

ON THE CYCLE STRUCTURE OF MALLOWS PERMUTATIONS¹

BY ALEXEY GLADKICH AND RON PELED

Tel Aviv University

We study the length of cycles of random permutations drawn from the Mallows distribution. Under this distribution, the probability of a permutation $\pi \in \mathbb{S}_n$ is proportional to $q^{\text{inv}(\pi)}$ where $q > 0$ and $\text{inv}(\pi)$ is the number of inversions in π .

We focus on the case that $q < 1$ and show that the expected length of the cycle containing a given point is of order $\min\{(1 - q)^{-2}, n\}$. This marks the existence of two asymptotic regimes: with high probability, when n tends to infinity with $(1 - q)^{-2} \ll n$ then all cycles have size $o(n)$ whereas when n tends to infinity with $(1 - q)^{-2} \gg n$ then macroscopic cycles, of size proportional to n , emerge. In the second regime, we prove that the distribution of normalized cycle lengths follows the Poisson–Dirichlet law, as in a uniformly random permutation. The results bear formal similarity with a conjectured localization transition for random band matrices.

Further results are presented for the variance of the cycle lengths, the expected diameter of cycles and the expected number of cycles. The proofs rely on the exact sampling algorithm for the Mallows distribution and make use of a special diagonal exposure process for the graph of the permutation.

CONTENTS

1. Introduction	1115
1.1. Sampling algorithm	1119
1.2. Relation with other models	1120
1.2.1. Permutons	1120
1.2.2. Band matrices	1120
1.2.3. Card shuffling	1121
1.2.4. Spatial random permutations	1122
1.3. Reader’s guide	1124
2. Notation and preliminaries	1124
3. The sampling algorithm and the arc chain	1125
3.1. Generating the graph of a Mallows permutation	1125
3.2. Diagonal exposure and the arc chain process	1127
3.3. The distribution of the arc chain	1130
3.4. The hitting time of zero	1135
3.5. Induced Mallows permutations and a stitching process	1138

Received January 2016; revised May 2017.

¹Supported by ISF Grants 1048/11 and 861/15 and IRG Grant SPTRF.

MSC2010 subject classifications. Primary 60C05, 05A05; secondary 60F05, 60K35, 82B23, 82B26, 60B99, 60B20.

Key words and phrases. Mallows permutations, cycle structure, Poisson–Dirichlet law, phase transition, macroscopic cycles, localization, delocalization, random band matrices.

The infinite case	1140
Proofs of Lemma 3.15 and Lemma 3.16	1143
4. Main theorems	1144
4.1. Expected number of cycles	1147
4.2. Expected cycle diameter	1148
Lower bound	1148
Upper bound	1149
4.3. Expected cycle length	1151
4.4. Variance of cycle length	1153
4.5. Poisson–Dirichlet law	1155
The Poisson–Dirichlet law in a space of multisets of reals	1155
Convergence criterion	1156
Proof of the Poisson–Dirichlet law	1157
Proof of part (ii) of Lemma 4.11	1160
Proof of part (i) of Lemma 4.11	1161
Limiting distribution of $\frac{1}{n} \mathcal{C}_{S_n} $	1163
5. Discussion and open questions	1164
The limiting distributions of the cycle length and cycle diameter	1164
Macroscopic cycles	1165
Microscopic cycles	1165
Intermediate regime	1166
Extensions of the parameter range	1166
Band models	1167
Higher dimensions and general graphs	1167
Acknowledgments	1168
References	1168

1. Introduction. The cycle structure of a random permutation picked uniformly from \mathbb{S}_n , the permutation group on n elements, is a classic topic in probability theory. Of the abundant literature on it, we mention two key facts: the distribution of the length of a cycle containing a given point is uniform on $\{1, \dots, n\}$. Moreover, the joint distribution of the lengths of the longest cycles in the permutation has an explicit limit; the sorted vector of cycle lengths, normalized by n , converges in distribution to the Poisson–Dirichlet distribution with parameter one.

