

EXPONENTIALLY SLOW MIXING ARISING FROM ENTROPIC REPULSION IN p -SOS MODEL

SEOKUN CHOI

ABSTRACT. We investigate the Glauber dynamics of the generalized $(2+1)$ -dimensional p -SOS model under a hard floor constraint. This setting induces entropic repulsion: the integer-valued interface height is forced to remain above the wall and consequently rises to a typical height $H(p, L)$ that depends on both the parameter p and the system size L . The phenomenon of entropic repulsion has been extensively studied in a variety of random interface models. In the classical SOS model ($p = 1$), [8, 9] derived an exponential lower bound for the mixing time, demonstrating that the Glauber dynamics mixes only after an exponentially long time in the low-temperature regime (large β , the inverse temperature). However, beyond this case, no rigorous lower bounds were previously known: even for the widely studied Discrete Gaussian model ($p = 2$), the metastable slowdown predicted by the entropic repulsion picture had remained an open problem. On the equilibrium side, [17] obtained sharp large-deviation principles and precise estimates of typical and maximal heights for all $1 \leq p < \infty$, but the dynamical consequences of these results had not been established.

Our main contribution is to close this gap by proving, for the first time, that exponentially slow mixing arising from entropic repulsion persists throughout the regime $1 < p < \infty$. Specifically, we establish an exponential lower bound, showing that the mixing time satisfies $\tau_{\text{mix}} \geq e^{cL}$ for some $c > 0$ depending on p and β . In addition, we provide a refined metastability analysis, proving that the hitting time of an intermediate level $aH(p, L)$ is at least $\exp(cL^{a^{d(p)}})$, where $d(p)$ is a positive function depending on p . The proof relies on an extension of the classical Peierls-type contour estimates, originally developed for $p = 1$, to the nonlinear p -SOS setting. Taken together, these results demonstrate that entropic repulsion induces uniformly slow mixing across the entire p -SOS family, thereby extending a phenomenon that had previously been established only for $p = 1$.

CONTENTS

1. Introduction	1
2. Model and Technical Framework	9
3. Equilibrium Properties under Entropic Repulsion	11
4. Proofs of Theorem 1 and Theorem 2	17
5. Appendix	25
References	27

1. INTRODUCTION

1.1. Background. The Solid-On-Solid (SOS) model, originally introduced by Temperley [20] in 1952 and also known as the Onsager-Temperley sheet, was proposed to describe the microscopic behavior of the crystal interface. In modern statistical mechanics, this model is viewed as an effective

description of the interface separating the positive and negative phases in the $(d+1)$ -dimensional Ising model, particularly in the low-temperature regime. In this setting, the SOS interface approximates the Ising interface by suppressing overhangs and bulk fluctuations, which are energetically costly at low temperatures.

We define the SOS model under general boundary conditions as follows. For a finite domain $\Lambda \subset \mathbb{Z}^2$, consider integer-valued height functions

$$\eta : \Lambda \rightarrow \mathbb{Z}, \quad \eta \equiv \phi \text{ on } \Lambda^c,$$

where $\phi : \Lambda^c \rightarrow \mathbb{Z}$ is a fixed boundary condition. The Hamiltonian is given by

$$H_\Lambda^\phi(\eta) = \sum_{\substack{x \in \Lambda, y \in \mathbb{Z}^2 \\ |x-y|=1}} |\eta(x) - \xi(y)|, \quad \xi(y) = \begin{cases} \eta(y), & y \in \Lambda, \\ \phi(y), & y \in \Lambda^c, \end{cases}$$

which includes both the interactions inside Λ and those with the prescribed boundary condition. The corresponding Gibbs measure is then

$$\hat{\pi}_\Lambda^\phi(\eta) = \frac{1}{Z_\Lambda^\phi} \exp(-\beta H_\Lambda^\phi(\eta)),$$

where $\beta > 0$ is the inverse temperature and Z_Λ^ϕ is the partition function ensuring normalization. In the special case of zero boundary condition $\phi \equiv 0$, we denote the Gibbs measure by $\hat{\pi}_\Lambda^0(\eta)$, and for convenience we shall simply write $\hat{\pi}_\Lambda(\eta)$.

The two-dimensional case $d = 2$ plays a distinguished role, being the unique dimension where a roughening transition takes place. For $d = 1$, the interface stays non-localized for all inverse temperatures $\beta > 0$, implying that the average height at a given site diverges as the system size increases [11, 19]. In contrast, for $d \geq 3$, Peierls-type arguments show that the interface is localized [7], with height fluctuations uniformly bounded in expectation. In two dimensions, however, a transition occurs between delocalized and localized behavior [5, 16]: at high temperatures the interface becomes rough, while at low temperatures it is localized. This critical phenomenon was first rigorously analyzed in the foundational works of Fröhlich and Spencer [14, 15].

To capture a broader range of interface behaviors, the classical SOS model [8, 17] can be extended to the so-called p -SOS model, defined for $p \geq 1$. Configurations follow the SOS model, that is, height functions taking integer values on $\Lambda \subset \mathbb{Z}^2$ together with a specified boundary condition $\phi : \Lambda^c \rightarrow \mathbb{Z}$.

The Hamiltonian is given by

$$H_\Lambda^{p,\phi}(\eta) = \sum_{\substack{x \in \Lambda, y \in \mathbb{Z}^2 \\ |x-y|=1}} |\eta(x) - \xi(y)|^p, \quad \xi(y) = \begin{cases} \eta(y), & y \in \Lambda, \\ \phi(y), & y \in \Lambda^c, \end{cases}$$

and the corresponding Gibbs measure is

$$\hat{\pi}_\Lambda^{p,\phi}(\eta) = \frac{1}{Z_\Lambda^{p,\phi}} \exp(-\beta H_\Lambda^{p,\phi}(\eta)), \tag{1}$$

where $\beta > 0$ is the inverse temperature, $p \in [1, \infty)$ is fixed, and $Z_\Lambda^{p,\phi}$ is the partition function ensuring normalization. In the special case of zero boundary condition $\phi \equiv 0$, we write $\hat{\pi}_\Lambda^{p,0}(\eta)$ and, for convenience, simply denote it by $\hat{\pi}_\Lambda^p(\eta)$.

This framework includes two notable special cases: when $p = 1$, the model reduces to the classical SOS model, while $p = 2$ corresponds to the Discrete Gaussian(DG) model. The parameter p thus interpolates between models with different types of gradient penalization, allowing for a unified

approach to interface fluctuation analysis. In this work, we aim to extend certain results previously established for the case $p = 1$ to the broader range $1 \leq p < \infty$. Recent developments show a growing interest in such generalized models [1, 2, 10, 17], as they offer a flexible framework for understanding a wide class of random interfaces.

One of the central phenomena in random interface (surface) models with hard constraints, for instance when a floor prevents penetration, is the *entropic repulsion*. When the interface is forced to stay above a fixed wall, the entropic term in the Gibbs distribution pushes the interface upward, in opposition to the energy term that favors minimizing gradients. To model this situation mathematically, one considers a sequence of expanding domains $\Lambda_L := \{1, 2, \dots, L\}^2 \subset \mathbb{Z}^2$ and studies the field η conditioned to be nonnegative on Λ_L . That is, one defines the conditional measure

$$\bar{\pi}_{\Lambda_L}^p(\cdot) := \hat{\pi}_{\Lambda_L}^p(\cdot \mid \Omega_{\Lambda_L}^+), \quad \text{where } \Omega_{\Lambda_L}^+ = \{\eta : \eta_x \geq 0 \text{ for all } x \in \Lambda_L\}.$$

This describes a random interface constrained to remain above a wall. A notable feature of this conditioning is that the interface's roughness remains essentially unchanged, while the entire interface shifts upward, keeping a finite distance from the wall. This behavior reflects the interface's tendency to preserve its freedom to fluctuate, which lies at the core of the entropic repulsion phenomenon. From a mathematical standpoint, entropic repulsion has been rigorously investigated in various settings. In particular, the phenomenon was first rigorously analyzed in the context of SOS, Discrete Gaussian, and related random interface models by Bricmont, El Mellouki, and Fröhlich (1986) [6], while subsequent works by Deuschel, Giacomin, and others [3, 4] investigated it in the setting of the Gaussian free field (GFF). Such behavior is observed not only in the classical SOS model but also in generalized p -SOS models and continuous deformations thereof.

Building on this foundation, more recent research has focused on sharper quantitative properties of the interface. In particular, the study of one-point large deviations and the equilibrium behavior in p -SOS models has become increasingly active [17]. In this context, it is natural to extend the investigation of entropic repulsion to the p -SOS setting, where a first step is to ensure the existence of the infinite-volume Gibbs measure. For sufficiently large β (low temperature), $\hat{\pi}^p$ is well-defined and unique, by standard low-temperature methods such as contour expansion and the Pirogov–Sinai theory. In the SOS ($p = 1$) and DG ($p = 2$) cases, this was proved in [6] (Section 2.1), where it is shown that the height distribution has an exponential tail and correlations decay exponentially. The same reasoning extends to general $p \geq 1$.

In the entropic repulsion regime, the p -SOS interface concentrates around a well-defined typical height, with nearly all heights lying close to the same level. The *typical height under entropic repulsion*, which depends on both the parameter p and the system size L , is defined based on [17] by

$$H(p, L) := \max \left\{ h \in \mathbb{N} : \hat{\pi}^p(\eta_0 \geq h) \geq \frac{5\beta}{L} \right\}. \quad (2)$$

Here and throughout, we write $a \asymp b$ if both $a = O(b)$ and $b = O(a)$. The large deviation rate of $\hat{\pi}^p(\eta_0 \geq h)$ and the corresponding typical height $H(p, L)$ for different values of p are summarized in Table 1, based on the results of [8, 9, 17]. While equilibrium properties of entropic repulsion are by now well understood for both the SOS model ($p = 1$) and its generalizations to $1 \leq p \leq \infty$ (see e.g. [17]), the dynamical behavior has remained much less explored. In particular, the phenomenon of slow mixing arising from entropic repulsion was established by Caputo et al. [9], who established an exponentially slow mixing time in the SOS case $p = 1$. Beyond this single case, no analogous lower bound was known, not even for the classical Discrete Gaussian model ($p = 2$). In this work, we focus on extending the lower bound on the mixing time under entropic repulsion, previously

obtained for the SOS model in [9], to the more general p -SOS model. Since our work builds directly on the framework of [9], several of our formulations and terminology follow their presentation, with adaptations to the generalized p -SOS setting.

TABLE 1. Large deviation rate of $\hat{\pi}^p(\eta_0 \geq h)$ and the corresponding typical height $H(p, L)$ for different values of p . Here c_p denotes a constant depending only on p , and $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow \infty$.

p	$-\log \hat{\pi}^p(\eta_0 \geq h)$	$H(p, L)$
1	$4\beta h + \varepsilon_\beta$	$\left\lfloor \frac{1}{4\beta} \log L \right\rfloor$
$1 < p < 2$	$(c_p \beta + o(1))h^p$	$\left(\frac{1+o(1)}{c_p \beta} \log L \right)^{\frac{1}{p}+o(1)}$
$p = 2$	$(2\pi\beta + o(1)) \frac{h^2}{\log h}$	$\left(\frac{1+o(1)}{4\pi\beta} \log L \log \log L \right)^{\frac{1}{2}+o(1)}$
$2 < p < \infty$	$\asymp \beta h^2$	$\asymp \left(\frac{\log L}{\beta} \right)^{\frac{1}{2}}$

1.2. p -SOS Model, Associated Gibbs Measures, and Glauber Dynamics. We next set up the Glauber dynamics (also known as heat bath dynamics or the Gibbs sampler) for the p -SOS model in $(2+1)$ dimensions, the process we shall analyze in this work. For clarity, we also collect the Gibbs measure notation that will be used throughout. Recall that $\Lambda_L := \{1, \dots, L\}^2 \subset \mathbb{Z}^2$ and let $\partial\Lambda_L$ denote its outer boundary. Given a boundary condition $\phi : \partial\Lambda_L \rightarrow \mathbb{Z}$, the Hamiltonian $H_{\Lambda_L}^{p,\phi}$ and the corresponding Gibbs measure $\hat{\pi}_{\Lambda_L}^{p,\phi}$ were introduced above. In addition to the unconstrained Gibbs measure $\hat{\pi}_{\Lambda_L}^{p,\phi}$, we shall work with two constrained versions:

- (i) With the constraint $\eta_x \geq 0$ for all $x \in \Lambda_L$ (no upper bound, i.e. $n_+ = \infty$), supported on

$$\Omega_{\Lambda_L}^+ := \{\eta \in \mathbb{Z}^{\Lambda_L} : \eta_x \geq 0 \ \forall x \in \Lambda_L\},$$

we denote the corresponding Gibbs measure by $\bar{\pi}_{\Lambda_L}^{p,\phi}$.

- (ii) With the restriction $0 \leq \eta_x \leq n_+$ for all $x \in \Lambda_L$, where $n_+ \in \mathbb{N}$ and $n_+ \geq \log L$, supported on

$$\Omega_{\Lambda_L, n_+}^+ := \{\eta \in \mathbb{Z}^{\Lambda_L} : 0 \leq \eta_x \leq n_+ \ \forall x \in \Lambda_L\},$$

we denote the corresponding Gibbs measure by $\pi_{\Lambda_L}^{p,\phi}$.

We next define the Glauber dynamics as a discrete-time Markov chain $(\eta(t))_{t \geq 0}$ on $\Omega_{\Lambda_L, n_+}^+$. At each step:

- A site $x \in \Lambda_L$ is chosen uniformly at random;
- The height η_x is resampled from $\pi_{\Lambda_L}^{p,\phi}(\cdot \mid \eta_{\Lambda_L \setminus \{x\}})$.

