Tracy–Widom Fluctuations in 2D Random Schrödinger Operators

Marcin Kotowski¹ , Bálint Virág²

¹ Faculty of Mathematics, Informatics, and Mechanics, University of Warsaw, ul. Banacha 2, 02-097 Warsaw, Poland. E-mail: mk249015@mimuw.edu.pl. URL: <http://www.math.toronto.edu/~marcin/>

² Department of Mathematics, University of Toronto, Bahen Centre, 40 St. George St., Toronto, ON M5S 2E4, Canada. E-mail: balint@math.toronto.edu. URL: <http://www.math.toronto.edu/~balint/>

Received: 21 June 2018 / Accepted: 14 February 2019
Published online: 24 April 2019 – © The Author(s) 2019

Abstract: We construct a random Schrödinger operator on a subset of the hexagonal lattice and study its smallest positive eigenvalues. Using an asymptotic mapping, we relate them to the partition function of the directed polymer model on the square lattice. For a specific choice of the edge weight distribution, we obtain a model known as the log-Gamma polymer, which is integrable. Recent results about the fluctuations of free energy for the log-Gamma polymer allow us to prove Tracy–Widom type fluctuations for the smallest eigenvalue of the random Schrödinger operator. We also relate the distribution of its k smallest positive eigenvalues to the nonintersecting partition functions of order k .

1. Introduction

In this paper, we study the correspondence between certain random Schrödinger operators defined on a subset of the 2D hexagonal lattice and a statistical physics model known as the directed log-Gamma polymer model. The directed log-Gamma polymer on a square lattice is obtained by putting random weights on the vertices of the lattice, drawn from the inverse Gamma distribution, and considering up-right paths connecting the opposite corners of the square, where each path is weighted by the product of its vertices. One is then interested in various statistical properties of such paths. The model has recently attracted considerable attention [[Sep12](#), [BCR13](#), [KQ16](#)], as it is integrable, i.e. allows explicit computations.

We construct a 2D random Schrödinger operator H and a mapping which maps its eigenvalues onto certain quantities in the directed polymer model, called the partition functions. Using results about the fluctuations of free energy for log-Gamma polymers [[KQ16](#)], we prove Tracy–Widom GUE fluctuations for the smallest positive eigenvalue of H (Theorem [1.1](#)). Moreover, we provide a description of higher eigenvalues in terms of partition functions related to non-intersecting paths. Such objects arise naturally in the technique known as geometric Robinson–Schoensted–Knuth correspondence [[COSZ14](#)].

Fig. 1. The lattice

To our knowledge, our work is the first rigorous proof of Tracy–Widom fluctuations of eigenvalues for a random Schrödinger operator. There has been some previous interest in the physics literature in studying the occurrence of the Tracy–Widom distribution in random Schrödinger operators. In [SOnP07] (see also subsequent works [SLDOn15, PMSO09]), the authors provide nonrigorous and numerical evidence for the Tracy–Widom distribution of the Green’s function in a random Schrödinger operator on the square lattice with random vertex weights (in physics called the Anderson model), also noting a heuristic connection to directed polymer models.

We consider a random Schrödinger operator defined on a hexagonal lattice in the shape of a rhombus. Formally, let G_n be a subset of the hexagonal lattice consisting of $2n - 1$ levels, with level k , for $k = 0, \dots, 2n - 2$, containing $\min\{k, 2n - k - 2\}$ hexagons. The first and last level contain only a single edge. An example of such lattice for $n = 4$ is shown in Fig. 1. We will call horizontal edges blue (thick) and the remaining edges red (thin).

We consider edges equipped with random real-valued weights, where the weight of an edge e is denoted by w_e . The random Schrödinger operator H_n , acting on functions $f : G_n \rightarrow \mathbb{R}$, is the weighted adjacency operator on G_n :

$$(H_n f)(v) = \sum_{e=(v,w)} w_e f(w)$$

where the sum is over all edges adjacent to v .

We consider two models defined on the lattice G_n :

1. **(i.i.d. model)** All edge weights are drawn independently at random from some distribution X that is nonzero almost surely and satisfies $\mathbb{E}e^{-t \log|X|}, \mathbb{E}e^{t \log|X|} < \infty$ for some $t > 0$
2. **(mixed model)** Pick some parameters $\gamma, \beta > 0$. The red (thin) edges are given weight 1. Each blue (thick) edge is independently assigned a weight drawn from the Gamma distribution $\Gamma(\gamma, \beta)$.

Let $N = n^2$. Since the underlying lattice is bipartite, it is easy to see that the spectrum of H_n consists of pairs of eigenvalues $\pm \lambda_1, \dots, \pm \lambda_N$, with $\lambda_1 \geq \dots \geq \lambda_N \geq 0$.

In the mixed model, we prove the following theorem about the smallest positive eigenvalue of H_n :

Theorem 1.1. *Let λ_N be the smallest positive eigenvalue of H_n in the mixed model with parameters $\gamma, \beta > 0$ such that $\log \beta > \Psi(\gamma)$ (so that the logarithms of inverses of blue edge weights have positive mean). Then we have as $n \rightarrow \infty$:*

$$\mathbb{P} \left(\frac{-\log \lambda_N - (\bar{f}_\gamma + 2 \log \beta)n}{n^{1/3}} \leq r \right) \rightarrow F_{\text{GUE}} \left(\left(\frac{\bar{g}_\gamma}{2} \right)^{-1/3} r \right)$$

where Ψ is the digamma function, F_{GUE} is the GUE Tracy–Widom distribution function, $\bar{f}_\gamma = -2\Psi(\gamma/2)$ and $\bar{g}_\gamma = -2\Psi''(\gamma/2)$.

The i.i.d. model is arguably more natural and we expect the theorem to hold also in that case:

Conjecture 1.2. *Theorem 1.1 holds also in the i.i.d. model for appropriate choice of constants.*

In both models, we make the following conjecture generalizing Tracy–Widom fluctuations also to higher eigenvalues.

Conjecture 1.3. *Let $\lambda_{N,k}$ be the k th smallest positive eigenvalue of H_n . Let*

$$\alpha_{n,k} = \frac{-\log \lambda_{N,k} - (\bar{f}_\gamma + 2 \log \beta)n}{(n\bar{g}_\gamma/2)^{1/3}}$$

as in Theorem 1.1, with the same assumptions about γ and β . Then for any k , as $n \rightarrow \infty$, the tuple $(\alpha_{n,1}, \dots, \alpha_{n,k})$ converges in distribution to the top k points of the Airy point process.

By considering submatrices of H_n , this conjecture can be extended to multiple space-time values of the (conjectured) scaling limit of last passage percolation. In particular, it should be possible to get the Airy sheet [QR14] as a limit in this model. This is in contrast with standard random matrix eigenvalue models, for which the Airy sheet is not expected to arise as a limit.

We make a step toward Conjecture 1.3 by proving that in the mixed model, the product of the bottom k eigenvalues is related, up to order $n^{1/3}$, to the partition functions for k -tuples of non-intersecting paths. Such objects appear naturally while studying exact formulas related to the geometric RSK correspondence [COSZ14]. It has been conjectured that the ratios of such non-intersecting partition functions converge to the Airy point process (see e.g. [LD16] for a similar conjecture for the continuum polymer).

Theorem 1.4. *Let $\lambda_N, \dots, \lambda_{N-k+1}$ be the k smallest positive eigenvalue of H_n in the mixed model with parameters $\gamma, \beta > 0$. For any fixed $k \geq 1$, let $Z_n^{(k)}$ be the non-intersecting partition function of order k for the square lattice corresponding to the mixed model (Definition 3.1). Then for any γ, β such that $\log \beta > \Psi(\gamma)$ and any $\delta > 0$:*

$$\mathbb{P} \left(n^{-1/3} \left| -\log \prod_{i=1}^k \lambda_{N-i+1} - \log |Z_n^{(k)}| \right| > \delta \right) \rightarrow 0.$$

The theorem holds also for more general models, see Theorem 3.7.

Let us briefly describe another model corresponding to a solvable polymer, namely the Beta polymer [BC17]. Let $\text{Beta}(\alpha, \beta)$ be the Beta distribution with parameters α and β . For every blue edge e , consider independent variables B_e with distribution

$Beta(\alpha, \beta)$ and pick a constant $c > \max\{\Psi(\alpha + \beta) - \Psi(\alpha), \Psi(\alpha + \beta) - \Psi(\beta)\}$. Put weight 1 on e , weight $X_{e_1} = cB_e$ on the red edge e_1 going up-right from e and weight $X_{e_2} = c(1 - B_e)$ on the red edge e_2 going down-right from e . For such a model, if we consider $\alpha = \beta = 1$, i.e. the uniform distribution, and a lattice of size $m \times n$ lattice with $\frac{m}{n} \rightarrow \theta \neq 1$, analogues of Theorems 1.1 and 1.4 also hold (see discussion at the end of Sects. 3.1, Sect. 3.2). The same is conjectured for $\alpha \neq \beta$, but in that case Tracy–Widom fluctuations are not known rigorously.