In this work, we study the cycle structure of a random permutation distributed according to the Mallows distribution. The Mallows distribution is a nonuniform distribution on permutations which was introduced by Mallows in statistical ranking theory [20]. It has recently been the focus of several studies in varied contexts including mixing times of Markov chains [3, 9], statistical physics [27, 28], learning theory [8], q -exchangeability [13, 14] and the problem of the longest increasing subsequence [2, 6, 21]. Borodin, Diaconis and Fulman ([7], Section 5) considered a class of models of random permutations (denoted P_θ there) for which the Mallows distribution is the prime example. They noted that many of the “usual questions” of applied probability and enumerative combinatorics remain open for such models and asked “Picking a permutation randomly from $P_\theta(\cdot)$, what is the distribution of

the cycle structure, longest increasing subsequence, ...?" Our work answers the part of this question pertaining to the cycle structure of the Mallows distribution.

The Mallows distribution on \mathbb{S}_n is parameterized by a real number $q > 0$ and is denoted $\mu_{n,q}$. It is given by

$$(1) \quad \mu_{n,q}[\pi] := \frac{q^{\text{inv}(\pi)}}{Z_{n,q}},$$

where

$$\text{inv}(\pi) := |\{(s, t) | s < t \text{ and } \pi_s > \pi_t\}|$$

denotes the number of inversions in π , and $Z_{n,q}$ is a normalizing constant, given explicitly by the following formula [26], Corollary 1.3.13,

$$Z_{n,q} = \prod_{i=1}^n (1 + q + \dots + q^{i-1}) = \prod_{i=1}^n \frac{1 - q^i}{1 - q}.$$

The Mallows distribution with parameter $q = 1$ coincides with the uniform distribution on \mathbb{S}_n . In this paper we restrict attention to the case that $0 < q < 1$ (a brief discussion of the case $q > 1$ is given in Section 5). As is well known, $\text{inv}(\pi)$ equals the minimal number of *adjacent* transpositions required to bring π to the identity. Thus, when $0 < q < 1$, the Mallows distribution gives higher weight to permutations, which are closer to the identity in an underlying one-dimensional geometry.

We are mainly interested in the properties of the Mallows distribution for q close to 1, usually as a function of n , although our results apply in the full range of $0 < q < 1$. Figure 1 depicts samples of the Mallows distribution. One simple feature of a Mallows distribution is that it typically displaces elements by a small amount. This is quantified in the following statement: there exists an absolute constant $c > 0$ such that if $\pi \sim \mu_{n,q}$ then for all $0 < q < 1$ and $1 \leq s \leq n$,

$$(2) \quad c \cdot \min\left\{\frac{q}{1-q}, n-1\right\} \leq \mathbb{E}|\pi_s - s| \leq \min\left\{\frac{2q}{1-q}, n-1\right\};$$

see [6] for a proof and related concentration bounds or [8, 14] for similar statements.

Thus the expected displacements are of order $o(n)$ when $\frac{1}{1-q} \ll n$ and it is natural to ask whether this also results in shorter cycles. Our first result determines the expected length of cycles.

NOTATION. For two quantities x, y , which may depend on other parameters such as n or q , we write $x \approx y$ if there exist absolute constants $c, C > 0$ such that $cy \leq x \leq Cy$.

For a permutation $\pi \in \mathbb{S}_n$ and $1 \leq s \leq n$ we let $\mathcal{C}_s = \mathcal{C}_s(\pi)$ be the orbit of s in π , that is, the set of points in the cycle of π which contains s .




FIG. 1. Graphs of samples of the Mallows distribution $\mu_{n,q}$ with $n = 1000$, $q = 0.99$ (left) and $q = 0.995$ (right). The red lines are at vertical distance $\frac{2}{1-q}$ from the diagonal. They delimit a region containing most of the points of the permutation; see also (2).

THEOREM 1.1 (Expected cycle length). *Let $n \geq 1$, $0 < q < 1$ and $\pi \sim \mu_{n,q}$. Then*

$$\mathbb{E}|C_s| \approx \min \left\{ \frac{1}{(1-q)^2}, n \right\} \quad \text{for all } 1 \leq s \leq n.$$

Thus the expected length of the cycle containing a given point transitions from being $o(n)$ when $\frac{1}{1-q} \ll \sqrt{n}$ to being $\Omega(n)$ in the complementary regime. The same is true also for the *maximal* cycle length in the permutation; see Claim 4.14. We say this transition marks the emergence of *macroscopic cycles* in the permutation.