Let $\mathbb{P}_\nu$ denote the law of the chain with initial distribution ν , and let μ_t^ν be its distribution at time t . When starting from a fixed configuration η , we write μ_t^η . This Markov chain is reversible with respect to $\pi_{\Lambda_L}^{p,\phi}$. The total variation mixing time is defined as

$$t_{\text{mix}} := \inf \left\{ t > 0 : \max_{\eta \in \Omega_{\Lambda_L, n_+}^+} \|\mu_t^\eta - \pi_{\Lambda_L}^{p,\phi}\|_{\text{TV}} \leq \frac{1}{4} \right\}.^1$$

gives an equivalent definition of the mixing time up to constant factors.

¹The particular choice of $1/4$ is inessential: any fixed $\varepsilon \in (0, 1)$ would give an equivalent definition of the mixing time up to constant factors.

1.3. Main result. We now turn to our main results concerning the Glauber dynamics of the p -SOS model under entropic repulsion. While equilibrium aspects of entropic repulsion are by now well understood, rigorous dynamical results had previously been limited to the SOS case $p = 1$ [9]. Our aim here is to establish exponential lower bounds on the mixing time, showing that entropic repulsion induces slow mixing uniformly across the entire regime $1 < p < \infty$.

Theorem 1. Fix $1 < p < \infty$. For sufficiently large inverse temperature β , there exists a constant $c := c(\beta, p) > 0$ such that the following holds. Consider the $(2 + 1)$ -dimensional p -SOS model on Λ_L with zero boundary condition, a floor at 0, and a ceiling at n_+ satisfying $n_+ \geq \log L$. For all sufficiently large $L \in \mathbb{N}$, the mixing time t_{mix} of the associated Glauber dynamics satisfies

$$e^{cL} \leq t_{\text{mix}}. \quad (3)$$

Beyond the global lower bound on the mixing time established in Theorem 1, one may ask how the dynamics actually approaches equilibrium in finer detail. In fact, starting from the flat configuration at height zero, the interface does not jump directly to its typical height $H(p, L)$, but rather moves through a cascade of intermediate metastable stages. At each stage the interface remains essentially flat, and upward transitions to the next level occur only through very rare large-deviation events, corresponding to the nucleation of a sufficiently large droplet at that height. This picture suggests that equilibration is not a smooth ascent, but rather a step-by-step mechanism in which droplets of increasing height appear successively and expand to cover the system.

Before presenting our second main result, let us emphasize the distinction with Theorem 1. While the global bound there applies uniformly to all $p \in (1, \infty)$, the refined analysis below reveals a dependence on the specific value of p . In particular, the dynamics requires an exponentially long time to lift the interface from the floor up to intermediate levels $aH(p, L)$, and the precise exponent reflects the strength of fluctuations encoded by p . This behavior is naturally interpreted in terms of metastability: the growth proceeds layer by layer, and each additional step must overcome an entropic barrier whose cost scales with the system size.

To formalize this dependence, we introduce the notation

$$d(p) := \begin{cases} p, & 1 < p < 2, \\ 2, & 2 \leq p < \infty. \end{cases} \quad (4)$$

Theorem 2. Fix $1 < p < \infty$ and sufficiently large inverse temperature β . There exists $c := c(\beta, p) > 0$ such that the following holds. Consider the Glauber dynamics $(\eta(t))_{t \geq 0}$ for the p -SOS model on $\Lambda = \{1, \dots, L\}^2$ with zero boundary condition, a floor at 0, and a ceiling at n_+ with $n_+ \geq \log L$. The dynamics starts from the all-zero configuration. For any $a \in (0, 1)$, define

$$\Omega_{p,a,L} = \left\{ \eta \in \Omega_{\Lambda_L, n_+}^+ : \#\{x \in \Lambda : \eta_x \geq aH(p, L)\} \geq \frac{9}{10}|\Lambda| \right\}.$$

Define

$$\tau_{p,a,L} = \min\{t \geq 0 : \eta(t) \in \Omega_{p,a,L}\}.$$

Then $\lim_{L \rightarrow \infty} \pi_{\Lambda}^p(\Omega_{p,a,L}) = 1$, and

$$\lim_{L \rightarrow \infty} \mathbb{P}\left(e^{cL^{a^{d(p)}}} \leq \tau_{p,a,L}\right) = 1.$$

1.3.1. Metastable States Induced by Entropic Repulsion. The metastable evolution of the interface under entropic repulsion can be naturally described in terms of the droplet picture. By a droplet we mean a finite region where the interface locally attains height $(h + 1)$, embedded in a background at

typical height h . In this sense, a droplet corresponds to a local phase $(h + 1)$ appearing within the dominant phase h , and its energetic cost is determined by the length of the surrounding contour. Thus, the nucleation probability of a droplet is directly governed by the geometric properties of its contour. This perspective is consistent with the classical analysis of metastability in the two-dimensional Ising model through droplet nucleation [18], as well as with the SOS model dynamics studied in [9]. In the present work, we extend this philosophy to the generalized p -SOS model. Inside the domain, the droplet interior typically remains close to the average height $aH(p, L)$, while near the boundary the interface gradually decreases due to the zero boundary condition. It is therefore more natural to regard the profile as smoothly interpolating between $aH(p, L)$ and 0, rather than dropping abruptly. The energetic barrier is then concentrated along the droplet edge, where the interface must absorb a height difference of order $aH(p, L)$.

Since the p -SOS Hamiltonian contains gradient terms of the form $|\nabla\eta|^p$, the energy cost associated with such a boundary discrepancy scales as $(aH(p, L))^p$. For $p = 1$, this cost is linear, making droplet nucleation comparatively easier, whereas for $p > 1$ the cost grows much faster, significantly suppressing the transition probability. As a consequence, for the same fraction $a \in (0, 1)$, the hitting time $\tau_{p,a,L}$ becomes substantially larger as p increases. This explains, on an intuitive level, why Theorem 2 exhibits a p -dependent exponent $d(p)$ in the lower bound.

In summary, the metastable states induced by entropic repulsion arise from the combined effect of entropy gain and boundary rigidity. This mechanism, already established in the Ising model [18] and in the classical SOS dynamics [9], persists in the p -SOS setting, where the rigidity effect is further amplified, leading to p -dependent time scales for mixing and equilibration.

1.4. Challenges and Strategy. In two dimensions, both discrete and continuous spin systems exhibit rich interface phenomena. When the spin values are unbounded, the discrete gradient model is typified by the SOS model, while its continuous counterpart is captured by the Gaussian Free Field (GFF). Each of these models possesses a well-developed set of structural and probabilistic properties. In what follows, we discuss the difficulties in generalizing from these well-understood frameworks and the conceptual obstacles encountered when transferring ideas from the SOS model and the GFF to the broader p -SOS context.

1.4.1. Continuous vs. Discrete gradient models. We examine the difference between discrete and continuous models from the perspective of gradient models. A prototypical example of a discrete gradient model is the integer-valued SOS model, while the Gaussian Free Field (GFF) represents the canonical case of a continuous gradient model.

In the SOS model, the height function takes integer values, and its structural properties make contour representations particularly effective. Analytical tools such as cylinders, level lines, and Peierls-type arguments are available, and these geometric approaches are well-suited to describe structural phenomena such as layering transitions, wetting, and entropic repulsion. In contrast, the GFF involves a real-valued height field and is well-analyzed through analytic tools such as Green functions, harmonic functions, and random walk representations. However, extending results from the SOS model to the generalized p -SOS model poses fundamental difficulties. For $p = 1$, the contour representation is well-defined and powerful, but for $1 < p < \infty$, the energy term becomes nonlinear, and such contour-based analyses are no longer directly applicable. Consequently, there exist significant technical barriers to generalizing the phenomena observed in the SOS model to the p -SOS setting.

1.4.2. *Technical Challenges in the p -SOS Model.* The p -SOS model is defined by a nonlinear energy term of the form

$$\sum_{x \sim y} |\eta_x - \eta_y|^p,$$

where $p > 1$. This nonlinearity leads to a fundamentally different mathematical structure compared to the classical SOS model ($p = 1$), introducing several technical challenges.

Many analytical techniques used in the SOS model—such as contour representation and Peierls-type arguments—rely on the subadditivity of the Hamiltonian and the ability to decompose height configurations into geometric objects. A key difficulty in extending Peierls-type arguments to the p -SOS model arises from the nonlinearity of the energy function. In the classical $p = 1$ SOS model, many arguments rely on the ability to shift the heights inside a given contour by a fixed amount without affecting the energy. For example, for any non-negative integers $x \leq y$ and any $h > 0$, one has

$$|y| - |x| = |y - x| = |y + h - (x + h)| = |y + h| - |x + h|,$$

so the energy difference remains invariant under such shifts.

However, for $p > 1$, this invariance no longer holds due to the nonlinearity of the function $x \mapsto |x|^p$. A simple derivative argument shows that for $x \leq y$ and $h > 0$,

$$|y + h|^p - |x + h|^p \geq p|x + h|^{p-1}|y - x|,$$

which implies that the energy difference depends on both x and the shift h . This lack of shift-invariance complicates the use of classical Peierls-type constructions, where contour-based mappings are frequently used to estimate large deviation probabilities.

As a consequence, while the one-point large deviation principle (LDP) for the SOS model can be established relatively easily (see Proposition 3.9 in [9]), the extension to the p -SOS model requires more sophisticated analytical tools. This is evident in [17], where Theorems 3.1, 5.1, and 5.5 involve techniques from harmonic and p -harmonic analysis to overcome the nonlinear nature of the energy landscape in the p -SOS setting. Consequently, extending some phenomena from SOS to the p -SOS model remains challenging. In particular, to estimate the lower bound on the mixing time under entropic repulsion, one needs a Peierls-type inequality adapted to this setting. We extended the inequality presented in Proposition 4.1 of [17] to the general p -SOS model. This allowed us to establish that the mixing time under entropic repulsion exhibits an exponential lower bound.

1.4.3. *Strategy.* To prove the exponential lower bound on the mixing time, we analyze the time it takes for the dynamics to reach a “typical” region of configurations that dominates the equilibrium measure. The argument follows two main steps. In the first step, we identify a target set of configurations that, at equilibrium, carries almost the entire probability mass. Intuitively, these are the configurations where the interface has already risen close to its entropically favored height. Demonstrating that this set has near-unit probability ensures that reaching it is essentially equivalent to reaching equilibrium. In the second step, we design an initial distribution supported on atypical configurations, where the interface is still close to the floor and lacks any macroscopic upward fluctuations. Starting from such an initial state, the dynamics must overcome significant energetic and entropic barriers to enter the typical region. We then estimate the probability that the dynamics makes this transition within an exponential time window. This requires two ingredients: controlling the likelihood of leaving the atypical region too quickly, and bounding the probability of creating the critical upward fluctuations that drive the system toward equilibrium. Both steps are shown to have exponentially small probabilities. Finally, the analysis is first carried out in the simpler

“floorless” setting, where technicalities are easier to manage, and then transferred back to the original model with a hard floor using standard comparison estimates. This strategy captures the core mechanism behind the slow mixing: the dynamics must wait for a rare fluctuation to initiate the transition to the typical height level, resulting in an exponentially large mixing time.

1.4.4. Resolution. We establish a general Peierls-type inequality for the p -SOS model in Theorem 3.3, extending classical arguments from [5]. This result generalizes the standard contour-based methods from the SOS case ($p = 1$) to the nonlinear setting $p > 1$. In particular, Theorem 3.3 includes a third term of the form $e^{-4\beta h}|\gamma|\log|\gamma|$, which, while slightly rougher than the bound presented in Proposition 4.1 of [17], has the advantage of being applicable uniformly across all p -SOS models. This term becomes negligible for sufficiently large β and h , and in fact vanishes naturally throughout the proof without affecting the core argument. Thus, this extension provides a robust foundation for applying Peierls-type techniques in the p -SOS setting, playing a crucial role in the subsequent analysis of the exponential lower bound on the mixing time.

The second key component concerns the definition and role of the typical height $H(p, L)$ under entropic repulsion. In the SOS case $p = 1$, Caputo et al. [9] showed that the typical height satisfies $H(1, L) = \frac{1}{4\beta} \log L + o(1)$. This was later sharpened in [8], where it was proved that $H(1, L) = \lfloor (1/4\beta) \log L \rfloor$, thus providing a precise characterization without any $o(1)$ error term. By contrast, for the p -SOS model, no such precise expression for $H(p, L)$ is available. Instead, [17] provides only a rough estimate. In our approach, we adopt the probabilistic criterion used there, defining $H(p, L)$ as the smallest height such that

$$\hat{\pi}_{\Lambda}^p(\eta_0 \geq H(p, L)) \geq \frac{5\beta}{L}.$$

This definition allows us to construct a suitable initial configuration without requiring an explicit formula for $H(p, L)$, thereby enabling the derivation of a lower bound on the mixing time. Moreover, this strategy offers a practical and computationally efficient way to extend techniques developed for the SOS model to the more general p -SOS setting.

1.5. Open Problems. This work represents a first step toward understanding the behavior of Glauber dynamics under entropic repulsion in the p -SOS model. However, several important questions remain open. Although we established an exponential lower bound on the mixing time, it is not yet clear whether this bound is sharp. In the classical SOS model, [17] proved an upper bound of $\exp(O(L \log L))$ by combining a refined canonical path argument with monotonicity, censoring techniques, and cluster expansion. In contrast, for the p -SOS model, it remains unclear whether a similar upper bound can be established. The difficulty lies in the lack of subadditivity and the increased complexity of energy penalties for large gradients. In this context, [10] developed a cluster expansion adapted to general p -SOS interfaces and provided tools that may be useful for deriving upper bounds, though a full mixing time analysis using these methods is still lacking.

A further open question concerns the sharpness of the exponential scale. While both lower and upper bounds of the form $\exp(cL)$ or $\exp(cL^a)$ have been established for the SOS model [9] and related settings, the precise value of the constant c has never been determined. It remains unclear whether c can be explicitly identified, or at least characterized, and whether this constant depends monotonically on p in the p -SOS model. Even understanding the qualitative behavior — for instance, whether c increases or decreases as p varies — would provide valuable insight into the dynamics under entropic repulsion.