In Theorems 1.1 and 1.4 we assume $\log \beta > \Psi(\gamma)$. We expect that for every $\gamma > 0$ there exists some $\beta_c < e^{\Psi(\gamma)}$ such that the above theorems hold for $\beta > \beta_c$ and do not hold for $\beta < \beta_c$. The reason is as follows (see Sect. 2 for the connection between eigenvalues and polymers). For β sufficiently small, the typical mean logarithmic weight of a polymer path is negative and outweighs the contribution to the partition function coming from the number of paths. Thus, the maximum over S, T in Theorem 2.5 will be realized for S equal to T , i.e. when the paths have length zero. It follows that the smallest positive eigenvalue will be, up to a factor of n^2 , equal to the inverse of the maximum of n^2 vertex weights, which has a different scaling with n than the exponential scaling in Theorem 1.1. The same discussion pertains to the scaling factor c in the definition of the Beta model.

We end this section with an outline of how the theorems are proved. In Sect. 2, we prove that the eigenvalues of the operator are equal to the square roots of the singular values of the directed weighted square lattice. These, in turn, happen to be related to the partition functions of the polymer model on the lattice (Theorem 2.5). Using this connection, in Sect. 3.1 we proceed to prove Theorem 1.4 using a technical combinatorial lemma whose proof is contained in Sect. A. Then, in Sect. 3.2, we prove Theorem 1.1 by exploiting known results about the fluctuations of the partition functions for the log-Gamma polymer.

2. Eigenvalues and Polymers

The results in this section are deterministic—we introduce the probabilistic part of the analysis in Sect. 3. In order to study the eigenvalues of a random Schrödinger operator on a graph G , we first prove a lemma allowing us to study instead singular values of a certain directed graph derived from G .

Lemma 2.1. *Let $G = (V, E)$ be a weighted bipartite graph on $2N$ vertices with bipartition $V = B \sqcup C$. Let w_e denote the weight of the edge e . Suppose that G admits a perfect matching $S \subseteq E$ with edges $e_i = (b_i, c_i)$, $b_i \in B$, $c_i \in C$, $i = 1, \dots, N$. Let $\tilde{G}$ be a weighted directed graph on N vertices, with vertex set S and with edges defined as follows. For each $e_i \in S$, we have a loop (e_i, e_i) with weight w_{e_i} . For each edge $f = (b_i, c_j) \notin S$, we have a directed edge (e_i, e_j) with weight w_f .*

Let A be the $2N \times 2N$ adjacency matrix of G and let $\tilde{A}$ be the $N \times N$ adjacency matrix of $\tilde{G}$. Then the eigenvalues λ_i of A are equal to $\pm\sigma_i$, where σ_i are the singular values of $\tilde{A}$.

Proof. Let $B = (b_1, \dots, b_n)$, $C = (c_1, \dots, c_n)$, ordered arbitrarily. Let us index the rows and columns of A with $(b_1, \dots, b_n, c_1, \dots, c_n)$. Then A has the block form:

$$A = \begin{pmatrix} 0 & \tilde{A} \\ \tilde{A}^T & 0 \end{pmatrix}.$$

Indeed, each edge (b_i, c_i) in G corresponds to the edge (e_i, e_i) in $\tilde{G}$, giving the diagonal entries of $\tilde{A}$. Each edge (b_i, c_j) for $i \neq j$ corresponds to an edge (e_i, e_j) in $\tilde{G}$, giving the off-diagonal entries. Clearly, the eigenvalues of A are equal to $\pm$ the square roots of eigenvalues of $\tilde{A}\tilde{A}^T$, which are simply the singular values of $\tilde{A}$. $\square$

We now construct a general mapping between singular values of a directed graph $\tilde{G}$ and partition functions of the polymer model on the same graph. The results are stated in generality, but will be used for directed graphs derived from the particular lattice G_n described in Sect. 1.

Let $\tilde{G}$ be a directed acyclic weighted graph on N vertices and let $\tilde{A}$ denote its adjacency matrix. Assume that every vertex has a loop with nonzero weight. This implies that $\tilde{A}$ is invertible. Indeed, consider the set of vertices with no incoming edges, which is nonempty since the graph is acyclic. Since the loop weights are nonzero, the equation $\tilde{A}f = 0$ implies that $f = 0$ at such vertices. We can then remove them and repeat until there are no vertices left, proving that $f \equiv 0$.

For $v, w \in \tilde{G}$, a path π from v to w is defined to be a sequence of edges connecting vertices $(v = u_1 \rightarrow u_2 \rightarrow \cdots \rightarrow u_n = w)$, where none of the edges are loops. We allow a path of length zero connecting a vertex v to itself. Let $\Pi(v, w)$ denote the set of all paths from v to w . We will say that a vertex v precedes w if there is a positive length path from v to w .

We define new weights on vertices and edges of $\tilde{G}$ in the following way. For a vertex u we put $w_u = (\tilde{A}_{u,u})^{-1}$ and for an edge $e = (u, v)$ we put $w_{u,v} = -\tilde{A}_{u,v}$. For a path $\pi = (u_1 \rightarrow \cdots \rightarrow u_n)$ let its weight $\text{wt}(\pi)$ be defined as:

$$\text{wt}(\pi) := \prod_{i=1}^{n-1} w_{u_i, u_{i+1}} \prod_{i=1}^n w_{u_i}. \quad (1)$$

Note that the weight of an empty path from u to itself is w_u .

For any $k \geq 1$, a pair of sequences of distinct vertices $S = (s_1, \dots, s_k)$, $T = (t_1, \dots, t_k)$ and a permutation σ of the set $\{1, \dots, k\}$, we let $\Pi_{\sigma, S, T}$ denote the set of all tuples of paths $\pi = (\pi_1, \dots, \pi_k)$ with π_i connecting s_i to $t_{\sigma(i)}$. For any such tuple we let $\text{wt}(\pi) := \prod_{i=1}^k \text{wt}(\pi_i)$. We let $\Pi_{\sigma, S, T}^{n.i.}$ denote the set of all such tuples with paths π_i non-intersecting (i.e. vertex disjoint).

Definition 2.2. Fix any $k \in \{1, \dots, N\}$ and two sequences of distinct vertices $S = (u_1, \dots, u_k)$, $V = (v_1, \dots, v_k)$. We define:

$$Z_{S, T}^{(k)} := \sum_{\sigma} \text{sgn}(\sigma) \sum_{\pi \in \Pi_{\sigma, S, T}^{n.i.}} \text{wt}(\pi).$$

For $k = 1$ we will simply write:

$$Z_{u, v} = \sum_{\pi \in \Pi_{\text{id}, \{v\}, \{u\}}} \text{wt}(\pi).$$

We put $Z_{u, v}$ equal zero if there are no paths from u to v .

For $u \in \tilde{G}$, let f_u denote the function defined on the vertices of $\tilde{G}$ by $f_u(v) = Z_{u, v}$. In particular, $f_u(v) = 0$ if v precedes u and $f_u(u) = w_u$. Let δ_u be the function equal to 1 on u and 0 otherwise.

Proposition 2.3. *The functions f_u satisfy $\tilde{A}f_u = \delta_u$.*

Proof. We clearly have:

$$(\tilde{A}f_v)(v) = \tilde{A}_{v,v}f_v(v) = 1.$$

For $v \neq w$, we have:

$$\begin{aligned} (\tilde{A}f_v)(w) &= \sum_{u \rightarrow w} \tilde{A}_{u,w}f_v(u) + \tilde{A}_{w,w}f_v(w) \\ &= \sum_{u \rightarrow w} \tilde{A}_{u,w} \sum_{\pi \in \Pi(v,u)} \text{wt}(\pi) + \tilde{A}_{w,w} \sum_{\pi \in \Pi(v,w)} \text{wt}(\pi) \\ &= -\tilde{A}_{w,w} \sum_{\tau \in \Pi(v,w)} \text{wt}(\tau) + \tilde{A}_{w,w} \sum_{\pi \in \Pi(v,w)} \text{wt}(\pi) = 0. \end{aligned}$$

□

The quantities $Z_{S,T}^{(k)}$ can be related to f_u using the well known Lindstrom-Gessel-Viennot formula [GV85] for expressing sums over non-intersecting paths as determinants:

Proposition 2.4. *For $S = (u_1, \dots, u_k)$, $T = (v_1, \dots, v_k)$ we have:*

$$Z_{S,T}^{(k)} = \det(f_{u_i}(v_j))_{i,j=1}^k. \quad (2)$$

Proof. The standard Lindstrom-Gessel-Viennot formula is usually formulated with weights only on the edges. To obtain 2 in the general case, consider a graph G' where for an edge (u, v) we put $w'_{u,v} = w_{u,v}w_u$ and $w'_u = 1$. By applying the standard Lindstrom-Gessel-Viennot formula to G' we obtain $Z_{S,T}^{(k)'} = \det(f'_{u_i}(v_j))_{i,j=1}^k$. The proof follows by noting that $Z_{S,T}^{(k)} = Z_{S,T}^{(k)'} \cdot \prod_{v \in T} w_v$ and $f_u(v) = f'_u(v) \cdot w_v$. □

Let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_N$ denote the singular values of $\tilde{A}^{-1}$.