Theorem 1.1 identifies a similarity between the uniform distribution and the Mallows distribution in the regime that macroscopic cycles exist, namely, that the expected cycle lengths in both distributions are of order n . The two distributions are quite different in many other respects, for instance, when $\frac{1}{1-q} \ll n$ they are distinguished even by their typical displacements as measured by (2). Our next result shows that as far as the lengths of the *long* cycles are concerned, the similarities extend much further than what may initially be expected: the two distributions give rise to the same limit law (see also Figure 2).

THEOREM 1.2 (Poisson–Dirichlet law). *Suppose that the sequence (q_n) satisfies*

$$0 < q_n < 1 \quad \text{and} \quad (1 - q_n)^2 \cdot n \rightarrow 0.$$

Let $\pi \sim \mu_{n,q_n}$ and let $\ell_1 \geq \ell_2 \geq \dots$ be the sorted lengths of cycles in π . Then, as $n \rightarrow \infty$,

$$\frac{1}{n}(\ell_1, \ell_2, \dots) \text{ converges in distribution to the Poisson–Dirichlet law}$$

with parameter one.




FIG. 2. Distribution of the length of the cycle containing a uniform random point. Obtained empirically with 1,000,000 samples.

In addition, for any sequence (s_n) satisfying $1 \leq s_n \leq n$, as $n \rightarrow \infty$,

$$\frac{1}{n}|\mathcal{C}_{s_n}| \text{ converges in distribution to the uniform distribution on } [0, 1].$$

Our results provide further information on the cycle lengths in the regime in which there are no macroscopic cycles. We show that the cycle lengths are not concentrated in the sense that their standard deviation has the same order of magnitude as their expectation.

THEOREM 1.3 (Variance of cycle length). *Let $n \geq 1, 0 < q < 1$ and $\pi \sim \mu_{n,q}$. Then*

$$\text{Var } |\mathcal{C}_s| \approx \min \left\{ \frac{q}{(1-q)^4}, (n-1)^2 \right\} \quad \text{for all } 1 \leq s \leq n.$$

The argument showing that the standard deviation is at least as large as the expectation, when q is bounded away from 0, bears something of a general nature and may be applicable to other spatial permutation models such as the interchange model; see Section 1.2.4 and Section 5 for more details on these models.

Our next theorem considers the diameter of cycles, showing that the cycles are dense in their support in the sense that their lengths are comparable to their diameters on average.

THEOREM 1.4 (Expected cycle diameter). *Let $n \geq 1, 0 < q < 1$ and $\pi \sim \mu_{n,q}$. Then, for all $1 \leq s \leq n$,*

$$(3) \quad \mathbb{E}[\max(\mathcal{C}_s) - \min(\mathcal{C}_s)] \approx \min \left\{ \frac{q}{(1-q)^2}, n-1 \right\}$$

and, moreover,

$$(4) \quad \mathbb{E}[\max(C_s) - s] \approx \min\left\{\frac{q}{(1 - q)^2}, n - s\right\},$$

$$(5) \quad \mathbb{E}[s - \min(C_s)] \approx \min\left\{\frac{q}{(1 - q)^2}, s - 1\right\}.$$

Given our previous theorems, naively, one may expect that a typical random Mallows permutation π has about $n / \min\{(1 - q)^{-2}, n\}$ cycles, as the cycle containing a given point typically has length of order $\min\{(1 - q)^{-2}, n\}$. However, such reasoning is known to be false even for a uniformly random permutation, in which the cycle containing a given point typically has length of order n , yet there are $\log n$ cycles on average. This phenomenon reflects the fact that while most cycles are short, most points lie in long cycles. Our last theorem clarifies that this is also the case for random Mallows permutations and gives the order of magnitude of the number of cycles.

THEOREM 1.5 (Expected number of cycles). *Let $n \geq 1, 0 < q < 1$ and $\pi \sim \mu_{n,q}$. Then*

$$\mathbb{E}[\text{number of cycles in } \pi] \approx (1 - q) \cdot n + \log(n + 1).$$

1.1. Sampling algorithm. Our results are based on an exact sampling algorithm for the Mallows distribution, which goes back to the original work of Mallows [20]. The algorithm allows us to sample a permutation $\pi \sim \mu_{n,q}$ sequentially as follows: Given π_1, \dots, π_{s-1} , the distribution of π_s is distributed on the remaining $n - s + 1$ values in a geometric progression. Precisely, if the remaining values are $j_1 < j_2 < \dots < j_{n-s+1}$ then

$$(6) \quad \mathbb{P}[\pi_s = j_k | \pi_1, \dots, \pi_{s-1}] = \frac{1 - q}{1 - q^{n-s+1}} \cdot q^{k-1}.$$

It is simple to verify the validity of this formula by noting that, given π_1, \dots, π_{s-1} , the assignment $\pi_s = j_k$ creates exactly $k - 1$ inversions between π_s and π_{s+1}, \dots, π_n ; precisely, if $\pi_s = j_k$ then necessarily $|\{t | t > s, \pi_t < \pi_s\}| = k - 1$.