Another natural direction is to extend the present results beyond the very low-temperature regime. While our analysis relies heavily on contour expansion techniques valid for large β , the behavior near the roughening transition β_R remains largely mysterious. In particular, obtaining rigorous mixing time estimates for p -SOS dynamics in the regime β just above β_R , where standard Peierls-type arguments break down, poses a significant challenge and requires new ideas beyond cluster expansion.

1.6. Organization. The remainder of the paper is organized as follows. In Section 2, we introduce the p -SOS model, the Glauber dynamics, and the main probabilistic tools, including the contour representation, the DLR formalism, and the FKG inequality. Section 3 presents the equilibrium properties of the model under entropic repulsion, such as the large deviation estimates and the typical height $H(p, L)$. In Section 4, we prove Theorems 1 and 2, establishing exponential lower bounds on the mixing time for the p -SOS model. Finally, Section 5 collects several technical computations that would otherwise interrupt the flow of the main text. In particular, detailed case-by-case analyses for $1 < p < 2$, $p = 2$, and $2 < p < \infty$ are deferred to the Appendix.

2. MODEL AND TECHNICAL FRAMEWORK

In this section we introduce the technical framework and probabilistic tools that will be used throughout the paper. We first recall the contour representation of the p -SOS model, which provides a convenient geometric description of height configurations and plays a central role in our large deviation and Peierls-type arguments. We then review the Dobrushin–Lanford–Ruelle (DLR) formalism for Gibbs measures, which allows us to formulate both finite and infinite-volume settings in a consistent way. Finally, we present an application of the FKG inequality, which will be repeatedly used to compare probabilities of monotone events under different boundary conditions.

2.1. Contour representation. We now introduce the notion of contours for the p -SOS model. Contours encode the level lines of the height function and allow us to describe deviations of the interface from its typical height in terms of geometric objects. In particular, they are the key tool in our large deviation estimates and in the analysis of metastable transitions, where the nucleation of critical droplets can be expressed through the formation of suitable contours.

Definition 1 (Contour). We denote by $\mathbb{Z}^{2*}$ the dual lattice of $\mathbb{Z}^2$, whose sites are placed at the centers of the unit squares of the original (primal) lattice. A *bond* in the dual lattice refers to a line segment connecting two neighboring sites in $\mathbb{Z}^{2*}$. Given two sites $x, y \in \mathbb{Z}^{2*}$, we say that they are *separated* by a dual bond e if the Euclidean distance from both x and y to e is exactly $1/2$. A pair of orthogonal dual bonds meeting at a vertex $x^* \in \mathbb{Z}^{2*}$ is called *linked* if both bonds lie on the same side of one of the two diagonal lines of slope ± 1 passing through x^* .

- (1) The bonds are distinct except for $e_0 = e_n$, forming a closed loop.
- (2) Each pair of consecutive bonds e_i, e_{i+1} shares a common vertex in the dual lattice $\mathbb{Z}^{2*}$.
- (3) If four bonds $e_i, e_{i+1}, e_j, e_{j+1}$ share a vertex $x^* \in \mathbb{Z}^{2*}$, then (e_i, e_{i+1}) and (e_j, e_{j+1}) must be *linked*, i.e., lie on the same side of one of the two 45° diagonals through x^* in the dual lattice.

We denote by $|\gamma|$ the length of γ , and by Λ_γ the set of sites in $\mathbb{Z}^2$ that are strictly enclosed by γ . We write $|\Lambda_\gamma|$ for the number of sites contained in the interior of γ .

The boundary associated with γ is defined as follows:

- $\partial\gamma$: the set of sites in $\mathbb{Z}^2$ that are adjacent to γ , either via Euclidean distance $1/2$ in the primal lattice from a dual bond in γ , or via distance $1/\sqrt{2}$ from a vertex in $\mathbb{Z}^{2*}$ where two non-linked bonds of γ meet.
- $\partial_\gamma^+ := \partial\gamma \cap \Lambda_\gamma$ and $\partial_\gamma^- := \partial\gamma \setminus \partial_\gamma^+$.
- $\eta|_{\partial_\gamma^+} \geq j$ to mean that $\eta_x \geq j$ for all $x \in \partial_\gamma^+$
- $\eta|_{\partial_\gamma^-} \leq j-1$ to mean that $\eta_y \leq j-1$ for all $y \in \partial_\gamma^-$.

Definition 2 (*h-contour*). Given a configuration η , a contour γ is said to be an *h-contour* if

$$\eta|_{\partial_\gamma^+} \geq h \quad \text{and} \quad \eta|_{\partial_\gamma^-} \leq h-1.$$

This represents a level drop across γ , from height at least h inside to at most $h-1$ outside.

Let $\mathcal{C}_{\gamma,h}$ denote the event that γ is an *h-contour* in configuration η , i.e.,

$$\mathcal{C}_{\gamma,h} := \left\{ \eta : \eta|_{\partial_\gamma^+} \geq h \quad \text{and} \quad \eta|_{\partial_\gamma^-} \leq h-1 \right\}.$$

Remark 2.1. For example, the elementary square contour of length 4 surrounding a site $x \in \mathbb{Z}^2$ is an *h-contour* if $\eta(x) \geq h$ and all its neighbors are at height at most $h-1$. This situation can also occur for multiple levels: if the inside height is at least $\max(h, h')$ and the outside height is at most $\min(h, h') - 1$, then the same geometric contour can simultaneously be an *h-contour* and an *h'-contour* for $h \neq h'$.

2.2. DLR Formalism and Gibbs Measures. Recall $\Lambda \subset \mathbb{Z}^2$ be a finite subset and denote by $\Omega_\Lambda = \mathbb{Z}^\Lambda$ the space of integer-valued configurations on Λ . A *Gibbs specification* is a family of probability kernels $\{\gamma_\Lambda^\phi\}_{\Lambda \subset \mathbb{Z}^2, \phi: \partial\Lambda \rightarrow \mathbb{Z}}$, where each γ_Λ^ϕ is defined via the Boltzmann–Gibbs distribution associated with the Hamiltonian of the model under consideration.

We consider the configuration space $\Omega = \mathbb{Z}^{\mathbb{Z}^2}$ equipped with the product σ -algebra $\mathcal{F}$ generated by local cylinder events. The set of probability measures on $(\Omega, \mathcal{F})$ is denoted by $\mathcal{M}_1(\Omega)$. A function $f: \Omega \rightarrow \mathbb{R}$ is said to be local if it depends only on finitely many coordinates, i.e., there exists $\Lambda \subset \mathbb{Z}^2$ such that $f(\eta) = f(\eta')$ whenever $\eta|_\Lambda = \eta'|_\Lambda$.

Definition 3 (DLR Condition). Let $\{\gamma_\Lambda^\phi\}$ be a Gibbs specification indexed by finite $\Lambda \subset \mathbb{Z}^2$ and boundary conditions $\phi: \partial\Lambda \rightarrow \mathbb{Z}$. A probability measure $\mu \in \mathcal{M}_1(\Omega)$ is said to satisfy the Dobrushin–Lanford–Ruelle (DLR) condition if, for every finite $\Lambda \subset \mathbb{Z}^2$ and every bounded measurable function $f: \Omega \rightarrow \mathbb{R}$,

$$\mu(f) = \int \gamma_\Lambda^{\eta|_{\partial\Lambda}}(f) \mu(d\eta).$$

This relation characterizes the consistency between the global measure μ and the family of local Gibbs specifications γ_Λ^ϕ . A probability measure that satisfies the DLR condition is referred to as an *infinite-volume Gibbs measure*. Importantly, such measures encode the local equilibrium structure of the system and allow a rigorous formulation of phase coexistence and correlation decay.

In the context of our study, we denote by $\hat{\pi}^p$ the infinite-volume Gibbs measure for the *p*-SOS model without floor or ceiling, and by $\bar{\pi}^p$ the corresponding measure under a hard floor at height zero. Both measures satisfy the DLR condition with respect to the appropriate finite-volume specifications. As we shall see in later sections, this structure is essential in justifying decompositions of probability expressions and applying recursive arguments based on domain restriction. For further background on the DLR formalism, we refer the reader to the standard reference by Friedli and Velenik [13].

2.3. FKG Inequality. The FKG inequality (named after Fortuin, Kasteleyn, and Ginibre) is a fundamental correlation inequality in the theory of lattice spin systems, first established in [12]. It

asserts that, under suitable conditions, increasing events are positively correlated. A Gibbs measure is said to have *attractive interactions* (or to be *ferromagnetic*) if the coupling constants are non-negative, so that neighboring variables tend to align. This monotonicity is precisely the condition under which the FKG inequality holds. Formally, a partial order can be defined on the configuration space Ω by pointwise comparison: for $\eta, \eta' \in \Omega$, we write $\eta \leq \eta'$ if and only if $\eta_x \leq \eta'_x$ for all $x \in \mathbb{Z}^2$. An event $E \subset \Omega$ is *increasing* if $\eta \in E$ and $\eta \leq \eta'$ imply $\eta' \in E$. Similarly, a function $f : \Omega \rightarrow \mathbb{R}$ is *nondecreasing* if $f(\eta) \leq f(\eta')$ whenever $\eta \leq \eta'$. For a Gibbs measure μ with attractive interactions, the FKG inequality states that for any pair of bounded nondecreasing local functions f and g ,

$$\mu(fg) \geq \mu(f)\mu(g),$$

where $(fg)(\eta) := f(\eta)g(\eta)$ denotes the pointwise product. In particular, if E and F are increasing events, then $f = \mathbf{1}_E$ and $g = \mathbf{1}_F$ yield $\mu(fg) = \mu(E \cap F) \geq \mu(E)\mu(F)$. In particular, if E and F are increasing events, their indicator functions $\mathbf{1}_E$ and $\mathbf{1}_F$ are nondecreasing, implying that the occurrence of one increases the likelihood of the other. A key feature of the FKG inequality is that it applies under arbitrary boundary conditions and general external fields, provided the interaction is attractive. Equivalently, in the setting of finite-volume Gibbs measures with boundary condition ϕ on a finite domain $\Lambda \subset \mathbb{Z}^2$, one has

$$\mu_\Lambda^\phi(fg) \geq \mu_\Lambda^\phi(f)\mu_\Lambda^\phi(g)$$

for all bounded nondecreasing local functions f and g . In our context, the FKG inequality will be used to compare probabilities under different boundary conditions and to deduce product-type lower bounds for certain increasing events. This property is particularly useful in the analysis of contour events, where the interior of a contour is influenced by its boundary values.

2.4. Exponential Decay of Correlations. For sufficiently large β , two-dimensional interface models such as the SOS, DG, and more generally the p -SOS model, exhibit exponential decay of truncated correlations between local observables. This was rigorously established in the classical work of Brandenberger and Wayne [5] via a convergent low-temperature cluster expansion.

Formally, let $f(\eta_{\Lambda_1})$ and $g(\eta_{\Lambda_2})$ be bounded local functions depending only on the disjoint finite sets $\Lambda_1, \Lambda_2 \subset \mathbb{Z}^2$. Then there exist constants $C_1, C_2 > 0$ (depending only on β and p) such that

$$|\hat{\pi}^p(fg) - \hat{\pi}^p(f)\hat{\pi}^p(g)| \leq C_1 \exp(-C_2 \text{dist}(\Lambda_1, \Lambda_2)),$$

where $\hat{\pi}^p$ denotes the infinite-volume Gibbs measure of the p -SOS model. Equivalently, truncated correlations of local observables decay exponentially in the distance between their supports.

3. EQUILIBRIUM PROPERTIES UNDER ENTROPIC REPULSION

The following result parallels Theorem 3.1 of [9] for the SOS model ($p = 1$) and Theorem 2 of [17] for the Discrete Gaussian model ($p = 2$). In the presence of entropic repulsion, nearly all heights in Λ_L concentrate around the typical height $H(p, L)$. In our case, we state a slightly weaker version, adapted from Proposition 5.1 of [10], to suit the p -SOS setting.

Theorem 3.1. *For any $p \geq 1$ and sufficiently large β , and for every $\varepsilon > 0$, there exist $K \in \mathbb{N}$ and $L_0 = L_0(\varepsilon, K, \beta, p)$ such that, for all $L \geq L_0$ and any box $\Lambda \subset \mathbb{Z}^2$ of side length L ,*

$$\bar{\pi}_\Lambda^p(\#\{v \in \Lambda : \eta_v \geq H(p, L) - K\}) \geq (1 - \varepsilon)|\Lambda| \geq 1 - \varepsilon.$$

In particular, the fraction of sites with height at least $H(p, L) - K$ tends to 1 in probability as $L \rightarrow \infty$.

For the mixing-time analysis, we shall work with a fixed ceiling height n_+ . However, the choice of ceiling has little effect on the stationary measure: the difference between the ceiling-constrained measure π_Λ^p (floor at 0, ceiling at n_+) and the floor-only measure $\bar{\pi}_\Lambda^{p,0}$ (no ceiling) is asymptotically negligible. This allows us to transfer estimates from $\bar{\pi}_\Lambda^{p,0}$ to $\pi_\Lambda^{p,0}$ at negligible cost. The next statement is the p -SOS analogue of Lemma 3.2 in [9].

Proposition 3.2. *Let $p \geq 1$ and $\beta \geq 1$ be sufficiently large. Then, for any event $A \subset \{0, \dots, n_+\}^\Lambda$, there exists a constant $c_1 > 0$ (as introduced in Proposition 3.5) such that*

$$\bar{\pi}_\Lambda^{p,0}(A) \leq \pi_\Lambda^{p,0}(A) \leq \left[1 + 2\left(c_1 L \log L e^{-2\beta n_+} + L^2 \hat{\pi}^p(\eta_0 \geq n_+/2) + L^{-4\beta+2}\right)\right] \bar{\pi}_\Lambda^{p,0}(A).$$

This inequality compares the full Gibbs measure $\pi_\Lambda^{p,0}$, defined with both floor and ceiling constraints, to the floor-only reference measure $\bar{\pi}_\Lambda^{p,0}$. The error terms arise from bounding the effect of the ceiling and the finite domain, and are small provided β and n_+ are large. The contour length term $c_1 L \log L$ corresponds to the outer boundary of the square domain Λ .