Theorem 2.5. *For any $k = 1, \dots, N$, we have:*

$$\max_{S,T} |Z_{S,T}^{(k)}| \leq \prod_{i=1}^k \sigma_i(\tilde{A}^{-1}) \leq \binom{N}{k}^2 \cdot \max_{S,T} |Z_{S,T}^{(k)}|$$

where the maximum ranges over all pairs of sequences of distinct vertices $S = (u_1, \dots, u_k)$, $T = (v_1, \dots, v_k)$.

Proof. We first use the following formula for the product of the singular values [Hog06]:

$$\prod_{i=1}^k \sigma_i(\tilde{A}^{-1}) = \max \{ |\det(U^* \tilde{A}^{-1} V)| : U, V \in \mathbb{C}^{N \times k}, UU^* = VV^* = I_k \}. \quad (3)$$

For a sequence of distinct vertices S of size k and a matrix $M \in \mathbb{C}^{N \times k}$, let M_S denote the submatrix obtained by taking rows with indices corresponding to S . With this notation I_S is the matrix having columns equal to δ_s for $s \in S$, i.e. the coordinate vectors corresponding to vertices in S . We have $M_S = I_S^* M$.

Fig. 2. The lattice G_n and the corresponding directed lattice $\tilde{G}_n$

For the lower bound, for any $S = (u_1, \dots, u_k)$, $T = (v_1, \dots, v_k)$ we plug $U = I_S$, $V = I_T$ into (3). Note that by Proposition 2.3, the matrix $\tilde{A}^{-1}$ expressed in the basis consisting of δ_u has the functions f_u as its columns. Thus, by Proposition 2.4 we have $\det((I_S)^* \tilde{A}^{-1} I_T) = Z_{S,T}^{(k)}$, from which the lower bound follows.

For the upper bound, for any U, V we use the Cauchy-Binet formula twice:

$$\begin{aligned} \det(U^* \tilde{A}^{-1} V) &= \sum_S \det(U_S^*) \det((\tilde{A}^{-1} V)_S) = \sum_S \det(U^* I_S) \det(I_S^* \tilde{A}^{-1} V) \\ &= \sum_{S,T} \det(U^* I_S) \cdot \det(I_S^* \tilde{A}^{-1} I_T) \cdot \det(I_T^* V). \end{aligned}$$

Clearly, we have $|\det(U^* I_S)|, |\det(I_T^* V)| \leq 1$, so:

$$\max_{U,V} |\det(U^* \tilde{A}^{-1} V)| \leq \binom{N}{k}^2 \cdot \max_{S,T} |\det(I_S^* \tilde{A}^{-1} I_T)| = \binom{N}{k}^2 \cdot \left| \max_{S,T} Z_{S,T} \right|.$$

□

We will now apply the construction above to the hexagonal lattice G_n from the previous section. In the case of G_n , the perfect matching in Lemma 2.1 consists of blue edges. The corresponding directed graph $\tilde{G}_n$ is a directed square lattice of size $n \times n$, so in the above notation $N = n^2$, with a loop added to each vertex. Both lattices are shown in Fig. 2.

Remark 2.6. In the mixed model, all edges of the directed square lattice have weights -1 . Since each $\tilde{A}_{u,u}$ was drawn independently from the Gamma distribution $\Gamma(\gamma, \beta)$, each loop u has a weight $w_u = (\tilde{A}_{u,u})^{-1}$ drawn independently from the inverse Gamma distribution $\Gamma^{-1}(\gamma, \beta)$ (Definition 3.8).

Remark 2.7. In the i.i.d. model, all edges of the hexagonal lattice have i.i.d. weights distributed as some random variable X . This implies that on the directed square lattice each edge weight $w_{u,v} = -\tilde{A}_{u,v}$ is distributed as $-X$ and each vertex weight $w_u = (\tilde{A}_{u,u})^{-1}$ is distributed as $\frac{1}{X}$.

We remark that in general, given any polymer model on the square lattice with weights on edges and vertices, one can construct a corresponding random Schrödinger operator by putting inverses of vertex weights on the blue edges and edge weights on the red edges.

The only caveat is that our results relating eigenvalues of the operator to the partition function of the polymer (Theorem 3.7) require that the mean logarithmic weight of a path is typically positive. By multiplying all weights on the red edges by a suitable constant, one can ensure this condition while shifting the free energy $\log Z_n$ by a deterministic constant.

3. Probabilistic Results

We now introduce the probabilistic part of the analysis for the square lattice with random edge and vertex weights. We do not require edge weights to be independent, only that for each path its edges are independent. The square lattice considered is the one from Fig. 2 rotated 45 degrees counterclockwise, so that the lower left corner is the point $(1, 1)$ and the upper right corner is the point (n, n) . We fix some $k \geq 1$ and let $S_0 = ((1, 1), (1, 2), \dots, (1, k))$, $T_0 = ((n, n - k + 1), (n, n - k + 2), \dots, (n, n))$.

For an edge e we let X_e denote its weight and for a vertex u we let X_u denote its weight.

Definition 3.1. For the directed square lattice from $(1, 1)$ to (n, n) , the non-intersecting partition function of order k is defined as:

$$Z_n^{(k)} := Z_{S_0, T_0}^{(k)} = \sum_{\pi=(\pi_1, \dots, \pi_k)} \prod_{i=1}^k \text{wt}(\pi_i)$$

where the summation is over all tuples of k vertex disjoint paths, with π_i connecting $(1, i)$ to $(n, n - k + i)$.

In this section, we prove two results. First, in Sect. 3.1, we show that if the mean logarithmic weight of a path is positive, then up to order $n^{1/3}$ the product of k top singular values of $\tilde{A}_n^{-1}$ is with high probability close the quantities $Z_n^{(k)}$ (Theorem 3.7). By known results about fluctuations of the polymer partition function, this then implies (Theorem 3.10) Tracy–Widom fluctuations of the smallest singular value of $\tilde{A}_n$, for the weights drawn from the inverse Gamma distribution.

Whenever we say that an event holds with high probability (w.h.p.), it will mean that the probability that it does not hold is superpolynomially small in n .

3.1. Eigenvalues and non-intersecting partition functions. The goal of this section is the proof of Proposition 3.5, which combined with the results from Sect. 2 implies Theorems 1.4 and 3.7.

It will be convenient to work in the case when all vertex weights are equal to 1. The proposition below states that we can move the all the weights to the edges while asymptotically changing logarithms of the partition functions by less than $n^{1/3}$.

Proposition 3.2. Suppose that for all vertices u we have $\mathbb{E}e^{-t \log |X_u|} < \infty$, $\mathbb{E}e^{t \log |X_u|} < \infty$ for some $t > 0$ and likewise for edges. For an edge (u, v) put $X'_{u,v} = X_{u,v} \cdot X_v$ and put $X'_u = 1$ for all vertices u . Note that the primed edge weights are not independent, but they are still independent along every path. For any $k \geq 1$, $\delta > 0$ and all S, T we have:

$$\mathbb{P}(n^{-1/3} \left| \log \left| Z_{S,T}^{(k)} \right| - \log \left| Z_{S,T}^{(k)'} \right| \right| > \delta) \rightarrow 0.$$

Proof. We have $Z_{S,T}^{(k)} = Z_{S,T}^{(k)'} \cdot \prod_{u \in S} X_u$, so $\log |Z_{S,T}^{(k)}| = \log |Z_{S,T}^{(k)'}| + \sum_{u \in S} \log |X_u|$. Since $\mathbb{E}e^{-t \log |X_u|}, \mathbb{E}e^{t \log |X_u|} < \infty$ for some $t > 0$, by Markov inequality for each X_u we have:

$$\mathbb{P}(|\log |X_u|| > \delta n^{1/3}) < Ce^{-t\delta n^{1/3}}$$

for some constant $C > 0$. By union bounding over polynomially many choices of S we can assume that w.h.p. for all choices of S we have $\sum_{u \in S} |\log |X_u|| < Ck\delta n^{1/3}$, which finishes the proof. $\square$

Thus, below we assume that all vertex weights are equal to 1. We will write X'_e for the distribution of the primed weight of an edge e . Below we will work under the assumption $\mathbb{E} \log |X'_e| > 0$, which implies that the mean logarithmic weight of a path is positive.