In our proofs, we develop more flexible versions of the above formula, allowing us to sample portions of the cycles of the permutation iteratively and control the evolution of these portions; see Section 3.2 and the beginning of Section 4.

There exist extensions of the Mallows distribution and the above formula (for $q < 1$) to infinite permutations; one-to-one and onto functions $\pi : \mathbb{N} \rightarrow \mathbb{N}$ or $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$. The extension to the case of \mathbb{N} is straightforward, one simply takes the limit $n \rightarrow \infty$ in (6) to obtain a geometric distribution; see Gnedin and Olshanski [13]. The extension to a two-sided infinite permutation, when the index set is \mathbb{Z} , is more complicated due to the fact that there is no natural initial position to start the sampling process from.

Generating methods for the two-sided infinite case were developed in [14]. In one of these methods, one samples two one-sided infinite Mallows permutations and uses a “stitching” mechanism to merge these into a two-sided infinite permutation. We also present a method for sampling a Mallows permutation “from an interior point”; see Section 3.5. The method is presented for finite n and may be used also for the two-sided infinite case via an approximation theorem from [14], Section 7.2. This method may serve as a bridge to transfer results from the finite n case to the two-sided infinite case.

1.2. *Relation with other models.* In this section, we briefly describe other models for which related results have been obtained or are conjectured.

1.2.1. *Permutons.* The regime of parameters in which $n \cdot (1 - q) \rightarrow \beta$ is also of special interest as in this case there is a limiting density to the empirical measure of the points in the graph of a Mallows permutation. Starr [27] obtained an explicit formula for the limiting density as a function of β . In modern terminology, the limiting density is called a *permuton*. Recently, Mukherjee [22] proved Poisson limit theorems for the lengths of *short* cycles for models converging to permutons, including the Mallows model as a special case. See also Kenyon, Král’, Radin and Winkler [17] for relations with permutons with fixed pattern densities.

1.2.2. *Band matrices.* In the study of random matrices, models of matrices with a band structure are of interest. We elaborate on one representative model: Let A be an $n \times n$ random matrix in which, for a given band width $0 < W \leq n$, the entries $A_{i,j}$, $|i - j| < W$, are independent and identically distributed standard Gaussian random variables and the other entries are set to zero. Define the symmetric band matrix H by

$$H := \frac{A + A^t}{\sqrt{2}}.$$

The main focus in these studies is on the eigenvalues and eigenvectors of H .

In one extreme case $W = 1$, meaning that the matrix H is diagonal, the eigenvectors are the standard basis vectors. The other extreme case, when $W = n$, results in the Gaussian Orthogonal Ensemble (GOE) distribution (up to scaling). In this case, the distribution of H is invariant under conjugation by orthogonal matrices, implying that the eigenvectors of H form a uniformly distributed orthonormal basis.

It is conjectured that random matrices of this kind undergo a localization/delocalization transition as the band width W increases beyond the threshold \sqrt{n} . More precisely, one expects that when $W \ll \sqrt{n}$, the eigenvectors are localized in the sense that most of their ℓ^2 mass lies on a set whose size is $o(n)$ (possibly even in an interval of such size), whereas if $W \gg \sqrt{n}$ the eigenvectors have their ℓ^2 mass approximately uniformly spread. Furthermore, in the second regime, it

is expected that the local eigenvalue statistics have the same limit as in the GOE case as n tends to infinity. Informally, we may say that the local eigenvalue statistics should have the mean-field limit in the delocalized regime. See the survey of Spencer [25] and references within for more on these topics.