The following result provides a Peierls-type upper bound on the probability of the occurrence of a given h -contour in the p -SOS model with a floor at zero. For the SOS case ($p = 1$), an analogous statement was established in Proposition 3.6 of [9], while for the Discrete Gaussian ($p = 2$), the corresponding estimate appears as Proposition 4.1 in [17]. Here, we extend these bounds to the general p -SOS setting, combining Peierls estimates with the entropic repulsion effect to obtain a bound valid for all $p \geq 1$.

Theorem 3.3. *Consider the p -SOS model ($p \geq 1$) in a finite connected subset $\Lambda \subset \mathbb{Z}^2$ with a floor at 0 and zero boundary condition on $\partial\Lambda$. Let $\gamma \subset \Lambda$ be a fixed closed dual circuit. For sufficiently large β , there exists a constant $c_1 > 0$ such that*

$$\bar{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h}) \leq \exp(-\beta|\gamma| + \hat{\pi}^p(\eta_0 \geq h)|\Lambda_\gamma| + c_1 e^{-4\beta h} |\gamma| \log |\gamma|).$$

To prove Theorem 3.3, we first establish a Peierls-type inequality for the p -SOS model without ceiling constraint. While the inequality is structurally similar to the one known for the standard SOS model, the nonlinear nature of the p -norm interaction requires certain technical modifications in the argument, which we present for completeness.

The next lemma provides the basic Peierls bound in the absence of both floor and ceiling, for the p -SOS model with arbitrary boundary conditions. This estimate will serve as a fundamental ingredient in the proofs of the main results stated above. For the SOS case ($p = 1$), the same bound appears in Lemma 3.7 of [9], and here we extend it to the general p -SOS setting ($p \geq 1$) without any change in the form of the estimate.

Lemma 3.4. *Let $\Lambda \subset \mathbb{Z}^2$ be finite and let $p \geq 1$. For any contour $\gamma \subset \Lambda$ and any $h \in \mathbb{Z}$,*

$$\hat{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h}) \leq \exp(-\beta|\gamma|). \quad (5)$$

Moreover, if $h' < h$ and γ, γ' are contours with $\Lambda_\gamma \subseteq \Lambda_{\gamma'}$, then

$$\hat{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h} | \mathcal{C}_{\gamma',h'}) \leq \exp(-\beta|\gamma|). \quad (6)$$

Proof of Lemma 3.4. Let $E_\gamma := \{e = \{x, y\} : x \sim y, e \text{ crosses } \gamma\}$ and for each $e \in E_\gamma$ orient it so that $x \in \Lambda_\gamma$ and $y \notin \Lambda_\gamma$. By the h -contour condition (Definition 2), we have $\eta_x \geq h$ and $\eta_y \leq h - 1$, hence

$$\Delta_e := \eta_x - \eta_y \geq 1 \quad (e \in E_\gamma).$$

Since $f(t) = |t|^p$ is convex on $\mathbb{R}$ and in particular $g(m) := m^p - (m-1)^p$ is nondecreasing for integers $m \geq 1$, we have the discrete convexity bound

$$\Delta_e^p - (\Delta_e - 1)^p \geq 1 \quad \text{for all integers } \Delta_e \geq 1.$$

Define the map $T : \mathbb{Z}^\Lambda \rightarrow \mathbb{Z}^\Lambda$ by

$$(T\eta)_v = \begin{cases} \eta_v - 1, & \text{if } v \in \Lambda_\gamma, \\ \eta_v, & \text{otherwise.} \end{cases} \quad (7)$$

If a fixed circuit γ is an h -contour of η , then the bijection taking $\eta \mapsto \eta - \mathbf{1}_{\Lambda_\gamma}$. Then $e = \{x, y\} \in E_\gamma$ with the above orientation,

$$|(T\eta)_x - (T\eta)_y|^p = |\eta_x - 1 - \eta_y|^p = (\Delta_e - 1)^p \leq \Delta_e^p - 1 = |\eta_x - \eta_y|^p - 1.$$

If $e \notin E_\gamma$ (both endpoints inside or both outside), then $(T\eta)_x - (T\eta)_y = \eta_x - \eta_y$ and the contribution to the Hamiltonian is unchanged. Summing over all edges,

$$H_p(T\eta) = \sum_{e \notin E_\gamma} |\eta_x - \eta_y|^p + \sum_{e \in E_\gamma} |(T\eta)_x - (T\eta)_y|^p \leq H_p(\eta) - |E_\gamma|.$$

By definition of a contour, $|E_\gamma| = |\gamma|$, hence

$$H_p(T\eta) \leq H_p(\eta) - |\gamma|.$$

Therefore

$$\hat{\pi}_\Lambda^{p,0}(T\eta) = \frac{1}{Z} \exp(-\beta H_p(T\eta)) \geq e^{\beta|\gamma|} \hat{\pi}_\Lambda^{p,0}(\eta).$$

Since T is injective on $C_{\gamma,h}$, summing over $\eta \in C_{\gamma,h}$ yields

$$\hat{\pi}_\Lambda^{p,0}(C_{\gamma,h}) \leq e^{-\beta|\gamma|},$$

which proves (5).

For (6), let γ, γ' satisfy $\Lambda_\gamma \subset \Lambda_{\gamma'}$ and $h' < h$. Then

$$\hat{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h} \mid \mathcal{C}_{\gamma',h'}) = \frac{\hat{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h} \cap \mathcal{C}_{\gamma',h'})}{\hat{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma',h'})}.$$

Define $T_\gamma \eta := \eta - \mathbf{1}_{\Lambda_\gamma}$. If $\eta \in \mathcal{C}_{\gamma,h} \cap \mathcal{C}_{\gamma',h'}$, then $\eta|_{\partial_\gamma^+} \geq h$ and $\eta|_{\partial_\gamma^-} \leq h-1$. Since $\Lambda_\gamma \subset \Lambda_{\gamma'}$, the transformation affects only Λ_γ (strictly inside γ') and lowers values by 1, hence $T_\gamma \eta \in \mathcal{C}_{\gamma',h'}$. Moreover,

$$H(T_\gamma \eta) \leq H(\eta) - |\gamma| \quad \Rightarrow \quad \hat{\pi}_\Lambda^{p,0}(T_\gamma \eta) \geq e^{\beta|\gamma|} \hat{\pi}_\Lambda^{p,0}(\eta).$$

Because T_γ is injective on $\mathcal{C}_{\gamma,h} \cap \mathcal{C}_{\gamma',h'}$ and its image is contained in $\mathcal{C}_{\gamma',h'}$, we obtain

$$\sum_{\eta \in \mathcal{C}_{\gamma,h} \cap \mathcal{C}_{\gamma',h'}} \hat{\pi}_\Lambda^{p,0}(\eta) \leq e^{-\beta|\gamma|} \sum_{\eta' \in \mathcal{C}_{\gamma',h'}} \hat{\pi}_\Lambda^{p,0}(\eta') = e^{-\beta|\gamma|} \hat{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma',h'}),$$

which yields (6). $\square$

Unlike the classical SOS model, where the linear structure of the energy allows rather explicit control of marginal distributions, the p -SOS model is more sensitive to boundary effects due to the nonlinear dependence of the energy on gradients. In particular, computing one-point marginals near the boundary becomes technically delicate. Nevertheless, by combining the tools introduced above with the exponential decay estimates of Brandenberger–Wayne [5], one can still derive useful approximations near the boundary. Although these bounds are not sharp, the associated error terms are exponentially small in the system size and therefore negligible for our purposes. This suffices to

establish a large deviation principle (LDP) for the p -SOS model close to the boundary of a finite domain.

Proposition 3.5. *Fix $p \geq 1$ and for sufficiently large $\beta \gg 1$. Let $\Lambda \subset \mathbb{Z}^2$ be finite, $v \in \Lambda$, and $h \geq 0$. Then there exists a constant $c_1 > 0$ such that*

$$\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h) \leq \begin{cases} c_1 e^{-4\beta h}, & \text{if } \text{dist}(v, \partial\Lambda) \leq \log |\partial\Lambda|, \\ \hat{\pi}^p(\eta_0 \geq h) + |\Lambda|^{-4}, & \text{if } \text{dist}(v, \partial\Lambda) \geq \log |\partial\Lambda|. \end{cases}$$

In particular, for every v within distance $\log |\Lambda|$ of the boundary, $\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h) \leq c e^{-4\beta h}$.

Proof of Proposition 3.5. To obtain the upper bound on the one-point tail of the finite-volume measure, we begin by estimating the probability $\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h)$ at an arbitrary site $v \in \Lambda$. As indicated in the proposition, the resulting bound exhibits the same exponential decay in h as in the SOS model. This is due to the derivation of a Peierls-type inequality for the p -SOS model, analogous in spirit to the one used in the SOS setting (see Lemma 3.4). However, this estimate may not be sufficiently sharp for our purposes, particularly when $\text{dist}(v, \gamma) \geq \log |\gamma|$. In such cases, we will instead rely on a sharper estimate derived from the correlation decay result of [5].

Case 1. We first consider the case where the vertex v is located in the bulk of the domain Λ , and we apply a Peierls-type argument. Suppose that $\eta_v \geq h$. Then there exists a sequence of h nested i -contours $(\gamma_1, \dots, \gamma_h)$, each surrounding v , such that $\gamma_{i+1} \subset \gamma_i$ for all $i = 1, \dots, h-1$. Let us denote the set of such sequences by

$$\mathcal{A}(h, v) := \{(\gamma_1, \dots, \gamma_h) : v \in \Lambda_{\gamma_h}, \Lambda_{\gamma_{i+1}} \subset \Lambda_{\gamma_i} \text{ for all } i\}.$$

Using the contour weight bound (see Lemma 5),

$$\hat{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h}) \leq \exp(-\beta |\gamma|),$$

iteratively, we obtain

$$\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h) \leq \sum_{(\gamma_1, \dots, \gamma_h) \in \mathcal{A}(h, v)} \exp\left(-\beta \sum_{i=1}^h |\gamma_i|\right).$$

We now estimate the number of such nested contour sequences. For contours whose enclosed area lies between 2 and a fixed constant M , the number of admissible sequences satisfies the bound

$$|\mathcal{A}_M(h, v)| \leq |\mathcal{A}_M(1, v)|^{M-1} (h + M)^{M-2}.$$

Moreover, for any M sufficiently large (independent of β), the total contribution from sequences involving contours with area exceeding M is at most $e^{-6\beta}$. To proceed, we partition the sequence of contours as follows: - $\gamma_1, \dots, \gamma_k$ are those with area greater than M ; - $\gamma_{k+1}, \dots, \gamma_l$ have area at most M ; - $\gamma_{l+1}, \dots, \gamma_h$ are unit-square contours (area equal to 1). Note that in the final group, each contour satisfies $|\gamma_i| = 4$. Combining these estimates, we obtain

$$\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h) \leq \sum_{0 \leq k \leq l \leq h} |\mathcal{A}_M(l-k, v)| e^{-6\beta l} e^{-4\beta(h-l)} \leq e^{-4\beta h} \sum_{0 \leq k \leq l \leq h} C(l-k+M)^{M-2} e^{-2\beta l}.$$

This final sum is bounded by $c_1 e^{-4\beta h}$ for some absolute constant c_1 , which completes the argument for the first case.

Case 2. We now turn to the second case, where the vertex v lies at distance at least $\log |\gamma|$ from the boundary of Λ , and hence also from the nearest h -contour γ . In order to obtain a sharper estimate, we make use of the exponential decay of correlations valid for the p -SOS model at low temperatures. Specifically, Theorem 2 in [5] asserts that for sufficiently large β there exist constants

$0 < C, c(\beta) < \infty$ such that, for any pair of bounded local observables $A_{\Lambda_1}, B_{\Lambda_2}$ with disjoint supports $\Lambda_1, \Lambda_2 \subset \mathbb{Z}^2$,

$$\left| \hat{\pi}^p(A_{\Lambda_1} B_{\Lambda_2}) - \hat{\pi}^p(A_{\Lambda_1}) \hat{\pi}^p(B_{\Lambda_2}) \right| \leq C e^{-c(\beta) \text{dist}(\Lambda_1, \Lambda_2)}.$$

By standard arguments, the same estimate holds uniformly in finite-volume measures $\hat{\pi}_\Lambda^{p,0}$, provided both supports lie in the bulk of Λ .

In our setting, we take $A_{\{v\}}(\eta) := \mathbf{1}_{\{\eta_v \geq h\}}$ and let $B_{\partial\Lambda}$ denote the indicator of the boundary configuration induced by the contour γ . Since v is at distance at least $\log|\gamma|$ from γ , the above bound yields

$$\left| \hat{\pi}_\Lambda^{p,0}(\eta_v \geq h) - \hat{\pi}^p(\eta_0 \geq h) \right| \leq C e^{-c(\beta) \log|\gamma|} = C |\gamma|^{-c} \leq |\gamma|^{-2\beta},$$

for sufficiently large β .

Finally, we combine the two regimes according to the distance from the boundary. If $\text{dist}(v, \partial\Lambda) \leq \log|\Lambda|$, Step 1 gives the uniform Peierls bound

$$\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h) \leq c_1 e^{-4\beta h}.$$

If $\text{dist}(v, \partial\Lambda) \geq \log|\Lambda|$, Step 2 yields

$$\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h) \leq \hat{\pi}^p(\eta_0 \geq h) + |\Lambda|^{-4}.$$

Putting the two cases together gives exactly the two-branch estimate in (3.5), as claimed. $\square$

The argument combines the tools introduced in Section 2.2 (DLR formalism) and Section 2.3 (FKG inequality), which will be used repeatedly to control correlations and to decompose probabilities with respect to appropriate boundary conditions.