Below we use notation introduced in Definition 2.2. All weights in the partition functions are now primed weights, distributed as X'_e . For any tuple of paths π , we let $E(\pi)$ denote the set of edges used by paths in π (if an edge is used by multiple paths we count it only once). Recall that $S_0 = ((1, 1), \dots, (1, k))$ and $T_0 = ((n, n-k+1), \dots, (n, n))$. The set of all non-intersecting tuples contributing to $Z_n^{(k)}$ is $\Pi_{\text{id}, S_0, T_0}^{n, i.}$.

The proof of Proposition 3.5 will follow from the lemma below, which is purely combinatorial and whose proof we defer to Sect. A. The lemma roughly says that any nonintersecting tuple connecting S to T can be modified into a nonintersecting tuple connecting S_0 to T_0 while removing only a constant number of edges and adding a constant number of path segments.

Lemma 3.3. *There exists a constant C depending only on k such that for any σ, S, T with $S \neq S_0$ or $T \neq T_0$ there exists a set of paths $\mathcal{P}$ which has size at most $C \cdot n^C$ and satisfies the following property. For every $\pi \in \Pi_{\sigma, S, T}^{n, i.}$ there exist a tuple $\pi' \in \Pi_{\text{id}, S_0, T_0}^{n, i.}$ such that $|E(\pi) \setminus E(\pi')| \leq C$ and $E(\pi) \Delta E(\pi')$ is a union of paths whose number is at most C and which are all elements of $\mathcal{P}$.*

We note that the lemma is obvious in the case $k = 1$, since it suffices to connect $S = \{s\}$ to $(1, 1)$ and $T = \{t\}$ to (n, n) with any two fixed paths.

We will use the following inequality, which is a straightforward consequence of a standard large deviation inequality for i.i.d variables [Dur10]. Note that we use the fact that edge weights are independent along every path.

Lemma 3.4. *Assume that $\mathbb{E}e^{-t \log |X'_e|} < \infty$ for some $t > 0$ and $\mathbb{E} \log |X'_e| = \mu > 0$. For any path ρ of length $m \leq n$ and $a < \mu$:*

$$\mathbb{P}\left(\sum_{e \in \rho} \log |X'_e| < ma\right) \leq e^{-I(a)m} \quad (4)$$

for some rate function I such that $I(a) > 0$ for $a < \mu$.

Proposition 3.5. *For any fixed $k \geq 1$ and $\delta > 0$, if $\mathbb{E} \log |X'_e| > 0$ and $\mathbb{E}e^{-t \log |X'_e|} < \infty$, $\mathbb{E}e^{t \log |X'_e|} < \infty$ for some $t > 0$, we have:*

$$\mathbb{P}\left(\frac{1}{n^{1/3}} \left| \log |Z_n^{(k)}| - \max_{S, T} \log |Z_{S, T}^{(k)}| \right| > \delta\right) \rightarrow 0.$$

Proof. We need to prove that for any $k \geq 1$ and $\delta > 0$, we have w.h.p. for some global constant D depending only on k :

$$\left| Z_n^{(k)} \right| \geq \max_{S,T} \left| Z_{S,T}^{(k)} \right| \cdot e^{-D\delta n^{1/3}}.$$

Letting:

$$Z_{\sigma,S,T} := \sum_{\pi \in \Pi_{\sigma,S,T}^{n.i.}} \text{wt}(\pi)$$

we clearly have:

$$\left| Z_{S,T}^{(k)} \right| \leq k! \cdot \max_{\sigma} \left| Z_{\sigma,S,T} \right|.$$

Thus, it suffices to prove that with high probability for all σ, S, T we have:

$$\left| Z_n^{(k)} \right| \geq \left| Z_{\sigma,S,T} \right| \cdot e^{-D\delta n^{1/3}}. \quad (5)$$

Consider a tuple $\pi \in \Pi_{\sigma,S,T}^{n.i.}$ contributing to $Z_{\sigma,S,T}$. By Lemma 3.3 there exists $\pi' \in \Pi_{\text{id},S_0,T_0}^{n.i.}$ such that $|E(\pi) \setminus E(\pi')| \leq C$ and $E(\pi) \Delta E(\pi')$ is a union of at most C paths which are elements of $\mathcal{P}$.

Let A be the event that all edges of the lattice have weights at most $e^{t\delta n^{1/3}}$ and let B be the event that all paths in $\mathcal{P}$ have weights at least $e^{-t\delta n^{1/3}}$. Note that both the number of edges in the lattice and the number of paths in $\mathcal{P}$ are polynomial in n .

Consider a path $\rho \in \mathcal{P}$. Let $x = \log \mathbb{E} e^{-t \log |X'_e|}$ and let $\varepsilon > 0$ be such that $\varepsilon x < t\delta/2$. If ρ has length $m < \varepsilon n^{1/3}$, by Markov inequality:

$$\mathbb{P} \left(\sum_{e \in \rho} \log |X'_e| < -t\delta n^{1/3} \right) \leq e^{mx - t\delta n^{1/3}} < e^{-1/2t\delta n^{1/3}}. \quad (6)$$

If ρ has length $m \geq \varepsilon n^{1/3}$, applying inequality (4) from Lemma 3.4 we have:

$$\mathbb{P} \left(\sum_{e \in \rho} \log |X'_e| < -t\delta n^{1/3} \right) < \mathbb{P} \left(\sum_{e \in \rho} \log |X'_e| < 0 \right) \leq e^{-\varepsilon n^{1/3} I(0)}. \quad (7)$$

Thus, by union bounding over a polynomial size family of events we have that B holds with high probability. Likewise, by Markov inequality and union bound over all edges of the lattice the event A also holds with high probability. Note that the events B depends only on $\mathcal{P}$, which depends only on σ, S, T and not on the tuple π .

Let $E(\pi) \Delta E(\pi') = \cup_{i=1}^m \pi_i$, where each path π_i belongs to $\mathcal{P}$ and $m \leq C$. Since A and B hold w.h.p., we have with high probability:

$$\text{wt}(\pi') \geq \text{wt}(\pi) \cdot e^{-2tC\delta n^{1/3}}. \quad (8)$$

We now need to sum Eq. (8) over all paths $\pi \in \Pi_{\sigma,S,T}^{n.i.}$. The map taking π to π' need not be injective. However, note that $E(\pi) \Delta E(\pi')$ is a union of at most C paths, all of which lie inside $\mathcal{P}$, which has size at most $C \cdot n^C$. Thus, each π' will have at most polynomially many preimages. Thus, summation of Eq. (8) over all possible $\pi \in \Pi_{\sigma,S,T}^{n.i.}$ proves the desired inequality (5). $\square$

Remark 3.6. In the case $\mathbb{E} \log |X'_e| = 0$, the logarithmic weight of a typical path is of order $\sqrt{n}$ and the proof does not apply. However, the same proof can be used to obtain a weaker statement, namely, for any $\varepsilon > 0$:

$$\mathbb{P} \left(\frac{1}{n^{1/2+\varepsilon}} \left| \log |Z_n^{(k)}| - \max_{S,T} \left| \log Z_{S,T}^{(k)} \right| \right| > \delta \right) \rightarrow 0.$$

By combining Proposition 3.5 and Theorem 2.5, we arrive at the following theorem.

Theorem 3.7. *For any fixed $k \geq 1$ and $\delta > 0$, if $\mathbb{E} \log |X'_e| > 0$ and $\mathbb{E} e^{-t \log |X'_e|} < \infty$, $\mathbb{E} e^{t \log |X'_e|} < \infty$ for some $t > 0$, we have:*

$$\mathbb{P} \left(n^{-1/3} \left| \log \prod_{i=1}^k \sigma_i(\tilde{A}^{-1}) - \log |Z_n^{(k)}| \right| > \delta \right) \rightarrow 0.$$

We now proceed to prove Theorem 1.4. To this end, let us first note the following properties of the inverse Gamma distribution.