Our results prove an analogous transition for the Mallows distribution. One may consider the permutation matrix H_π associated with a random permutation $\pi \sim \mu_{n,q}$. By (2), this matrix has an approximate band structure in the sense that few of its nonzero entries $(H_\pi)_{s,\pi_s}$ have $|\pi_s - s|$ greater than a constant multiple of the band width $W = \min\{\frac{1}{1-q}, n\}$ (in fact, the probability that $|\pi_s - s| \geq tW$ decays exponentially in t ; see [6], Theorem 1.1 and Figure 1). Such a matrix is orthogonal, having its eigenvalues on the unit circle. The eigenvalues and eigenvectors of H_π are determined by the cycle structure of π : associated with each cycle of length ℓ , one has the ℓ eigenvalues $\exp(\frac{2\pi i j}{\ell})$, $0 \leq j \leq \ell - 1$, and correspondingly ℓ eigenvectors, supported on the coordinates of the cycle and giving equal mass to all points of it. Thus a localization/delocalization transition corresponds to the emergence of cycles whose length is of order n . Theorem 1.1 shows that such a transition occurs as the band width increases beyond \sqrt{n} , paralleling the conjecture for random band matrices. Moreover, Theorem 1.2 shows that in the delocalized regime, the statistics of long cycles approach the Poisson–Dirichlet distribution, the limiting statistics for uniform random permutations, in analogy with the above prediction for the local eigenvalue statistics.

The reader is also referred to the survey of Olshanski [23] for other analogies between random permutations and random matrices, discussing, in particular, analogies between random permutations distributed according to the Ewens measure (see also Section 5) and deformations of Dyson’s circular ensemble of random matrices.


1.2.3. Card shuffling. There are many natural dynamics on permutations for which the uniform distribution is stationary. Diaconis and Shahshahani [10] consider the following natural card shuffling scheme: Start with a deck of n cards. At each step, choose two cards uniformly and independently and exchange their positions in the deck. How many steps does one need to perform in order for the deck to become almost perfectly shuffled? In a beautiful application of representation theory to the study of Markov chains, it is proved in [10] that the state of the deck after $\frac{1}{2}n \log n + cn$ such steps is close to uniform (in the total variation distance) when c is a large positive constant, and is far from uniform when c is a large negative constant. The latter bound follows from the analysis of the coupon collector problem: when c is a large negative constant there will be many cards in the deck which have not moved from their initial position, creating a permutation with many fixed points. Thus the result of [10] may be interpreted as saying that the number of short cycles is the main obstacle for a permutation to become approximately uniform in this card shuffling scheme.

Schramm [24] considered the above card shuffling scheme further, investigating the state of the deck after tn steps are performed. The analysis in [24] proceeds by drawing an associated graph on the vertex set $\{1, \dots, n\}$, in which an edge is put between i and j if the cards at positions i and j in the deck have been exchanged. This associated graph is distributed as an Erdős–Rényi random graph, allowing one to deduce from the standard literature that when $t \leq \frac{1}{2}$, all cycles in the random permutation have size $o(n)$. Schramm’s work focuses on the case that $t > \frac{1}{2}$ and proves that macroscopic cycles emerge in this regime (see also Berestycki [4] for a later simpler argument). Moreover, confirming a conjecture of Aldous, it is proved that the limiting joint distribution of these macroscopic cycles obeys the same Poisson–Dirichlet law observed for uniform permutations. Thus, although it takes about $\frac{1}{2}n \log n$ steps for the full permutation to become approximately uniform, it takes far fewer steps for macroscopic cycles to start emerging and the joint distribution of these macroscopic cycles converges very quickly to the limiting joint distribution. A similar fact is true for the Mallows model by our results: when q increases beyond the threshold $1 - \frac{1}{\sqrt{n}}$, although the Mallows permutation is still far from uniform [distinguished by its displacements, say, as in (2)], macroscopic cycles begin to emerge and their joint distribution converges to the Poisson–Dirichlet law.

In this context, we mention that the Mallows permutation also arises via a shuffling algorithm. As proved by Benjamini, Berger, Hoffman and Mossel [3], it arises as the stationary distribution of a biased card-shuffling algorithm. In this algorithm, one starts with a deck of cards numbered $\{1, \dots, n\}$ and at each iteration picks uniformly a pair of adjacent cards in the deck. One flips a coin with probability $p = \frac{1}{1+q}$ for heads and rearranges the two cards according to the coin result, in increasing order if heads and in decreasing order if tails. The iterations are done independently of one another.

1.2.4. *Spatial random permutations.* A *spatial random permutation* is a random permutation which is biased towards the identity in some underlying geometry. This broad idea covers many models, among them the Mallows distribution which is biased towards the identity in a one-dimensional geometry. In this section, we briefly describe two other models in this class for which related results have been proved.