Proof of Theorem 3.3. By Proposition 3.2, it is enough to establish the corresponding estimates for the measure $\bar{\pi}$ without a ceiling. We start with the estimate for $\mathcal{C}_{\gamma,h}$, combining a Peierls-type argument with FKG. Let T_γ be the map in (7), which lowers heights by 1 on Λ_γ and leaves the outside unchanged. For any configuration η such that $\eta \in \mathcal{C}_{\gamma,h}$ and $\eta|_{\Lambda_\gamma} > 0$, every edge crossing γ has a height drop of at least 1, hence

$$H^p(T_\gamma \eta) \leq H^p(\eta) - |\gamma|.$$

Therefore, for the Gibbs weights we have

$$e^{-\beta H^p(\eta)} \leq e^{-\beta|\gamma|} e^{-\beta H^p(T_\gamma \eta)}.$$

Since T_γ is injective on $\{\eta \in \mathcal{C}_{\gamma,h} : \eta|_{\Lambda_\gamma} > 0\}$, it follows that

$$\bar{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h} \cap \{\eta|_{\Lambda_\gamma} > 0\}) \leq e^{-\beta|\gamma|}. \quad (8)$$

To convert this into a bound under the floor measure, recall that

$$\bar{\pi}_\Lambda^{p,0}(A) = \hat{\pi}_\Lambda^{p,0}(A \cap \{\eta \geq 0\}) / \hat{\pi}_\Lambda^{p,0}(\eta \geq 0). \quad (9)$$

Conditioning on the configuration on $\partial_\gamma^\pm$ and using FKG/monotonicity, we obtain

$$\bar{\pi}_\Lambda^{p,0}(\eta \in \mathcal{C}_{\gamma,h}, \eta|_{\Lambda_\gamma} > 0) = \hat{\pi}_\Lambda^{p,0}(\mathbf{1}_{\{\eta_\gamma^- \leq h-1\}} \mathbf{1}_{\{\eta|_{\Lambda \setminus \Lambda_\gamma} \geq 0\}} \mathbf{1}_{\{\eta_\gamma^+ \geq h\}} \hat{\pi}_{\Lambda_\gamma}^{p,\eta_\gamma^+}(\eta|_{\Lambda_\gamma} > 0)) / \hat{\pi}_\Lambda^{p,0}(\eta \geq 0).$$

By monotonicity of the increasing event $\{\eta|_{\Lambda_\gamma} > 0\}$ in the outer boundary condition,

$$\hat{\pi}_{\Lambda_\gamma}^{p,\eta_\gamma^+}(\eta|_{\Lambda_\gamma} > 0) \geq \hat{\pi}_{\Lambda_\gamma}^{p,h}(\eta|_{\Lambda_\gamma} > 0) = \hat{\pi}_{\Lambda_\gamma}^{p,0}(\eta|_{\Lambda_\gamma} > -h).$$

Applying FKG to the increasing events $\{\eta_x < h\}$ over $x \in \Lambda_\gamma$ yields

$$\hat{\pi}_{\Lambda_\gamma}^{p,0}(\eta|_{\Lambda_\gamma} > -h) \geq \prod_{x \in \Lambda_\gamma} (1 - \hat{\pi}_{\Lambda_\gamma}^{p,0}(\eta_x \geq h)).$$

Consequently,

$$\pi_{\Lambda}^{p,0}(\eta \in \mathcal{C}_{\gamma,h}, \eta|_{\Lambda_\gamma} > 0) \geq \pi_{\Lambda}^{p,0}(\mathcal{C}_{\gamma,h}) \prod_{x \in \Lambda_\gamma^{\text{in}}} (1 - \hat{\pi}_{\Lambda_\gamma}^{p,0}(\eta_x \geq h)). \quad (10)$$

We now bound the product in (10). By the one-point bounds from Proposition 3.5 (applied inside $\Lambda_\gamma^{\text{in}}$), for sufficiently large β ,

$$\hat{\pi}_{\Lambda_\gamma}^{p,0}(\eta_x \geq h) \leq \begin{cases} c_1 e^{-4\beta h}, & \text{if } \text{dist}(x, \gamma) \leq \log |\gamma|, \\ \hat{\pi}^p(\eta_0 \geq h) + |\Lambda_\gamma|^{-2}, & \text{otherwise.} \end{cases}$$

Thus,

$$\prod_{x \in \Lambda_\gamma} (1 - \hat{\pi}_{\Lambda_\gamma}^{p,0}(\eta_x \geq h)) \geq \exp(-c_1 e^{-4\beta h} |\gamma| \log |\gamma|) \exp(-\hat{\pi}^p(\eta_0 \geq h) A(\gamma)). \quad (11)$$

Combining this with (8) and (10), we obtain

$$\pi_{\Lambda}^{p,0}(\mathcal{C}_{\gamma,h}) \leq \exp(-\beta |\gamma| + \hat{\pi}^p(\eta_0 \geq h) A(\gamma) + c_1 e^{-4\beta h} |\gamma| \log |\gamma|),$$

which is the desired bound. $\square$

We are now ready to prove Proposition 3.2 by combining the results established so far. Throughout this subsection we fix the box $\Lambda_L = \{1, \dots, L\}^2$. We write $\pi_{\Lambda_L}^{p,0}$ for the p -SOS measure with a floor at 0 and a ceiling at n_+ , and $\bar{\pi}_{\Lambda_L}^{p,0}$ for the floor-only measure (no ceiling).

Proof of Proposition 3.2. Let $Z_{\Lambda_L}^{p,0}$ and $\bar{Z}_{\Lambda_L}^{p,0}$ be their partition functions, respectively. We have the identity

$$\pi_{\Lambda_L}^{p,0}(A) = \bar{\pi}_{\Lambda_L}^{p,0}(A) \frac{\bar{Z}_{\Lambda_L}^{p,0}}{Z_{\Lambda_L}^{p,0}}.$$

Since $\bar{Z}_{\Lambda_L}^{p,0} \geq Z_{\Lambda_L}^{p,0}$, it follows that

$$\bar{\pi}_{\Lambda_L}^{p,0}(A) \leq \pi_{\Lambda_L}^{p,0}(A).$$

For the reverse inequality, note that

$$\frac{Z_{\Lambda_L}^{p,0}}{\bar{Z}_{\Lambda_L}^{p,0}} = \bar{\pi}_{\Lambda_L}^{p,0}(\eta \leq n_+), \quad \text{hence} \quad \pi_{\Lambda_L}^{p,0}(A) = \frac{\bar{\pi}_{\Lambda_L}^{p,0}(A)}{\bar{\pi}_{\Lambda_L}^{p,0}(\eta \leq n_+)} = \frac{\bar{\pi}_{\Lambda_L}^{p,0}(A)}{1 - \bar{\pi}_{\Lambda_L}^{p,0}(\bigcup_{v \in \Lambda_L} \{\eta_v > n_+\})}.$$

By a union bound,

$$\pi_{\Lambda_L}^{p,0}(A) \leq \frac{\bar{\pi}_{\Lambda_L}^{p,0}(A)}{1 - \sum_{v \in \Lambda_L} \bar{\pi}_{\Lambda_L}^{p,0}(\eta_v > n_+)}. \quad (12)$$

We now estimate the sum of ceiling-exceedance probabilities under $\bar{\pi}_{\Lambda_L}^{p,0}$. Split Λ_L into a boundary layer and an interior:

$$\Lambda_{\text{bd}} := \{v \in \Lambda_L : \text{dist}(v, \partial \Lambda_L) \leq \log |\Lambda_L|\}, \quad \Lambda_{\text{int}} := \Lambda_L \setminus \Lambda_{\text{bd}}.$$

By Proposition 3.5 (boundary estimate with $h = n_+/2$) there exists $c = c(\beta, p)$ such that, for all $v \in \Lambda_{\text{bd}}$,

$$\bar{\pi}_{\Lambda_L}^{p,0}(\eta_v > n_+) \leq c e^{-2\beta n_+}.$$

As $|\partial\Lambda_L| = O(L)$ and $\log|\Lambda_L| = \log(L^2) = 2\log L$, we obtain

$$\sum_{v \in \Lambda_{\text{bd}}} \bar{\pi}_{\Lambda_L}^{p,0}(\eta_v > n_+) \leq c L \log L e^{-2\beta n_+}. \quad (13)$$

For $v \in \Lambda_{\text{int}}$, Proposition 3.5 (bulk estimate) and monotonicity give

$$\bar{\pi}_{\Lambda_L}^{p,0}(\eta_v > n_+) \leq \hat{\pi}^p(\eta_0 \geq n_+/2) + |\Lambda_L|^{-2\beta}.$$

Summing over $|\Lambda_{\text{int}}| \leq L^2$ sites yields

$$\sum_{v \in \Lambda_{\text{int}}} \bar{\pi}_{\Lambda_L}^{p,0}(\eta_v > n_+) \leq L^2 \hat{\pi}^p(\eta_0 \geq n_+/2) + L^2 \cdot |\Lambda_L|^{-2\beta} = L^2 \hat{\pi}^p(\eta_0 \geq n_+/2) + L^{-4\beta+2}. \quad (14)$$

Combining (13) and (14),

$$\sum_{v \in \Lambda_L} \bar{\pi}_{\Lambda_L}^{p,0}(\eta_v > n_+) \leq c L \log L e^{-2\beta n_+} + L^2 \hat{\pi}^p(\eta_0 \geq n_+/2) + L^{-4\beta+2}.$$

Plugging this into (12) and using $(1-x)^{-1} \leq 1+2x$ for x small (here ensured by sufficiently large β and $n_+ \geq \log L$) gives

$$\pi_{\Lambda_L}^{p,0}(A) \leq \left[1 + 2(c L \log L e^{-2\beta n_+} + L^2 \hat{\pi}^p(\eta_0 \geq n_+/2) + L^{-4\beta+2}) \right] \bar{\pi}_{\Lambda_L}^{p,0}(A).$$

Together with $\bar{\pi}_{\Lambda_L}^{p,0}(A) \leq \pi_{\Lambda_L}^{p,0}(A)$, this proves the proposition. $\square$

4. PROOFS OF THEOREM 1 AND THEOREM 2

4.1. Proof of Theorem 1.

Proof of Theorem 1. We aim to prove that

$$e^{cL} \leq t_{\text{mix}}.$$

In order to establish this lower bound, it suffices to show that the dynamics requires exponentially long time to reach the typical region of configurations. By Theorem 3.1, let

$$h = H(p, L) - K,$$

where K is the constant appearing in Theorem 3.1. Define the event

$$\mathcal{B} = \left\{ \eta : \#\{x \in \Lambda_L : \eta_x \geq h+1\} \geq \frac{1}{2}L^2 \right\},$$

and let $\tau_{\mathcal{B}}$ denote the hitting time of $\mathcal{B}$. Since

$$\pi_{\Lambda}^{p,0}(\mathcal{B}) = 1 - o(1), \quad (15)$$

we conclude that if

$$\tau_{\mathcal{B}} \geq e^{cL},$$

then necessarily $\tau_{\text{mix}} \geq e^{cL}$. Thus, to complete the proof of the mixing-time lower bound, it suffices to show that there exists a sufficiently small constant $c > 0$ such that

$$\min_{\eta} \mathbb{P}^{\eta}(\tau_{\mathcal{B}} < e^{cL}) = o(1). \quad (16)$$

In other words, starting from the worst possible initial configuration, the dynamics will, with high probability, require at least exponential time to reach the typical region $\mathcal{B}$.

Since $\mathcal{B}$ is an increasing event, we can reduce the problem to analyzing the dynamics starting from the minimal configuration $\sqcup$ (all heights equal to zero). Indeed, for any initial law ν we have

$$\min_{\eta} \mathbb{P}^{\eta}(\tau_{\mathcal{B}} < e^{cL}) = \mathbb{P}^{\sqcup}(\tau_{\mathcal{B}} < e^{cL}) \leq \mathbb{P}^{\nu}(\tau_{\mathcal{B}} < e^{cL}),$$

where the equality follows from the monotonicity of the dynamics with respect to the initial condition.

Before defining the set A , let us explain the idea behind it. To prove the lower bound on the mixing time, we must show that reaching the typical set $\mathcal{B}$ requires creating a large fluctuation, which is exponentially unlikely. To capture this, we restrict the dynamics to start from an atypical region A where the interface remains close to the floor and no h -contour surrounds a region of area larger than δL^2 . In other words, A excludes configurations containing droplets that occupy a positive fraction of the system, so that any trajectory reaching $\mathcal{B}$ must first generate such a large contour at exponential cost. This motivates the following definition of A .

Having specified the atypical set A , we now choose the initial distribution ν as the equilibrium measure conditioned on A . To make this precise, let $\delta \in (0, \frac{1}{4})$ be a sufficiently small constant such that, the following holds:

$$\delta < \left[(\beta - \log \mu_2 - \varepsilon_L) \cdot \frac{2\bar{\pi}_\Lambda^{p,0}(\eta_0 \geq H+1)}{5\beta\bar{\pi}_\Lambda^{p,0}(\eta_0 \geq h)} \right]^2, \quad (17)$$

where

$$\varepsilon_L := c_1 \log(2(L+1)^2) \cdot e^{-4\beta c_1 \sqrt{\frac{\log L}{\beta}}},$$

and μ_2 denotes the connective constant in $\mathbb{Z}^2$. The choice of δ ensures that the energetic cost of creating large h -contours dominates over entropic effects, which will be crucial for preventing the dynamics from quickly escaping the initial configuration set.