Definition 3.8. A random variable X has inverse Gamma distribution with parameters $\gamma, \beta > 0$, denoted $\Gamma^{-1}(\gamma, \beta)$, if its probability distribution is supported on positive reals with density:

$$\mathbb{P}(X \in dx) = \frac{\beta^\gamma}{\Gamma(\gamma)} x^{-\gamma-1} \exp \left(-\frac{\beta}{x} \right) dx.$$

Remark 3.9. Let $X \sim \Gamma^{-1}(\gamma, \beta)$ and let Ψ be the digamma function. Then $\mathbb{E} \log X = -\Psi(\gamma) + \log \beta$ and $\text{Var} \log X = \Psi'(\gamma)$. In particular, for $\log \beta > \Psi(\gamma)$ we have $\mathbb{E} \log X > 0$. Also, for small enough $t > 0$ we have $\mathbb{E} e^{-t \log X} < \infty$, $\mathbb{E} e^{t \log X} < \infty$.

Proof of Theorem 1.4. By Remark 2.6, in the mixed model the dual graph corresponds to the directed square lattice with inverse Gamma vertex weights. By Proposition 3.2, if instead we transfer the weights to the edges, the difference between the partition functions of the vertex weighted model and the edge weighted model, scaled by $n^{-1/3}$, converges to zero in probability. By Remark 3.9, for $\log \beta > \Psi(\gamma)$ the inverse Gamma logarithmic edge weights X'_e , obtained from Proposition 3.2, satisfy the assumptions of Theorem 3.7, so the theorem holds also for the vertex weighted case. The proof follows by invoking Lemma 2.1. $\square$

The same results also hold for the Beta model defined in Sect. 1. Recall that weights in that model are defined, for some constant $c > 0$, as $X_{e_1} = cB_e$ and $X_{e_2} = c(1 - B_e)$ for the red edges outgoing from a given blue edge e , where B_e has distribution $\text{Beta}(\alpha, \beta)$. The corresponding polymer model has weights on the edges going out of a given vertex equal to cB_e and $c(1 - B_e)$. Note that edge weights are not independent, but are still independent along every path. Since $\mathbb{E} \log B_e = \Psi(\alpha) - \Psi(\alpha + \beta)$ and $\mathbb{E} \log(1 - B_e) = \Psi(\beta) - \Psi(\alpha + \beta)$, by taking $c > \max\{\Psi(\alpha + \beta) - \Psi(\alpha), \Psi(\alpha + \beta) - \Psi(\beta)\}$ we can ensure that $\mathbb{E} \log X_{e_i} > 0$ and one can also check that $e^{t \log X_{e_i}}, e^{-t \log X_{e_i}} < \infty$ for some $t > 0$. The only difference from Proposition 3.5 is that for $\alpha \neq \beta$ the vertical and horizontal edges are not distributed identically. The only required change in the proof is modifying inequalities (6), (7) to allow edges from either distribution, which is straightforward to do.

3.2. Fluctuations of the smallest eigenvalue in the exactly solvable case. We now establish the fluctuations of the smallest eigenvalue in the mixed model (Theorem 1.1). Let us recall the definition of the discrete log-Gamma polymer [BCR13]. Let Γ_n be the square lattice where each vertex has a weight w_u drawn independently from $\Gamma^{-1}(\gamma)$ and all edges have weight 1. The log-Gamma polymer partition function is then:

$$Z_n = \sum_{\pi: (1,1) \rightarrow (n,n)} \prod_{u \in \pi} w_u.$$

By Remark 2.6, the lattice $\tilde{G}_n$ obtained in the mixed model is the same as in the log-Gamma polymer with $\beta = 1$, except that edges have weights -1 instead of 1. However, since every path $\pi : (1, 1) \rightarrow (n, n)$ has an even number of edges, the partition functions of the two models will be equal. Below we put $N := n^2$.

Theorem 3.10. *Let $\sigma_N(\tilde{A}_n)$ be the smallest singular value of the $N \times N$ adjacency matrix $\tilde{A}_n$ of the lattice $\tilde{G}_n$. Then for all $\gamma, \beta > 0$ such that $\log \beta > \Psi(\gamma)$ we have:*

$$\mathbb{P} \left(\frac{-\log \sigma_N(\tilde{A}_n) - (\bar{f}_\gamma + 2 \log \beta)n}{n^{1/3}} \leq r \right) \rightarrow F_{\text{GUE}} \left(\left(\frac{\bar{g}_\gamma}{2} \right)^{-1/3} r \right)$$

where Ψ is the digamma function, F_{GUE} is the GUE Tracy–Widom distribution function, $\bar{f}_\gamma = -2\Psi(\gamma/2)$ and $\bar{g}_\gamma = -2\Psi''(\gamma/2)$.

Proof. By Theorem 2.1 of [KQ16], we have for any $\gamma > 0$ and $\beta = 1$:

$$\mathbb{P} \left(\frac{\log Z_n - \bar{f}_\gamma n}{n^{1/3}} \leq r \right) \rightarrow F_{\text{GUE}} \left(\left(\frac{\bar{g}_\gamma}{2} \right)^{-1/3} r \right). \quad (9)$$

We note that this theorem was first proved in [BCR13], under the stronger assumption that $\gamma < \gamma'$ for some γ' . Consider the mixed model with $\gamma, \beta > 0$ such that $\log \beta > \Psi(\gamma)$. Denote the partition function for that model by Z_n^β . By Remark 3.9, the assumptions of Theorem 3.7 for $k = 1$ are satisfied, so the random variable $n^{-1/3} \left| \log \sigma_1(\tilde{A}_n^{-1}) - \log |Z_n^\beta| \right|$ converges to 0 in probability. It suffices to see that the mixed model for arbitrary β corresponds to the same polymer as for $\beta = 1$, except that all vertex weights are multiplied by β . It follows that $\log Z_n^\beta = \log Z_n + \log \beta \cdot 2n$, which in conjunction with (9) and $\sigma_N(\tilde{A}_n) = (\sigma_1(\tilde{A}_n^{-1}))^{-1}$ finishes the proof. $\square$

Theorem 1.1 follows from Theorem 3.10 by invoking Lemma 2.1 and noting that in the mixed model, the dual graph corresponds exactly to the lattice $\tilde{G}_n$ from Theorem 3.10.

As noted in Sect. 1, thanks to the result about fluctuations of the Beta polymer from [BC17], a similar result holds also for the model with Beta distributed weights with $\alpha = \beta = 1$, i.e. the uniform distribution, and when the lattice is rectangular of size $m \times n$, with $\frac{m}{n} \rightarrow \kappa \neq 1$. By taking $c > e$ we ensure that $\mathbb{E} \log |X_{e_i}| > 0$. By invoking Theorem 2.15 of [BC17], the same proof as in Theorem 3.10 gives, for appropriate choice of constants, an analogous theorem for the Beta model. This is also conjectured for $\alpha \neq \beta$, but not known rigorously.

We end with a brief discussion of the i.i.d model. The i.i.d. model also corresponds to a polymer model on the lattice, where the vertex and edge weights are independent given as in Remark 2.7 and the weight of a path is the product of the weights of all edges and vertices it contains. By Proposition 3.2, one can consider a model with weights

only on the edges which in this case satisfies $\mathbb{E} \log |X'_e| = 0$. If one could establish Proposition 3.5 in the case $\mathbb{E} \log |X'_e| = 0$ and a result analogous to Theorem 2.1 of [KQ16] for such an i.i.d. polymer, these would imply that an analogue of Theorem 3.10, and therefore Theorem 1.1, holds also for the i.i.d. model.

Acknowledgements. We would like to thank Guillaume Barraquand for bringing to our attention the results about the Beta polymer model and for very useful remarks regarding stating our results for general parameters γ, β . We would like to thank Michal Kotowski and an anonymous referee for helpful comments on the manuscript. Balint Virag was supported by the Canada Research Chair program, the NSERC Discovery Accelerator grant, the MTA Momentum Random Spectra research group, and the ERC consolidator grant 648017 (Abert).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

A. Appendix: Combinatorial Results on Non-intersecting Paths

The goal of this section is to prove Lemma 3.3.

A naive approach to the proof would be as follows. Given sets S, T and a nonintersecting tuple π connecting S to T , we would like to connect each point in S to a point in S_0 , and likewise for T and T_0 , with some path so as to obtain a nonintersecting tuple connecting S_0 to T_0 . However, it is easy to give examples where this cannot be done, e.g. when one of the points in S is cut off from the origin by some paths in π . Therefore, a more careful approach is needed.

We start with an informal outline of the proof. Observe that if π contains intersecting paths, one can switch them to make them noncrossing with each other (see Fig. 3). One can then push the paths away from each other to remove the intersections (see Fig. 5). However, in the process we can lose an unbounded number of edges (e.g. consider two zigzag paths touching each other at each corner). The key idea to prevent this is to connect every missing vertex from S_0 to a vertex from S using a path τ which is simple, i.e. goes only right and then up. All modifications of paths will occur in a neighborhood of τ of fixed radius. Since τ is simple, every path in such a neighborhood can make only a bounded, independent of n , number of turns, which in turn guarantees that we always remove only a bounded number of edges.