Let $G = (V, E)$ be a finite or infinite bounded-degree graph. The *interchange process* (also called the *stirring process* in some of the literature) gives a dynamics on permutations in \mathbb{S}_V , one-to-one and onto functions $\pi : V \rightarrow V$, which is associated to the structure of the graph. Each edge of the graph is endowed with an independent Poisson process of rate 1. An edge is said to *ring* at time t if an event of its Poisson process occurs at that time. Starting from the identity permutation $\pi^0 \in \mathbb{S}_V$, the interchange process, introduced by Tóth [30], is the permutation-valued stochastic process (π^t) obtained by performing a transposition along each edge at each time that it rings.



(a) Graph of a sample of the interchange process in a one-dimensional geometry with $n = 1000$ and $t = 10000$. (b) Graph of a sample of the Mallows distribution $\mu_{n,q}$ with $n = 1000$ and $q = 0.99$.

FIG. 3. Comparison between the Mallows distribution and the interchange model.

The interchange process on the complete graph coincides with a continuous time version of the Diaconis–Shashahani card shuffling algorithm discussed in the Section 1.2.3. Special attention has been given to the case that the graph $G = \mathbb{Z}^d$, where the interchange process is related to the magnetization of the quantum Heisenberg ferromagnet [30]. In particular, the following conjecture of Bálint Tóth has attracted significant attention but remains unresolved: When $d = 2$, for any $t > 0$, all cycles of π^t are finite almost surely. In contrast, when $d \geq 3$ and t is sufficiently large, π^t has an infinite cycle almost surely.

Besides the case of the complete graph, results on the existence of long cycles in the interchange process are currently available only for trees, by Angel [1] and Hammond [15, 16], and for the hypercube graph, by Kotecký, Miłoś and Ueltschi [18].

Recently, a quantitative study of the interchange process in a one-dimensional geometry, $V = \{1, \dots, n\}$ with i adjacent to $i + 1$, was performed by Kozma and Sidoravicius. Here, as each $(\pi_s^t)_{t>0}$ is a simple random walk, the typical displacement $|\pi_s^t - s|$ is of order $\min\{\sqrt{t}, n\}$. Thus the graph of π^t has a band structure similar to the graph of a Mallows permutation (see Figure 3), and the two models seem graphically similar when one takes

$$(7) \quad q = 1 - \frac{1}{1 + \sqrt{t}}.$$

In a work in preparation, Kozma and Sidoravicius prove that the expected length of the cycle containing a given point in π^t has order $\min\{t + 1, n\}$. This result, whose mathematical details were completed before our work began, is analogous to our Theorem 1.1 when making the assignment (7).

A second model of spatial random permutations, related to the Feynman–Kac representation of the ideal Bose gas in quantum statistical mechanics, has also

received significant attention; see [5] and references within. In this model, one samples a random collection of points (x_1, \dots, x_n) in a finite box $\Lambda \subset \mathbb{R}^d$ and a random permutation π on these points. The distribution is such that permutations with large displacements $\pi(x_i) - x_i$ have lower density. In the physical context, the emergence of macroscopic cycles in the model is related to the phenomenon of Bose–Einstein condensation. In recent work, Betz and Ueltschi [5] (see also Sütő [29]) have shown that the following phase transition takes place in the model when $d \geq 3$: Define the density of points per unit area $\rho = \frac{n}{|\Lambda|}$. There exists a critical density ρ_c such that, with probability tending to 1 as n tends to infinity, if the density is fixed to a value $\rho < \rho_c$ then all cycles have length $o(n)$, whereas if it is fixed to a value $\rho > \rho_c$ then macroscopic cycles, of size proportional to n , emerge. Moreover, in the second regime, the distribution of suitably normalized cycle lengths converges in distribution to the Poisson–Dirichlet law.