Rearranging the inequality (17), we obtain

$$\lambda := \sqrt{\delta} \cdot \frac{5\beta\bar{\pi}_\Lambda^{p,0}(\eta_0 \geq h)}{2\bar{\pi}_\Lambda^{p,0}(\eta_0 \geq H+1)} + \varepsilon_L < \beta - \log \mu_2. \quad (18)$$

This bound on λ guarantees that the exponential decay rate in the Peierls-type estimate remains sufficiently large. We then take as the initial distribution ν the conditional measure $\pi_\Lambda^{p,0}(\cdot | A)$, where A is the event that there exists no h -contour γ with area exceeding δL^2 , that is,

$$A = \bigcap_{\gamma: |\Lambda_\gamma| > \delta L^2} (\mathcal{C}_{\gamma,h})^c.$$

In other words, we restrict the initial configuration to those without macroscopic h -contours, making it entropically unlikely to reach the target set $\mathcal{B}$ in subexponential time.

In the sequel, we denote by ∂A the internal boundary of A , defined as

$$\partial A := \{ \eta \in A : p(\eta, \eta') > 0 \text{ for some } \eta' \notin A \},$$

where $p(\cdot, \cdot)$ is the transition probability of the dynamics.

Intuitively, ∂A represents the set of configurations in A that can exit A in a single update. Let $\tau_{\partial A}$ be the hitting time of ∂A .

Up to the hitting time $\tau_{\partial A}$, if the chain starts from a configuration in $A \setminus \partial A$, its evolution is identical to that of a Glauber dynamics constrained to remain in A , whose reversible distribution is exactly $\nu = \pi_\Lambda^{p,0}(\cdot | A)$. Applying a union bound over integer times $0 \leq t \leq e^{cL}$, we arrive at

$$\begin{aligned} \mathbb{P}^\nu(\tau_{\mathcal{B}} < e^{cL}) &\leq \mathbb{P}^\nu(\tau_{\partial A} < e^{cL}) + \mathbb{P}^\nu(\tau_{\mathcal{B}} < e^{cL}, \tau_{\partial A} \geq e^{cL}) \\ &\leq e^{cL} [\nu(\partial A) + \nu(\mathcal{B})]. \end{aligned} \quad (19)$$

We next establish exponential decay for the two terms inside the bracket. To keep the exposition concise, we defer the technical estimates to the lemmas below. Throughout, we retain the notation

introduced earlier (in particular, $A, \tilde{A}, \delta, \varepsilon_L$ defined above). The following two lemmas summarize the required bounds on the boundary term $\nu(\partial A)$ and on the probability of the typical set $\nu(\mathcal{B})$.

Lemma 4.1. *For sufficiently large β , there exists $c_2 = c_2(\beta) > 0$ such that*

$$\nu(\partial A) \leq e^{-c_2 L}.$$

Lemma 4.2. *For sufficiently large β , there exists $c_3 = c_3(\beta) > 0$ such that*

$$\nu(\mathcal{B}) \leq e^{-c_3 L}.$$

From Lemma 4.1, we already know that the boundary contribution $\nu(\partial A)$ decays exponentially fast in L . Similarly, Lemma 4.2 ensures the same exponential decay for $\nu(\mathcal{B})$. Hence, both terms inside the bracket in (19) are exponentially small, and it remains to translate these bounds from the auxiliary floor-only measure to the original measure with both floor and ceiling.

Recall that for any measurable $A \subset \{0, \dots, n_+\}^\Lambda$, one has

$$\bar{\pi}_\Lambda^{p,0}(A) \leq \pi_\Lambda^{p,0}(A) \leq \left[1 + 2(c_1|\gamma| \log |\gamma| e^{-2\beta n_+} + |\Lambda| \bar{\pi}^p(\eta_0 \geq h) + |\Lambda|^{-2\beta+1})\right] \bar{\pi}_\Lambda^{p,0}(A),$$

where $\bar{\pi}_\Lambda^{p,0}$ denotes the floor-only p -SOS measure.

In the regime $\log L \leq n_+$ and for sufficiently large β , the correction terms in the square bracket are negligible:

$$c_1|\gamma| \log |\gamma| \cdot e^{-2\beta n_+} \leq 2c_1(L+1)^2 \log[2(L+1)^2] \cdot e^{-2\beta \log L} \leq L^3 \cdot L^{-2\beta} \leq L^{-2},$$

and by appendix 5.1, we get

$$\bar{\pi}^p(\eta_0 \geq n_+/2) \leq L^{-3}. \quad (20)$$

Thus, for large enough β ,

$$\pi_\Lambda^{p,0}(A) = (1 + o(1)) \bar{\pi}_\Lambda^{p,0}(A).$$

As a consequence, any exponentially small probability bound established under $\bar{\pi}_\Lambda^{p,0}$. In particular, this implies that all exponentially small probability bounds established under $\bar{\pi}_\Lambda^{p,0}$ carry over directly to the full measure $\pi_\Lambda^{p,0}$. In conclusion, the second term on the right-hand side of (19) is $o(1)$ whenever $c < \min(c_2, c_3)$, completing the proof. $\square$

We begin by estimating the first term in equation (19). The following lemma provides the required bound.

Proof of Lemma 4.1. Next, introduce the set

$$\tilde{A} := \bigcap_{\gamma: |\Lambda_\gamma| > \frac{\delta}{5} L^2} (\mathcal{C}_{\gamma,h})^c,$$

which consists of configurations with no h -contour of area larger than $\delta L^2/5$. Since modifying a single site can create at most four disjoint h -contours, it follows that $\partial A \subset A \setminus \tilde{A}$. Consequently,

$$\nu(\partial A) = \frac{\pi_\Lambda^{p,0}(\partial A)}{\pi_\Lambda^{p,0}(A)} \leq \frac{\pi_\Lambda^{p,0}(A \setminus \tilde{A})}{\pi_\Lambda^{p,0}(A)}.$$

We now aim to bound the ratio

$$\frac{\pi_\Lambda^{p,0}(A \setminus \tilde{A})}{\pi_\Lambda^{p,0}(A)} \quad (21)$$

and show that it decays exponentially in L . To this end, consider a contour γ with relative area $|\Lambda_\gamma|/L^2 \in (\frac{\delta}{5}, \delta)$. Define the map T_γ as in (7), acting on a configuration η by lowering every height

inside γ by one unit, i.e.,

$$(T_\gamma \eta)_v = \begin{cases} \eta_v - 1, & v \in \Lambda_\gamma, \\ \eta_v, & \text{otherwise.} \end{cases}$$

Whenever $\eta \in A \cap \mathcal{C}_{\gamma,h}$, the transformed configuration $T_\gamma \eta$ still belongs to A , which implies

$$\pi_\Lambda^{p,0}(A) \geq \sum_{\eta \in A \cap \mathcal{C}_{\gamma,h}, \eta|_{\Lambda_\gamma} > 0} \pi_\Lambda^{p,0}(T_\gamma \eta).$$

In the proof of Theorem 3.3, by equation (8), this sum can be bounded above by

$$e^{\beta|\gamma|} \cdot \pi_\Lambda^{p,0}(\eta|_{\Lambda_\gamma} > 0 \mid A, \mathcal{C}_{\gamma,h}) \cdot \pi_\Lambda^{p,0}(A \cap \mathcal{C}_{\gamma,h}). \quad (22)$$

It remains to control the middle factor. Similarly, in the proof of Theorem 3.3, from equation (11) we obtain a product-form lower bound, obtained via independence across sites and uniform one-point estimates, which yields

$$\pi_\Lambda^{p,0}(\eta|_{\Lambda_\gamma} > 0 \mid \mathcal{C}_{\gamma,h}) \geq \exp(-|\Lambda_\gamma| \bar{\pi}^p(\eta_0 \geq h) - c_1 |\gamma| \log |\gamma| e^{-4\beta h}). \quad (23)$$

Combining (22), (23), we deduce

$$\pi_\Lambda^{p,0}(\mathcal{C}_{\gamma,h} \mid A) \leq \exp(-\beta|\gamma| + |\Lambda_\gamma| \bar{\pi}^p(\eta_0 \geq h) + c_1 |\gamma| \log |\gamma| e^{-4\beta h}), \quad (24)$$

which is the desired contour probability bound.

By the planar isoperimetric inequality on $\mathbb{Z}^2$, any contour γ satisfies

$$|\Lambda_\gamma| \leq \frac{|\gamma|^2}{16}.$$

When $|\Lambda_\gamma| < \delta L^2$, this yields

$$\begin{aligned} |\Lambda_\gamma| \cdot \bar{\pi}^p(\eta_0 \geq h) &\leq \sqrt{\delta L^2} \cdot \sqrt{\frac{|\gamma|^2}{16}} \cdot \bar{\pi}^p(\eta_0 \geq h) \\ &= \sqrt{\delta} \cdot \frac{1}{2} L |\gamma| \cdot \bar{\pi}^p(\eta_0 \geq h) \\ &\leq \left[(\beta - \log \mu_2 - \varepsilon_L) \cdot \frac{2\bar{\pi}^p(\eta_0 \geq H+1)}{5\beta \bar{\pi}^p(\eta_0 \geq h)} \right] \cdot L |\gamma| \cdot \frac{1}{2} \cdot \bar{\pi}^p(\eta_0 \geq h) \\ &\leq (\beta - \log \mu_2 - \varepsilon_L) \cdot \frac{5\beta}{5\beta L} \cdot L |\gamma| \\ &= (\beta - \log \mu_2 - \varepsilon_L) |\gamma|. \end{aligned} \quad (25)$$

On the other hand, recall that $\log L \leq n_+$, so that

$$\begin{aligned} c_1 |\gamma| \log |\gamma| \cdot e^{-4\beta h} &\leq c_1 |\gamma| \log(2(L+1)^2) \cdot e^{-4\beta h} \quad (\text{since } |\gamma| \leq 2(L+1)^2) \\ &\leq c_1 |\gamma| \log(2(L+1)^2) \cdot e^{-4\beta \log L} \\ &= c_1 |\gamma| \cdot \exp(\log \log(2(L+1)^2) - 4\beta \log L) \\ &= o(|\gamma|). \end{aligned} \quad (26)$$

Hence, the right-hand side of (24) is bounded above by $e^{-(\beta-\lambda)|\gamma|}$. A union bound over contours γ with $|\Lambda_\gamma| > (\delta/5)L^2$ then proves (21).

In conclusion, the first term in the right-hand side of (19) is $o(1)$ if $c < c_2$. $\square$

We next turn to the second term in equation (19). The following lemma establishes the corresponding bound.

Proof of Lemma 4.2. First, we claim that for any short h -contour γ and any $v \in \Lambda_\gamma$, where “macroscopic” means that $|\gamma| < \log^2 L$, we have

$$\bar{\pi}_\Lambda^{p,0}(\eta_v \geq h+1 \mid \mathcal{C}_{\gamma,h}) \leq \frac{1}{4}. \quad (27)$$

Indeed, let $\mathcal{A} = \{\gamma' : v \in \Lambda_{\gamma'} \subseteq \Lambda_\gamma\}$. Then, by applying (7) from Lemma 3.4, we obtain

$$\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h+1 \mid \mathcal{C}_{\gamma,h}) \leq \sum_{\gamma' \in \mathcal{A}} \hat{\pi}_\Lambda^{p,0}(C_{\gamma',h+1} \mid \mathcal{C}_{\gamma,h}) \leq \sum_{\gamma' \in \mathcal{A}} e^{-\beta|\gamma'|} \leq \frac{1}{8},$$

for β sufficiently large, since the number of contours $\gamma' \in \mathcal{A}$ of length k is at most $k\mu_2^k$, using the fact that each such γ' crosses a horizontal line within distance k to the right of v . To transfer this estimate to the setting with a floor, observe that by Remark 3.8 in [9], we have

$$\bar{\pi}_\Lambda^{p,0}(\eta_v \geq h+1 \mid \mathcal{C}_{\gamma,h}) = \frac{\hat{\pi}_\Lambda^{p,0}(\eta_v \geq h+1, \eta|_{\Lambda_\gamma} \geq 0 \mid \mathcal{C}_{\gamma,h})}{\hat{\pi}_{\Lambda_\gamma}^{p,h}(\eta|_{\Lambda_\gamma} \geq 0 \mid \eta|_{\partial\gamma} \geq h)}. \quad (28)$$

We have already established that the numerator is at most $1/8$. For the denominator, by monotonicity,

$$\hat{\pi}_{\Lambda_\gamma}^{p,h}(\eta \upharpoonright_{\Lambda_\gamma} \geq 0) = \hat{\pi}_{\Lambda_\gamma}^{p,0}(\eta \upharpoonright_{\Lambda_\gamma} \geq -h) \geq 1 - \hat{\pi}^{p,0}(\eta \geq 0)|\Lambda_\gamma| - |\gamma| \log |\gamma| e^{-4\beta h}.$$

using Proposition 3.5 and a union bound over the sites in Λ_γ . Since $|\Lambda_\gamma| \leq |\gamma|^2 = O(\log^4 L)$, the last term is $L^{-1+o(1)}$, so the denominator in (28) is at least $7/8$.