The above approach is formalized as follows. Lemma A.4 states that every tuple of paths can be uncrossed. Lemma A.9 states that noncrossing paths can be pushed away to make them nonintersecting while losing only a bounded number of edges. For technical reasons this requires ensuring that paths intersect properly (Definition A.1), which is handled by Lemma A.8. Finally, Lemma A.10 combines the previous lemmas to inductively connect missing points from S_0 to points from $S \setminus S_0$. The lemma requires that there are no points in S too close to the boundary, which is ensured by Lemma A.11. The section ends with the proof of Lemma 3.3.

For a path τ we let $N(\tau)$ be the neighborhood of τ of radius $6k$, i.e. the set of all vertices within ℓ_1 distance at most $6k$ of some vertex on τ . For a sequence S of initial vertices let $\text{ord}(S) := \min\{i \mid (1, i) \notin S\}$ be the y coordinate of the lowest vertex $(1, i)$ not contained in S and likewise for a sequence T of terminal vertices let $\text{ord}(T) := \max\{i \mid (n, i) \notin T\}$.

Definition A.1. For two paths, π_1 from s to t and π_2 from s' to t' , such that $\pi_1 \cap \pi_2 \neq \emptyset$ we say that they **intersect properly** if $s', t' \notin \pi_1$ and $s, t \notin \pi_2$.

We define a partial order on the vertices of the lattice, letting $(x, y) > (x', y')$ if $y > y'$ and $x < x'$. Consider two paths π_1, π_2 . If there exist vertices $x \in \pi_1, y \in \pi_2$ such that $x > y$ and for all $x' \in \pi_1, y' \in \pi_2$ if x' and y' are comparable, then also $x' > y'$, we will write $\pi_1 \succ \pi_2$. We write $\pi_1 \succeq \pi_2$ if $\pi_1 \succ \pi_2$ or $\pi_1 = \pi_2$.

Definition A.2. A crossing intersection between π_1 and π_2 is a maximal connected subset of vertices $C \subseteq \pi_1 \cap \pi_2$ such that π_1 enters C from the left and exits from the right and π_2 enters C from the bottom and exits up. We say that π_1 and π_2 are **crossing** if $\pi_1 \cap \pi_2$ contains at least one crossing intersection. We call a tuple of paths noncrossing if none of the paths in the tuple are crossing.

Remark A.3. If π_1 and π_2 intersect properly and are not crossing, then $\pi_1 \succ \pi_2$ or vice versa.

Lemma A.4. Fix σ, S, T and consider a tuple of paths $\pi = (\pi_1, \dots, \pi_k)$, where π_i connects s_i to $t_{\sigma(i)}$. There exists a permutation σ' and a tuple $\pi' = (\pi'_1, \dots, \pi'_k)$, where π'_i connects s_i to $t_{\sigma'(i)}$, such that π' uses exactly the same multiset edges as π and paths in π' are noncrossing. Moreover, there is no pair $j < j'$ such that $\pi_{j'} \prec \pi_j$. We will say that the tuple π' is obtained by **uncrossing** π .

Proof. We will prove the following statement, from which the lemma follows. Consider any sequences S, T of vertices of length k and a multiset of edges P such that for every vertex not in $S \cup T$ the number of ingoing edges is equal to the number of outgoing edges and the total number of edges outgoing from S minus the number of edges ingoing into S is k . We claim that there exists a tuple of paths $\pi = (\pi_1, \dots, \pi_k)$ and a permutation σ such that π_j connects s_j to $t_{\sigma(j)}$, the tuple π uses exactly the edges from P and the paths π_j are noncrossing.

We proceed by induction with respect to k . For $k = 1$ the statement is obvious. Suppose we have already proved the lemma for $k - 1$. For $i = 2, \dots, 2n$ let L_i be the line $x + y = i$. In each strip between L_i and L_{i+1} we list the edges from P from the southeasternmost one to the northwesternmost one. We assign each edge a label equal to the number of edges preceding it (multiple copies of an edge are assigned consecutive numbers). Consider the northeasternmost vertex $s \in S$ among all vertices from $x \in S$ such that the diagonal line started at x and going southeast intersects no other vertices

Fig. 3. Switching paths to make them non-crossing

Fig. 4. Lemma A.8

used by the edges from P . There exists a path ρ starting at s and ending at some $t \in T$ which uses only edges labelled 0. Indeed, suppose this was not the case and follow the edges labelled 0 starting from s . Take the first vertex x such that there is no outgoing edge from x with label 0. This implies there is a vertex y southeast of x with an outgoing edge labelled 0. The vertex y must belong to S and is to the northeast of s , contradicting the choice of s .

After removing ρ , by the inductive hypothesis the remaining edges can be arranged into a noncrossing tuple of paths $\tilde{\pi} = (\tilde{\pi}_2, \dots, \tilde{\pi}_k)$ and the tuple $\tilde{\pi}$ can be reordered in such a way that there is no pair $j < j'$ such that $\tilde{\pi}_{j'} \prec \tilde{\pi}_j$. The tuple $\pi' = (\rho, \tilde{\pi}'_1, \dots, \tilde{\pi}'_{k-1})$ is also noncrossing since ρ used only edges labelled 0, so it cannot cross any of the paths $\tilde{\pi}'_i$. For the same reason the tuple π' will be ordered in the desired way. $\square$

Definition A.5. Let $\pi = (\pi_1, \dots, \pi_k)$ be a tuple of paths. Let s, t be the starting and ending vertices of the path π_j for some j . Let P be the set of all paths that start at s , end at t and use edges from the set $\cup_{j=1}^k \pi_j$. We define $\min(s, t; \pi)$ to be the smallest path, with respect to $\succ$, among all the paths in P .

Note that this is well defined since for any two paths $\rho, \rho' \in P$ there exists a path $\rho'' \in P$ such that $\rho \succeq \rho'', \rho' \succeq \rho''$. The path ρ'' can be obtained by uncrossing ρ, ρ' (Lemma A.4) and taking the southeasternmost of the resulting paths. Since every set of paths in P has a lower bound, there exists a unique smallest element.

Definition A.6. We say a path makes a turn at a vertex v if it enters v from the left and exits up or enters from below and exits right. We call a path τ **simple** if it makes at most one turn.

Remark A.7. Let τ be a simple path. If a path is contained inside the neighborhood $N(\tau)$, the number of turns it makes is bounded by some constant D depending only on k . This implies that the number of possible paths lying inside $N(\tau)$ is bounded by $D' \cdot n^{D''}$ for some constants D', D'' depending only on k .

In the following lemmas we fix some $\sigma, S^{(0)}, T^{(0)}$ and a tuple $\pi^{(0)} \in \Pi_{\sigma, S^{(0)}, T^{(0)}}$. We also fix a simple path τ that starts at a vertex $(1, m)$ for some m . For any tuple $\pi = (\pi_1, \dots, \pi_k)$, a constant C and any $0 \leq i \leq k$ we define the following conditions:

Fig. 5. Pushing paths away to make them non-intersecting

- (1) for all $j \geq k - i + 1$ the paths π_j do not intersect any paths in π
- (2) there are no j, j' such that $j > j'$ and $\pi_j < \pi_{j'}$
- (3) $|E(\pi^{(0)}) \setminus E(\pi)| \leq C$
- (4) all the vertices at which paths in π intersect and all the edges in $E(\pi) \Delta E(\pi^{(0)})$ lie within distance $5(i + 1)$ of τ ; also, the set $E(\pi) \Delta E(\pi^{(0)})$ is a union of at most C paths
- (5) for all $j \leq m$ the vertex $(1, j)$ lies on exactly one path from π and for all $j > m$ the vertex $(1, j)$ lies on at most one path from π

Lemma A.8. *Let $\pi \in \Pi_{\sigma, S, T}$ for some σ, S, T be a noncrossing tuple that satisfies Conditions (1)–(5) for some i and constant C . Then there exists a noncrossing tuple $\pi' \in \Pi_{\sigma', S', T'}$ for some σ', S', T' that also satisfies Conditions (1)–(5) for i and some constant C' depending only on C and k such that all the paths in π' intersect properly.*

Proof. Let $(s_j, t_{\sigma}(j))$ be the pair of starting and terminal vertices of the path π_j . Suppose there are some $j \neq j'$ such that $s_{j'}$ lies on π_j (the case of $t_{\sigma}(j')$ is dealt with analogously). Note that by Condition (5) $s_{j'} \neq (1, y)$ for any y . Let s' be defined as the first vertex on $\pi_{j'}$ which does not belong to any path other than $\pi_{j'}$. If s' exists, let $\pi'_{j'}$ be the subpath of $\pi_{j'}$ started at s' . We replace $s_{j'}$ with s' and $\pi_{j'}$ with $\pi'_{j'}$. If there is no s' , we replace both $s_{j'}$ and $t_{\sigma}(j')$ with any vertex in $N(\tau)$ southeast of $\pi_{j'}$ which does not belong to any path and take $\pi'_{j'}$ to be a path of length zero. Note that since $\pi_{j'}$ intersected some other path, by Condition (1) we must have $j' \geq i$, so to maintain Condition (2) we can reorder the paths π_j for $j \geq i$ and this does not violate Condition (1).

