems and the Coulomb gas. *Comm. Math. Phys.*, 81(4):527–602, 1981.
- [22] S. Ganguly and R. Gheissari. Local and global geometry of the 2D Ising interface in critical prewetting. *Ann. Probab.*, 49(4):2076–2140, 2021.
- [23] R. Gheissari and E. Lubetzky. Approximate domain markov property for rigid Ising interfaces. *J. Statist. Phys.*, 190(5):99, 2023.
- [24] R. Gheissari and E. Lubetzky. Entropic repulsion of 3D Ising interfaces conditioned to stay above a floor. *Electron. J. Probab.*, 28:1–44, 2023.
- [25] D. A. Huse and D. S. Fisher. Dynamics of droplet fluctuations in pure and random Ising systems. *Phys. Rev. B*, 35:6841–6846, May 1987.
- [26] D. Ioffe, S. Ott, S. Shlosman, and Y. Velenik. Critical prewetting in the 2D Ising model. *Ann. Probab.*, 50(3):1127–1172, 2022.
- [27] D. Ioffe and Y. Velenik. Low-temperature interfaces: prewetting, layering, faceting and Ferrari-Spohn diffusions. *Markov Process. Related Fields*, 24(3):487–537, 2018.
- [28] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989.
- [29] R. Kotecký and D. Preiss. Cluster expansion for abstract polymer models. *Comm. Math. Phys.*, 103(3):491–498, 1986.
- [30] H. Lacoïn. Wetting and layering for solid-on-solid I: Identification of the wetting point and critical behavior. *Comm. Math. Phys.*, 362(3):1007–1048, 2018.

- [31] H. Lacoïn. Wetting and layering for solid-on-solid II: Layering transitions, Gibbs states, and regularity of the free energy. *J. Éc. polytech. Math.*, 7:1–62, 2020.
- [32] P. Lammers. A dichotomy theory for height functions. Preprint, [arXiv:2211.14365](https://arxiv.org/abs/2211.14365).
- [33] J. L. Lebowitz and A. E. Mazel. A remark on the low-temperature behavior of the SOS interface in half-space. *J. Statist. Phys.*, 84(3-4):379–397, 1996.
- [34] D. Levin and Y. Peres. *Markov Chains and Mixing Times (2nd ed.)*. American Mathematical Society, Providence, RI, 2017.
- [35] T. M. Liggett, R. H. Schonmann, and A. M. Stacey. Domination by product measures. *Ann. Probab.*, 25(1):71–95, 1997.
- [36] E. Lubetzky, F. Martinelli, A. Sly, and F. L. Toninelli. Quasi-polynomial mixing of the 2D stochastic Ising model with “plus” boundary up to criticality. *J. Eur. Math. Soc. (JEMS)*, 15(2):339–386, 2013.
- [37] F. Martinelli. On the two-dimensional dynamical Ising model in the phase coexistence region. *J. Statist. Phys.*, 76(5):1179–1246, 1994.
- [38] F. Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 93–191. Springer, Berlin, 1999.
- [39] F. Martinelli and E. Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region. I. The attractive case. *Comm. Math. Phys.*, 161:447–486, 1994.
- [40] F. Martinelli, E. Olivieri, and R. H. Schonmann. For 2-D lattice spin systems weak mixing implies strong mixing. *Comm. Math. Phys.*, 165(1):33–47, 1994.
- [41] Y. Peres and P. Winkler. Can extra updates delay mixing? *Comm. Math. Phys.*, 323(3):1007–1016, 2013.
- [42] 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.
- [43] R. H. Schonmann. Slow droplet-driven relaxation of stochastic Ising models in the vicinity of the phase coexistence region. *Comm. Math. Phys.*, 161(1):1–49, 1994.
- [44] R. H. Schonmann. Theorems and conjectures on the droplet-driven relaxation of stochastic Ising models. In *Probability and phase transition (Cambridge, 1993)*, volume 420 of *NATO Adv. Sci. Inst. Ser. C: Math. Phys. Sci.*, pages 265–301. Kluwer Acad. Publ., Dordrecht, 1994.
- [45] R. H. Schonmann and S. B. Shlosman. Wulff droplets and the metastable relaxation of kinetic Ising models. *Comm. Math. Phys.*, 194(2):389–462, 1998.
- [46] H. N. V. Temperley. Statistical mechanics and the partition of numbers. II. The form of crystal surfaces. *Proc. Cambridge Philos. Soc.*, 48:683–697, 1952.
- [47] Y. Velenik. Entropic repulsion of an interface in an external field. *Probability Theory and Related Fields*, 129(1):83–112, May 2004.

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