Hence,

$$\bar{\pi}_\Lambda^{p,0}(\eta_v \geq h+1 \mid \mathcal{C}_{\gamma,h}) \leq \frac{1/8}{7/8} \leq \frac{1}{4},$$

as claimed in (27). With inequality (27) at hand, we now apply the same type of concentration argument that leads to

$$\bar{\pi}_\Lambda^{p,0}\left(\sum'_\gamma \# \{x \in \Lambda_\gamma : \eta_x \geq h+1\} \geq \frac{1}{2}L^2\right) \leq \exp\left(-c_2 L^{2-o(1)}\right), \quad (29)$$

but now with full detail. For each admissible h -contour γ (indexed in the primed sum), define the random variable

$$X_\gamma := \# \{x \in \Lambda_\gamma : \eta_x \geq h+1\}.$$

Conditioned on the event that all relevant γ are h -contours (i.e., $\eta_x \geq h$ on Δ_γ^+), the random variables $\{X_\gamma\}_\gamma$ are mutually independent, since the associated domains Λ_γ are disjoint. By inequality (27), the expectation of each X_γ satisfies

$$\mathbb{E}[X_\gamma] \leq \frac{1}{4}|\Lambda_\gamma|.$$

Furthermore, since X_γ counts the number of sites in Λ_γ where $\eta_x \geq h+1$ and each height η_x is trivially bounded by $\log L$ due to the ceiling, we have the uniform bound

$$X_\gamma \leq |\Lambda_\gamma| \leq (\log L)^4,$$

where we used that $|\gamma| \leq \log^2 L$ implies $|\Lambda_\gamma| \leq |\gamma|^2 \leq (\log L)^4$. Letting $Y := \sum'_\gamma X_\gamma$ be the total number of sites (across all such contours) at height at least $h+1$, we observe

$$\mathbb{E}[Y] \leq \frac{1}{4} \sum |\Lambda_\gamma| \leq \frac{1}{4}L^2.$$

Since the $\{X_\gamma\}$ are independent and uniformly bounded by $(\log L)^4$, we may apply Hoeffding's inequality to obtain

$$\mathbb{P}(Y \geq \frac{1}{2}L^2) \leq \exp\left(-\frac{2\left(\frac{1}{4}L^2\right)^2}{N(\log L)^8}\right),$$

where $N \leq L^2$ is the number of disjoint h -contours in the sum. Thus,

$$\mathbb{P}(Y \geq \tfrac{1}{2}L^2) \leq \exp\left(-\frac{L^4}{8L^2(\log L)^8}\right) = \exp\left(-\frac{L^2}{8(\log L)^8}\right).$$

Replacing constants appropriately, we recover inequality (29), where the summation Σ' is over all short h -contours γ .

Let $B = B(\eta)$ denote the collection of h -contours γ of a configuration η such that $|\gamma| > \log^2 L$:

$$B := \{\gamma : \gamma \text{ is an } h\text{-contour of } \eta \text{ with } |\gamma| > \log^2 L\}.$$

Suppose that for some $\theta > \log \mu_2$, we have the following uniform bound for all such contours γ with small area, as ensured by an inequality of the form (24):

$$\bar{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h}) \leq \exp(-\theta|\gamma|).$$

Indeed, in our setting, combining the estimate

$$\bar{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h} \mid A) \leq \exp(-\beta|\gamma| + \bar{\pi}^p(\eta_0 \geq h)|\Lambda_\gamma| + 2c|\gamma| \log |\gamma| \cdot e^{-4\beta h})$$

with the isoperimetric inequality $|\Lambda_\gamma| \leq \frac{|\gamma|^2}{16}$ and applying bounds (25) and (26), we obtain that for contours γ with $|\Lambda_\gamma| \leq \delta L^2$ and for $\delta > 0$ sufficiently small,

$$\bar{\pi}_\Lambda^{p,0}(\mathcal{C}_{\gamma,h} \mid A) \leq \exp(-(\beta - \lambda)|\gamma|) = \exp(-\theta|\gamma|),$$

where $\theta := \beta - \lambda > \log \mu_2$ provided that β is sufficiently large. Let $\gamma_1, \dots, \gamma_m$ be m disjoint h -contours with $|\gamma_i| > \log^2 L$ and disjoint interiors $\{\Lambda_{\gamma_i}\}_{i=1}^m$. Then, by independence over disjoint supports and the exponential bound on each contour, we have:

$$\bar{\pi}_\Lambda^{p,0}\left(\bigcap_{i=1}^m \mathcal{C}_{\gamma_i,h}\right) \leq \exp\left(-\theta \sum_{i=1}^m |\gamma_i|\right). \quad (30)$$

We now bound the total probability that such a collection of m disjoint long h -contours appears, by summing over all such contour configurations. For each γ_i of length $k_i = |\gamma_i|$, there are at most L^2 possible choices for its origin (starting point) and at most R^{k_i} possible choices for a self-avoiding path of length k_i . Therefore, for a fixed m -tuple of contour lengths $(k_1, \dots, k_m)$ with each $k_i > \log^2 L$, the number of such contour configurations is at most

$$\prod_{i=1}^m L^2 R^{k_i}.$$

Combining this with the exponential probability bound (30), we obtain

$$\sum_{k_1, \dots, k_m > \log^2 L} \left(\prod_{i=1}^m L^2 R^{k_i} e^{-\theta k_i} \right) = \left(\sum_{k > \log^2 L} L^2 R^k e^{-\theta k} \right)^m.$$

Finally, summing over all $m \geq 1$, we obtain the total probability bound:

$$\bar{\pi}_\Lambda^{p,0}(\exists \text{ disjoint long } h\text{-contours } \gamma_1, \dots, \gamma_m) \leq \sum_{m=1}^{\infty} \left(\sum_{k > \log^2 L} L^2 R^k e^{-\theta k} \right)^m.$$

Since $\theta > \log R$, the inner sum satisfies

$$\sum_{k > \log^2 L} L^2 R^k e^{-\theta k} \leq L^2 \sum_{k > \log^2 L} e^{-ck} = O\left(e^{-c \log^2 L}\right),$$

for some $c = \theta - \log R > 0$. Therefore, the entire series converges and satisfies

$$\sum_{m=1}^{\infty} \left(e^{-c \log^2 L} \right)^m = O \left(e^{-c \log^2 L} \right). \quad (31)$$

We now choose the threshold T to match the event in (31), which corresponds to the number of vertices v in the interiors of long h -contours γ such that $\eta_v \geq h + 1$ exceeding $\frac{1}{2}L^2$. By the isoperimetric inequality, each such γ satisfies $|\Lambda_\gamma| \leq \frac{1}{16}|\gamma|^2$. To ensure that the total area covered is at least $\frac{1}{2}L^2$, the total contour length must satisfy

$$\sum_{\gamma \in B} |\gamma| \geq CL \quad \text{for some constant } C > 0.$$

Therefore, we set $T = CL$ and deduce from the previous exponential tail bound:

$$\bar{\pi}_\Lambda^{p,0} \left(\sum_{\gamma \in B} |\gamma| \geq T \right) \leq \exp(-cT) = \exp(-cCL).$$

Combining this with the isoperimetric inequality again, we obtain that

$$\sum_{\gamma \in B} \# \{v \in \Lambda_\gamma : \eta_v \geq h\} \geq \frac{1}{2}L^2 \quad \Rightarrow \quad \sum_{\gamma \in B} |\gamma| \geq c'L,$$

for some constant $c' > 0$. Hence,

$$\bar{\pi}_\Lambda^{p,0} \left(\sum_{\gamma \in B}'' \# \{v \in \Lambda_\gamma : \eta_v \geq h + 1\} \geq \frac{1}{2}L^2 \right) \leq \exp(-c_3L), \quad (32)$$

where the summation $\sum''$ is over all long h -contours γ , and $c_3 > 0$ is an absolute constant. Hence, combining inequalities 29 and 32, we obtain the bound 4.2. $\square$

4.2. Proof of Theorem 2. Fix $a \in (0, 1)$ and $p \in (1, \infty)$. Our goal is to prove an exponential lower bound on the hitting time

$$\tau_{p,a} := \min\{t \geq 0 : \eta(t) \in \Omega_{p,a}\}, \quad \Omega_{p,a} := \{\eta : \#\{x : \eta_x \geq aH(p, L)\} \geq (1 - \varepsilon_\beta)|\Lambda|\}.$$

The argument parallels the proof of Theorem 1, with the only difference that we now work at the intermediate height

$$h := aH(p, L) - 1,$$

so the initial law is supported on configurations forbidding *large* h -contours at a lower level than in Theorem 1. This lowers the critical droplet area scale, and the precise exponent depends on p .

Recall the definition

$$d(p) := \begin{cases} p, & 1 < p < 2, \\ 2, & 2 \leq p < \infty. \end{cases} \quad (33)$$

The appearance of $d(p)$ stems from the cost of forming a macroscopic droplet at height $aH(p, L)$. For $1 < p < 2$, the boundary gradient term scales like $(\nabla \eta)^p$, so raising the interface in a region of area A entails an energetic penalty of order A^p . For $p \geq 2$, fluctuations exhibit Gaussian behavior and the dominant cost becomes quadratic, leading to $d(p) = 2$. Thus the transition at $p = 2$ in (33). By the one-point large deviation bounds (Table 1), for every $\varepsilon > 0$ and all L large enough (with β large),

$$\hat{\pi}^p(\eta_0 \geq aH(p, L) - 1) \leq L^{-a^{d(p)+\varepsilon}}. \quad (34)$$

The detailed proof of (34) and the role of $d(p)$ is given in Lemma 4.3 and Appendix 5.2.

Lemma 4.3. *For sufficiently large β and L , and fixed $a \in (0, 1)$. for p -SOS model on $\mathbb{Z}^2$. Then for every $\varepsilon > 0$ and all L large enough,*

(1) *If $1 < p < 2$, then*

$$\hat{\pi}^p(\eta_0 \geq aH(p, L) - 1) \leq L^{-a^p + \varepsilon}.$$

(2) *If $2 \leq p < \infty$, then*

$$\hat{\pi}^p(\eta_0 \geq aH(p, L) - 1) \leq L^{-a^2 + \varepsilon}.$$

Proof of Theorem 2. We aim to establish an exponential lower bound on the hitting time $\tau_{p,a}$, which depends on the parameters $a \in (0, 1)$ and $p \in (1, \infty)$. The proof closely follows that of Theorem 1. In particular, we will prove (16) under the modified choice

$$h := aH(p, L) - 1.$$

With this change, the only difference from Theorem 1 is the initial law for the Glauber dynamics: we condition on the same atypical event (4.1) as before, but evaluated at the lower height level $h = aH(p, L) - 1$. In effect, we forbid *large* h -contours at a lower threshold, which in the p -SOS model shifts the critical droplet *area* scale to $\delta L^{2a^{d(p)}}$ (see Lemma 4.3 for the origin of $d(p)$). For comparison, in the SOS case $p = 1$ the critical area scale is δL^{2a} , as shown in [9].

Fix a small constant $\delta = \delta(\beta, p, a) > 0$ and define the restricted event

$$A_{p,a} := \bigcap_{\gamma: |\Lambda_\gamma| > \delta L^{2a^{d(p)}}} (\mathcal{C}_{\gamma,h})^c, \quad (35)$$

i.e., $A_{p,a}$ is the event that there exists no h -contour γ whose enclosed area exceeds $\delta L^{2a^{d(p)}}$. We start the dynamics from the equilibrium measure conditioned on $A_{p,a}$, namely $\nu := \pi(\cdot | A_{p,a})$. The target set is

$$B := \left\{ \eta : \#\{x : \eta_x \geq h + 1\} \geq \frac{1}{2} L^2 \right\}.$$

By the same union-bound as in (19), we obtain

$$\mathbb{P}_\nu(\tau_B < e^{cL^{a^{d(p)}}}) \leq e^{cL^{a^{d(p)}}} \left\{ \nu(\partial A_{p,a}) + \nu(B) \right\}. \quad (36)$$

Consequently, if the right-hand side is $o(1)$, then $\tau_B \geq e^{cL^{a^{d(p)}}}$ with probability $1 - o(1)$. Since $B \subset \Omega_{p,a}$ and $\pi(\Omega_{p,a}) \rightarrow 1$ at equilibrium, the desired lower bound follows.

By the isoperimetric inequality $|\Lambda_\gamma| \leq |\gamma|^2/16$ and a Peierls-type estimate, we have

$$\pi(\mathcal{C}_{\gamma,h}) \leq \exp\left\{ -\beta|\gamma| + |\Lambda_\gamma| \hat{\pi}^p(\eta_0 \geq h) + c_0 e^{-4\beta h} |\gamma| \log |\gamma| \right\}.$$

Invoking Lemma 4.1 together with (34), and for contours satisfying $|\Lambda_\gamma| \leq \delta L^{2a^{d(p)}}$, we obtain

$$C_0 |\Lambda_\gamma| \hat{\pi}^p(\eta_0 \geq h) \leq C_0 \delta L^{2a^{d(p)}} \cdot L^{-a^{d(p)} + \varepsilon} = C_0 \delta L^{a^{d(p)} + \varepsilon}.$$

Choosing $\delta > 0$ and $\varepsilon > 0$ sufficiently small (and then L sufficiently large), the perimeter term $-\beta|\gamma|$ dominates both the area correction and the logarithmic factor $e^{-4\beta h} |\gamma| \log |\gamma|$. Summing over admissible contours and applying a union bound yields

$$\nu(\partial A_{p,a}) \leq \exp\left\{ -c_1 L^{a^{d(p)}} \right\}, \quad (37)$$

for some $c_1 = c_1(\beta, p, a) > 0$.

On the event $A_{p,a}$ every h -contour has area at most $\delta L^{2a^{d(p)}}$. To realize B , one needs (at least) a cluster of many such small h -contours whose *total* length is $\gtrsim L^{a^{d(p)}}$. Inside each contour, the

one-point cost of raising the height from h to $h+1$ is $L^{-a^{d(p)}+o(1)}$ by (34). Arguing as in Lemma 4.2, and combining FKG/monotonicity with an independence approximation and a large-deviation estimate (e.g., Hoeffding–Azuma), we obtain

$$\nu(B) \leq \exp\{-c_2 L^{a^{d(p)}}\}. \quad (38)$$

Substituting (37) and (38) into (36) yields

$$\mathbb{P}_\nu(\tau_B < e^{cL^{a^{d(p)}}}) \leq e^{cL^{a^{d(p)}}} \left(e^{-c_1 L^{a^{d(p)}}} + e^{-c_2 L^{a^{d(p)}}} \right) = o(1).$$

By monotonicity, the same lower bound holds for the chain started from the all-zero configuration. Finally, since $\pi(\Omega_{p,a}) \rightarrow 1$, we conclude that

$$\mathbb{P}(\tau_{p,a} \geq e^{cL^{a^{d(p)}}}) \rightarrow 1 \quad (L \rightarrow \infty).$$

□

5. APPENDIX

In this appendix we collect technical estimates and detailed computations that are used throughout the paper. In particular, several probabilistic bounds depend sensitively on the value of p , and verifying these requires a separate case-by-case analysis. To keep the main arguments streamlined, we defer these calculations here. We also provide complete proofs of auxiliary lemmas that were only stated in the main text.