We claim that the resulting tuple of paths π' uses all but boundedly many edges of $\pi_{j'}$, independent of n . Let:

$$s_{j'} = x_0, x_1, \dots, x_l = s'$$

be the sequence of vertices on $\pi_{j'}$ between $s_{j'}$ and s' . Suppose that we lose an edge (x_n, x_{n+1}) used by $\pi_{j'}$. Then for some path $\rho \in \pi$ we have $x_n \in \pi_{j'} \cap \rho$ and $x_{n+1} \notin \rho$ (otherwise the edge is still contained in ρ). All the paths are noncrossing, so, apart possibly of the first edge of $\pi_{j'}$, either $\pi_{j'}$ or ρ makes a turn at x_n . By Condition (4) the intersection at x_n , hence also the turn, happens inside $N(\tau)$. This implies that the total number of edges lost is at most the number of turns made by any path inside $N(\tau)$, which by Remark A.7 is bounded by some constant D independent of n , plus one.

Since we have removed at most $D + 1$ edges, Condition (3) still holds with the constant C increased by $D + 1$. The remaining conditions are clearly unchanged and the resulting tuple is still noncrossing. We can repeat the above step and after at most $2k$ such steps all the paths will intersect properly. The Condition (3) will be satisfied with $C' = C + 2k(D + 1)$. $\square$

Lemma A.9. *Let $\pi \in \Pi_{\sigma,S,T}$ for some σ, S, T be a noncrossing tuple that satisfies Conditions (1)–(5) for some i and constant C and all intersecting paths in π intersect properly. Then there exists a tuple $\pi' \in \Pi_{\sigma,S,T}$ that satisfies Conditions (1)–(5) for $i + 1$ and some constant C' depending only on C and k .*

Proof. To pass from i to $i + 1$ we need to ensure Condition (1) for i . Suppose that π_{k-i} intersects some path π_j with $j < k - i$. For every such j , we do the following. Informally, we want to push the path π_j right and down at the intersections to remove the intersections, see Fig. 5. To make this notion precise, let $P\pi_{k-i}$ denote the path π_{k-i} with the edges touching the bottom or right boundary are removed and the remaining edges shifted one vertex down and one vertex to the right. We let $\pi'_j := \min(s_j, t_{\sigma(j)}; (\pi_j, P\pi_{k-i}))$, defined as in Definition A.5. The path π'_j starts and ends at the same vertices as π_j . We claim that:

- π'_j and π_{k-i} do not intersect
- $\pi'_j \cup \pi_{k-i}$ uses all but boundedly many edges of π_j , independent of n

The paths π_{k-i} and π_j intersect properly and are noncrossing, which by Remark A.3 and Condition (2) implies that $\pi_{k-i} \succ \pi_j$. Assume by contradiction that $\pi'_j \cap \pi_{k-i} \neq \emptyset$ and take $u = (x, y)$ to be the first vertex in $\pi'_j \cap \pi_{k-i}$. The paths go together from u to some vertex $v = (x', y')$, possibly $v = u$, at which they exit to different vertices. We claim that π_{k-i} enters u from the left and π'_j enters u from the bottom. If not, the path π'_j would contain a vertex w to the left of u and π_{k-i} a vertex w' below u , so $w > w'$. Since $\pi_{k-i} \succ \pi_j$, this would imply $(w, u) \notin \pi_j$, so the edge (w, u) must come from $P\pi_{k-i}$, which is clearly impossible. In the same fashion we prove that leaves v up and π'_j leaves v to the right. Let ρ be the subpath of π'_j connecting $(x, y - 1)$ to $(x' + 1, y')$. Let $P\rho$ be ρ shifted one vertex down and right, so clearly $P\rho \subseteq P\pi_{k-i}$. By considering the path π_j with the path connecting (x, y) to (x', y') replaced by $P\rho$, we reach a contradiction with the minimality of $\pi'_j := \min(s_j, t_{\sigma(j)}; (\pi_j, P\pi_{k-i}))$.

We now show that $\pi'_j \cup \pi_{k-i}$ use all but boundedly many edges of π_j . We first claim the only edges lost are the ones between some $u \in \pi_j \setminus \pi_{k-i}$ and its neighbor $v \in \pi_j \cap \pi_{k-i}$. Suppose there is a horizontal edge $(u, v) \in \pi_j \setminus (\pi_{k-i} \cup \pi'_j)$ such that u, v do not belong to π_{k-i} . Let u', v' be the vertices u, v shifted one right and down. It is easy to see that none of the vertices u, v, u', v' are available for the path π'_j . Since $\pi_j \succ \pi'_j$, in order to cross the northwest-southeast diagonal line through u the path π'_j must pass through the set $\{x : x < u\}$ using a vertex belonging to $P\pi_{k-i}$. This would imply the existence of a vertex $w \in \pi_{k-i}$ such that $w < u$, contradicting $\pi_{k-i} \succ \pi_j$. The case of a vertical edge is handled in the same way.

Now, consider an edge lost between a vertex $u \in \pi_j \setminus \pi_{k-i}$ and its neighbor $v \in \pi_j \cap \pi_{k-i}$. Since the paths are noncrossing, one of the paths makes a turn at u or v . Thus, the number of edges lost is bounded by twice the number of turns made by any path inside $N(\tau)$, which is bounded by some constant D by Remark A.7. By the same reasoning, the set $E(\pi'_j) \setminus E(\pi_j)$ is a union of a set of paths of size bounded by $2D$, since

it consists of path segments starting or ending at vertices where one of the paths makes a turn.

We first check that by replacing π_j with π'_j we have not violated Condition (1). Suppose that for some $l \geq k - i + 1$ the paths π'_j and π_l intersect at some vertex x . We must have $x \notin \pi_j$, as otherwise $\pi_j \cap \pi_l \neq \emptyset$, contradicting Condition (1). We must then have $x \in P\pi_{k-i}$, which means there is a vertex $x' \in \pi_{k-i}$ such that $x' > x$. Since the paths are noncrossing, by Remark A.3 this would imply $\pi_{k-i} \succ \pi_l$, contradicting Condition (2).

We check that by replacing π_j with π'_j we have not violated Condition (2). Suppose to the contrary that $\pi'_j \succ \pi_l$ for some $l > j$. This means that for some $x \in \pi'_j$ and $y \in \pi_l$ we have $x > y$. Since π_j and π'_j have the same starting and terminal vertices and $\pi_j \geq \pi'_j$, the path π_j must cross the region $\{z : z \geq x\}$, in particular, it contains a vertex z such that $z \geq x > y$, proving that $\pi_j \succ \pi_l$ and contradicting Condition (2). The proof that Condition (2) holds for $l < j$ proceeds in the same fashion.

Note that after this procedure the path π_{k-i} still intersects properly and is noncrossing with other paths that it intersects. We can thus repeat the above procedure for all $j < k - i$ such that π_{k-i} and π_j intersect, obtaining a new tuple π' in which π_{k-i} does not intersect any paths. The Condition (1) holds now for $i + 1$. Condition (3) holds now for $i + 1$ since we have removed at most $2kD$ edges and likewise for Condition (4). Note that every path has been moved by distance at most 1 and all the edges added or removed were within distance $5i$ of τ . Since Condition (4) held for i , this implies that now it holds for $i + 1$. The resulting tuple π' thus satisfies all the required conditions for $i + 1$. $\square$

Lemma A.10. Fix σ, S, T , with either $S \neq S_0$ or $T \neq T_0$ up to reordering. Suppose that $S \setminus S_0$ does not contain any vertices closer than $k + 1$ to the bottom boundary and $T \setminus T_0$ does not contain any vertices closer than $k + 1$ to the top boundary. There exists a simple path τ with the following properties. For any non-intersecting tuple $\pi \in \Pi_{S,T,\sigma}^{n,i}$ there exist σ', S', T' and a non-intersecting tuple $\pi' \in \Pi_{S',T',\sigma'}^{n,i}$ such that $\text{ord}(S') > \text{ord}(S)$ (or, if $S = S_0$, $\text{ord}(T') < \text{ord}(T)$), $|E(\pi) \setminus E(\pi')|$ is bounded independently of n and $E(\pi) \Delta E(\pi')$ is a union of paths whose number is bounded independently of n and which all lie within $N(\tau)$.