5.1. Proof of inequality 20. We claim that for any p and whenever $\log L \leq n_+$,

$$\bar{\pi}^p(\eta_0 \geq n_+/2) \leq L^{-3}.$$

We use the large-deviation estimates in Table 1. Throughout, $o(1)$ -terms are with respect to $h, n_+ \rightarrow \infty$ and can be absorbed for large L (with β fixed and sufficiently large).

Case $p = 1$. From Table 1,

$$-\log \bar{\pi}^1(\eta_0 \geq h) = 4\beta h + \varepsilon_\beta, \quad \varepsilon_\beta \rightarrow 0 \text{ as } \beta \rightarrow \infty.$$

Thus

$$\bar{\pi}^1(\eta_0 \geq n_+/2) \leq \exp(-2\beta n_+ + o(n_+)) \leq \exp(-2\beta \log L + o(\log L)) \leq L^{-3}$$

for all large L , since β is fixed and sufficiently large.

Case $1 < p < 2$. Here

$$-\log \bar{\pi}^p(\eta_0 \geq h) = (c_p \beta + o(1)) h^p,$$

so

$$\bar{\pi}^p(\eta_0 \geq n_+/2) \leq \exp\left(-\frac{c_p \beta + o(1)}{2^p} n_+^p\right) \leq \exp\left(-\frac{c_p \beta}{2^p} (\log L)^p + o((\log L)^p)\right).$$

Since $p > 1$ and β is fixed and large, the exponent dominates $3 \log L$ for large L , yielding the claim.

Case $p = 2$. We have

$$-\log \bar{\pi}^2(\eta_0 \geq h) = (2\pi\beta + o(1)) \frac{h^2}{\log h}.$$

Hence

$$\bar{\pi}^2(\eta_0 \geq n_+/2) \leq \exp\left(-\frac{2\pi\beta + o(1)}{4} \frac{n_+^2}{\log n_+}\right) \leq \exp\left(-\frac{\pi\beta + o(1)}{2} \frac{(\log L)^2}{\log \log L}\right) \leq L^{-3}$$

for all large L , since $(\log L)^2 / \log \log L \gg \log L$.

Case $2 < p < \infty$. Here $-\log \bar{\pi}^p(\eta_0 \geq h) \asymp \beta h^2$, so for some $c'_p > 0$,

$$\bar{\pi}^p(\eta_0 \geq n_+/2) \leq \exp\left(-\frac{c'_p \beta}{4} n_+^2\right) \leq \exp\left(-\frac{c'_p \beta}{4} (\log L)^2\right) \leq L^{-3}$$

for all large L .

In all cases, the bound $\bar{\pi}^p(\eta_0 \geq n_+/2) \leq L^{-3}$ holds for sufficiently large L under the standing assumption that β is sufficiently large (depending only on p).

5.2. Proof of lemma 4.3.

Proof. We shall use Table 1 to derive an upper bound on $\hat{\pi}^p(\eta_0 \geq aH(p, L) - 1)$. Throughout, $o(1)$ denotes a quantity that vanishes as $L \rightarrow \infty$, and $a \in (0, 1)$ is fixed.

There exist constants $c_p = c_p(\beta) > 0$ such that, as $h \rightarrow \infty$,

$$-\log \hat{\pi}^p(\eta_0 \geq h) = \begin{cases} 4\beta h + \varepsilon_\beta, & p = 1, \\ (c_p \beta + o(1)) h^p, & 1 < p < 2, \\ (2\pi\beta + o(1)) \frac{h^2}{\log h}, & p = 2, \\ \asymp \beta h^2, & 2 < p < \infty, \end{cases} \quad h \rightarrow \infty,$$

where $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow \infty$, and $f \asymp g$ means $f = O(g)$ and $g = O(f)$.

Consequently, the typical height $H(p, L)$ satisfies

$$H(p, L) \sim \begin{cases} \frac{1}{4\beta} \log L, & p = 1, \\ \left(\frac{1}{c_p \beta} \log L\right)^{1/p}, & 1 < p < 2, \\ \left(\frac{1}{4\pi\beta} \log L \log \log L\right)^{1/2}, & p = 2, \\ \asymp \sqrt{\frac{\log L}{\beta}}, & 2 < p < \infty. \end{cases}$$

Case 1: $1 < p < 2$. By the one-point LDP and the definition of $H = H(p, L)$,

$$\hat{\pi}^p(\eta_0 \geq h) = \exp\left(-(c_p \beta + o(1)) h^p\right), \quad \exp\left(-(c_p \beta + o(1)) H^p\right) \geq \frac{5\beta}{L}.$$

Setting $h = aH - 1$ and noting that $(aH - 1)^p = a^p H^p (1 + o(1))$, we obtain

$$\hat{\pi}^p(\eta_0 \geq aH - 1) = \exp\left(-(c_p \beta + o(1)) (aH - 1)^p\right) \leq \exp\left(-(c_p \beta a^p + o(1)) H^p\right) \leq \left(\frac{5\beta}{L}\right)^{a^p} \leq (5\beta)^{a^p} L^{-a^p + o(1)}.$$

Hence, for any $\varepsilon > 0$ and all sufficiently large L ,

$$\hat{\pi}^p(\eta_0 \geq aH(p, L) - 1) \leq L^{-a^p + \varepsilon}.$$

Case 2: $p = 2$. By the one-point LDP for the discrete Gaussian case and the definition of $H = H(2, L)$,

$$\hat{\pi}^2(\eta_0 \geq h) = \exp\left(-(2\pi\beta + o(1)) \frac{h^2}{\log h}\right), \quad (2\pi\beta) \frac{H^2}{\log H} \geq \log \frac{5\beta}{L}.$$

For $h = aH - 1$ we have $\log(aH - 1) = \log H + \log(a - \frac{1}{H}) = \log H (1 + o(1))$ and $(aH - 1)^2 = a^2 H^2 (1 + o(1))$, hence

$$\frac{(aH - 1)^2}{\log(aH - 1)} = \frac{a^2 H^2}{\log H} (1 + o(1)).$$

Therefore,

$$\hat{\pi}^2(\eta_0 \geq aH - 1) = \exp\left(- (2\pi\beta + o(1)) \frac{(aH - 1)^2}{\log(aH - 1)}\right) = \exp\left(- a^2(1 + o(1)) \log L\right) = L^{-a^2 + o(1)}.$$

Consequently, for any $\varepsilon > 0$ and all sufficiently large L ,

$$\hat{\pi}^2(\eta_0 \geq aH(2, L) - 1) \leq L^{-a^2 + \varepsilon}.$$

Case 3: $2 < p < \infty$. Assume the two-sided bounds: for all sufficiently large h ,

$$C_1 \beta h^2 \leq I(h) := -\log \hat{\pi}^p(\eta_0 \geq h) \leq C_2 \beta h^2, \quad (39)$$

for some constants $0 < C_1 \leq C_2 < \infty$. Choose the typical height $H = H(p, L)$ so that

$$I(H(p, L)) = -\log \hat{\pi}^p(\eta_0 \geq H) \leq \log \frac{L}{5\beta}. \quad (40)$$

From (39)–(40) it follows that

$$\frac{1}{C_2 \beta} \log L (1 + o(1)) \leq H^2 \leq \frac{1}{C_1 \beta} \log L (1 + o(1)).$$

Fix $a \in (0, 1)$ and set $h = aH - 1$. Then

$$I(aH - 1) \geq C_1 \beta (aH - 1)^2 = C_1 \beta a^2 H^2 (1 + o(1)),$$

and therefore

$$\begin{aligned} \hat{\pi}^p(\eta_0 \geq aH - 1) &\leq \exp\left(- C_1 \beta a^2 H^2 (1 + o(1))\right) \\ &\leq \exp\left(- a^2 \frac{C_1}{C_2} \log L (1 + o(1))\right) = L^{-a^2(C_1/C_2) + o(1)}. \end{aligned}$$

Here $o(1) \rightarrow 0$ as $L \rightarrow \infty$.

Since the assumption (39) expresses $I(h) \asymp \beta h^2$ as $h \rightarrow \infty$, for any $\delta > 0$ and all sufficiently large h we may write

$$(1 - \delta) c_p \beta h^2 \leq I(h) \leq (1 + \delta) c_p \beta h^2.$$

Taking $C_1 = (1 - \delta) c_p$ and $C_2 = (1 + \delta) c_p$ yields

$$\frac{C_1}{C_2} \geq \frac{1 - \delta}{1 + \delta} \geq 1 - 2\delta.$$

Letting L be large (equivalently $H \rightarrow \infty$), choose $\delta = \delta(L) \rightarrow 0$ slowly to obtain

$$\hat{\pi}^p(\eta_0 \geq aH - 1) \leq L^{-a^2(1 - 2\delta) + o(1)}.$$

Given any $\varepsilon > 0$, by taking δ and $o(1)$ smaller than $\varepsilon/4$ and $\varepsilon/2$, respectively, we conclude that

$$\hat{\pi}^p(\eta_0 \geq aH(p, L) - 1) \leq L^{-a^2 + \varepsilon}$$

for all sufficiently large L . □

REFERENCES

- [1] Roland Bauerschmidt, Jiwoon Park, and Pierre-François Rodriguez. The discrete gaussian model, i. renormalisation group flow at high temperature. *The Annals of Probability*, 52(4):1253–1359, 2024.

- [2] Roland Bauerschmidt, Jiwoon Park, and Pierre-François Rodriguez. The discrete gaussian model, ii. infinite-volume scaling limit at high temperature. *The Annals of Probability*, 52(4):1360–1398, 2024.
- [3] Erwin Bolthausen, Jean-Dominique Deuschel, and Giambattista Giacomin. Entropic repulsion and the maximum of the two-dimensional harmonic. *The Annals of Probability*, 29(4):1670–1692, 2001.
- [4] Erwin Bolthausen, Jean-Dominique Deuschel, and Ofer Zeitouni. Entropic repulsion of the lattice free field. *Communications in mathematical physics*, 170(2):417–443, 1995.
- [5] R Brandenberger and CE Wayne. Decay of correlations in surface models. *Journal of Statistical Physics*, 27(3):425–440, 1982.
- [6] Jean Bricmont, A El Mellouki, and Jürg Fröhlich. Random surfaces in statistical mechanics: Roughening, rounding, wetting,... *Journal of statistical physics*, 42(5):743–798, 1986.
- [7] Jean Bricmont, Jean-Raymond Fontaine, and Joel L Lebowitz. Surface tension, percolation, and roughening. *Journal of Statistical Physics*, 29(2):193–203, 1982.
- [8] Pietro Caputo, Eyal Lubetzky, Fabio Martinelli, Allan Sly, and Fabio L Toninelli. Scaling limit and cube-root fluctuations in sos surfaces above a wall. *Journal of the European Mathematical Society*, 18(5):931–995, 2016.
- [9] Pietro Caputo, Eyal Lubetzky, Fabio Martinelli, Allan Sly, and Fabio Lucio Toninelli. Dynamics of (2+1)-dimensional sos surfaces above a wall: Slow mixing induced by entropic repulsion. *The Annals of Probability*, 42:1516–1589, 2014.
- [10] Pietro Caputo, Fabio Martinelli, and Fabio Lucio Toninelli. Entropic repulsion in $-\nabla \phi - p$ surfaces: a large deviation bound for all $p \geq 1$. *Bollettino dell’Unione Matematica Italiana*, 10(3):451–466, 2017.
- [11] Michael E Fisher. Walks, walls, wetting, and melting. *Journal of Statistical Physics*, 34(5):667–729, 1984.
- [12] Cees M Fortuin, Pieter W Kasteleyn, and Jean Ginibre. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics*, 22(2):89–103, 1971.
- [13] Sacha Friedli and Yvan Velenik. *Statistical mechanics of lattice systems: a concrete mathematical introduction*. Cambridge University Press, 2018.
- [14] Jürg Fröhlich and Thomas Spencer. Kosterlitz-thouless transition in the two-dimensional plane rotator and coulomb gas. *Physical Review Letters*, 46(15):1006, 1981.
- [15] Jürg Fröhlich and Thomas Spencer. The berezinskii-kosterlitz-thouless transition (energy-entropy arguments and renormalization in defect gases). In *Scaling and Self-Similarity in Physics: Renormalization in Statistical Mechanics and Dynamics*, pages 29–138. Springer, 1983.
- [16] G Gallavotti, A Martin-Löf, and S Miracle-Solé. Some problems connected with the description of coexisting phases at low temperatures in the ising model. In *Statistical mechanics and mathematical problems*, pages 162–204. Springer, 2007.
- [17] Eyal Lubetzky, Fabio Martinelli, and Allan Sly. Harmonic pinnacles in the discrete gaussian model. *Communications in Mathematical Physics*, 344:673–717, 2016.
- [18] Roberto H Schonmann and Senya B Shlosman. Wulff droplets and the metastable relaxation of kinetic ising models. *Communications in mathematical physics*, 194(2):389–462, 1998.
- [19] Harold Neville Vazeille Temperley. Combinatorial problems suggested by the statistical mechanics of domains and of rubber-like molecules. *Physical Review*, 103(1):1, 1956.
- [20] Harold NV Temperley. Statistical mechanics and the partition of numbers ii. the form of crystal surfaces. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 48, pages 683–697. Cambridge University Press, 1952.

(Seokun Choi) DEPARTMENT OF MATHEMATICAL SCIENCES, KAIST, SOUTH KOREA

Email address: seokunchoi@kaist.ac.kr