Proof. We treat the case $S \neq S_0$, as the case of T is analogous.

Let $(1, m)$ be the vertex in $S_0 \setminus S$ with smallest m . Pick the path τ to be the one connecting $(1, m)$ to any vertex $s_j \in S \setminus S_0$ and such that it goes right and then up (possibly only up if $s_j = (1, t)$ for some t). This is possible since we assume that no vertices in S lie within the first k levels from the bottom. We modify the set S by removing s_j and inserting $(1, m)$. We replace the path π_j , connecting s_j to $t_{\sigma(j)}$, with $\tau \cdot \pi_j$, the concatenation of τ and π_j , which now starts at $(1, m)$. In this way, we obtain a new tuple that, in general, may contain intersecting paths.

We will construct the desired tuple π' using an inductive procedure. We let $\pi^{(0)} := \pi$. At step $0 \leq i \leq k$ we will have sets $S^{(i)}, T^{(i)}$ and a possibly intersecting tuple $\pi^{(i)} = (\pi_1^{(i)}, \dots, \pi_k^{(i)})$ connecting $S^{(i)}$ to $T^{(i)}$ and satisfying Conditions (1)–(5) for i . Applying Lemmas A.8 and A.9 will allow us to construct $\pi^{(i+1)}$ satisfying the conditions for $i + 1$.

To obtain $\pi^{(1)}$, we do the following. Condition (5) clearly holds for vertices $(1, l)$ with $l < m$ since the initial tuple π was nonintersecting. To ensure that it holds also for $(1, m)$, let j be such that $(1, m)$ is the starting vertex of the path $\pi_j^{(0)}$ and suppose that $(1, m)$ also lies on a path $\pi_{j'}^{(0)}$ with $j \neq j'$. Because π was nonintersecting, the path

Fig. 6. Switching paths $\pi_j^{(0)}$ and $\pi_{j'}^{(0)}$

$\pi_{j'}^{(0)}$ must start at $(1, m - 1)$. Recall that the path τ , which is an initial segment of $\pi_j^{(0)}$, connects $(1, m)$ to the vertex s_j .

If $s_j \neq (1, l)$ for any l , we switch paths so that the new path $\tilde{\pi}_{j'}^{(0)}$ starts with the path $(1, m - 1) \rightarrow (2, m - 1) \rightarrow (2, m)$ and then follows the edges of $\pi_j^{(0)}$, while $\tilde{\pi}_j^{(0)}$ starts at $(1, m)$ and follows the edges of $\pi_{j'}^{(0)}$. This is shown in Fig. 6.

If $s_j = (1, l)$ for some $l > m$, consider the endpoint t of the path $\pi_{j'}^{(0)}$. If $t = (1, l')$ for some $l' \geq m$, we must have $l' < l$ because in the tuple π the paths $\pi_j^{(0)}$ and the subpath of $\pi_{j'}^{(0)}$ starting at s_j were nonintersecting. We replace the vertex t by $(1, m - 1)$ and replace $\pi_{j'}^{(0)}$ with a path of length zero. If $t \neq (1, l')$, let $(2, a)$ be the first vertex on $\pi_{j'}^{(0)}$ in the second column. We replace the segment of $\pi_{j'}^{(0)}$ connecting $(1, m - 1)$ to $(2, a)$ with a path that first goes right and then up to $(2, a)$.

We then apply Lemma A.4 to make the tuple noncrossing and ensure that there are no $j' > j$ such that $\pi_{j'}^{(0)} < \pi_j^{(0)}$. The procedure above ensures that in all the cases considered the Conditions (3) and (4) hold. In this way all the Conditions (1)–(5) for $i = 0$ are satisfied.

To pass from $\pi^{(i)}$ to $\pi^{(i+1)}$, we apply the following steps. We first apply Lemma A.4 to obtain a tuple $\pi^{(i)'}$ with all the paths non-crossing. Then, we apply Lemma A.8 to obtain a tuple $\pi^{(i)''}$ in which all the paths intersect properly. We can then apply Lemma A.9, obtaining a tuple $\pi^{(i)'''}$ which satisfies Conditions (1)–(5) for $i + 1$. We can then put $\pi^{(i+1)} := \pi^{(i)'''}$.

After performing the inductive procedure, we obtain a tuple $\pi^{(k)}$ that satisfies all the required conditions. We put $S' := S^{(k)}$, $T' := T^{(k)}$, $\pi' := \pi^{(k)}$. By Condition (1) the tuple is nonintersecting. By Condition (3) $|E(\pi) \setminus E(\pi')|$ is bounded independently of n . By Condition (4) $E(\pi) \Delta E(\pi')$ is a union of paths whose number is bounded independently of n and which all lie within $N(\tau)$. Finally, since the vertices $(1, l)$ for $l \leq m$ were never removed, we have $\text{ord}(S') > \text{ord}(S)$. $\square$

Lemma A.11. *For any σ, S, T and a tuple $\pi \in \Pi_{S, T, \sigma}^{n, i}$ there exist σ', S', T' and a tuple $\pi' \in \Pi_{S', T', \sigma'}^{n, i}$ such that $S' \setminus S_0$ does not contain any vertices closer than $k + 1$ to the bottom boundary, $T' \setminus T_0$ does not contain any vertices closer than $k + 1$ to the top boundary, $|E(\pi) \setminus E(\pi')|$ is bounded independently of n and $E(\pi) \Delta E(\pi')$ is a union of*

$\mathcal{P}$ to be the set of all possible paths contained in $\cup_{i=1}^m N(\tau_i) \cup N_b$. By Remark A.7, the set $\mathcal{P}$ has size at most $D' \cdot n^{D''}$ for some constants D', D'' depending only on k . The proof is finished by taking $C := \max\{D, D'', 2kC'\}$. $\square$

References

- [BC17] Barraquand, G., Corwin, I.: Random-walk in beta-distributed random environment. *Probab. Theory Relat. Fields* **167**(3), 1057–1116 (2017)
- [BCR13] Borodin, A., Corwin, I., Remenik, D.: Log-gamma polymer free energy fluctuations via a fredholm determinant identity. *Commun. Math. Phys.* **324**(1), 215–232 (2013)
- [COSZ14] Corwin, I., O’Connell, N., Seppäläinen, T., Zygouras, N.: Tropical combinatorics and whittaker functions. *Duke Math. J.* **163**(3), 513–563 (2014)
- [Dur10] Durrett, R.: *Probability: Theory and Examples*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge (2010)
- [GV85] Gessel I, Viennot, G (1985) Binomial determinants, paths, and hook length formulae. *Advances in Mathematics*. 58: 300–321
- [Hog06] Hogben, L.: *Handbook of Linear Algebra*. Discrete Mathematics and Its Applications, 1st edn. Chapman & Hall, Boca Raton (2006)
- [KQ16] Krishnan, A., Quastel, J.: Tracy–widom fluctuations for perturbations of the log-gamma polymer in intermediate disorder. *Ann. Appl. Probab.* **28**, 3736 (2016)
- [LD16] De Luca, A., Le Doussal, P.: Mutually avoiding paths in random media and largests eigenvalues of random matrices. *Phys. Rev. E* **95**, 030103 (2016)
- [PMSO09] Prior, J., Somoza, A.M., Ortuño, M.: Conductance distribution in two-dimensional localized systems with and without magnetic fields. *Eur. Phys. J. B* **70**, 513–521 (2009)
- [QR14] Quastel, J., Remenik, D.: Airy processes and variational problems. In: Ramírez, A.F., Arous, G.B., Ferrari, P.A., Newman, C.M., Sidoravicius, V., Vares, M.E. (eds.) *Topics in Percolative and Disordered Systems*, pp. 121–171. Springer, New York (2014)
- [Sep12] Seppäläinen, T.: Scaling for a one-dimensional directed polymer with boundary conditions. *Ann. Probab.* **40**(1), 19–73 (2012)
- [SLDOn15] Somoza, A.M., Le Doussal, P., Ortuño, M.: Unbinding transition in semi-infinite two-dimensional localized systems. *Phys. Rev. B* **91**, 155413 (2015)
- [SONP07] Somoza, A.M., Ortuño, M., Prior, J.: Universal distribution functions in two-dimensional localized systems. *Phys. Rev. Lett.* **99**, 116602 (2007)

Communicated by A. Borodin