1.3. Reader’s guide. Section 2 introduces notation and preliminary facts used throughout the paper.

Section 3 develops flexible sampling methods for the Mallows distribution and studies closely related random processes: In Section 3.1, we develop tools for sampling the graph of a random Mallows permutation sequentially. These are used in Section 3.2 to introduce a “diagonal” exposure procedure for the graph. There, we also define the “arc chain” of a permutation, which tracks the number of open arcs (incomplete cycles) throughout the diagonal exposure process, and analyze its basic properties. Concentration bounds for the arc chain process are developed in Section 3.3 and used significantly in later proofs. Section 3.4 provides bounds on the time it takes an arc chain to reach zero, of use in the regime where q is bounded away from 1. Analogous bounds on return times also appeared in the recent work [2] of Basu and Bhatnagar where a related Markov chain is introduced. Section 3.5 considers the distribution of rectangular subsets of the graph of a Mallows permutation. These provide the starting point for a method to sample a Mallows permutation “from a mid-point”, which is further extended to a sampling method for the two-sided infinite case, when $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$.

Our main theorems are proved in Section 4: Section 4.1 is dedicated to the proof of Theorem 1.5, regarding the number of cycles. In Section 4.2, we prove Theorem 1.4 on the diameter of cycles by providing deviation bounds for the distribution of the maximal and minimal element of the cycle containing a given point. In Section 4.3, we prove Theorem 1.1 regarding the length of cycles. Section 4.4 is dedicated to the proof of Theorem 1.3, which provides bounds on the variance of the cycle lengths. Theorem 1.2 on the Poisson–Dirichlet law is proved in Section 4.5.

We conclude in Section 5 with a discussion and a selection of open questions.

2. Notation and preliminaries.

- Throughout the rest of the paper, n is a positive integer whilst $q \in (0, 1)$ is a real parameter:

- For two quantities $x, y \geq 0$, which may depend on other parameters such as n or q , we write $x \lesssim y$ if there exists an absolute constant $c > 0$ such that $x \leq c \cdot y$. Note that $x \approx y$ is equivalent to $x \lesssim y$ and $y \lesssim x$.
- \mathbb{N} is the set of positive integers while $[n] := \{i \in \mathbb{N} \mid i \leq n\} = \{1, 2, \dots, n\}$.
- $\mathbb{1}_A$ and $\mathbb{1}\{A\}$ denote the indicator random variable of an event A .
- Throughout the paper, we denote by $\xi = \xi_q$ the following:

$$(8) \quad \xi := \min \left\{ i \in \mathbb{N} \mid q^i \leq \frac{1}{2} \right\} = \left\lceil \log_q \frac{1}{2} \right\rceil \approx \frac{1}{1-q}.$$

- In order to avoid cumbersome expressions we will use an abbreviated notation when referring to subsets of \mathbb{Z}^2 . We write, for instance,

$$\{x < a, y < b\} \quad \text{instead of} \quad \{(x, y) \in \mathbb{Z}^2 \mid x < a, y < b\}$$

and analogous expressions involving other subsets of \mathbb{Z}^2 .

- We introduce two useful symmetries of the Mallows distribution $\mu_{n,q}$, that is, bijections $\mathbb{S}_n \leftrightarrow \mathbb{S}_n$ that preserve $\mu_{n,q}$. The inverse symmetry is induced by the inversion map

$$(9) \quad \mu_{n,q}[\pi] = \mu_{n,q}[\pi^{-1}].$$

The reversal symmetry is defined via the reversal map $r : s \mapsto n + 1 - s$ by

$$(10) \quad \mu_{n,q}[\pi] = \mu_{n,q}[r \circ \pi \circ r],$$

where we note that $\pi_i = j$ if and only if $(r \circ \pi \circ r)(n + 1 - i) = n + 1 - j$. The fact that the two maps $\pi \mapsto \pi^{-1}$ and $\pi \mapsto r \circ \pi \circ r$ preserve the Mallows distribution follows simply by checking that they preserve the number of inversions. These two symmetries will prove useful as they also preserve the cycle structure. Specifically, if \mathcal{C} is a cycle of π , then \mathcal{C}^{-1} and $r \circ \mathcal{C} \circ r$ are cycles of π^{-1} and $r \circ \pi \circ r$, respectively.

3. The sampling algorithm and the arc chain. In this section, we present a sampling algorithm for the Mallows distribution which will be fundamental in our analysis. We further identify a Markov chain associated to this sampling algorithm, termed the arc chain and explore its basic properties.

3.1. *Generating the graph of a Mallows permutation.* In Section 1.1, a method is presented for sampling the values (π_s) of a Mallows permutation iteratively. Here, we explain a related method, which generates the graph of the permutation

$$\Gamma_\pi := \{(s, \pi_s) \mid s \geq 1\}$$

in an iterative manner, allowing to expose portions of the graph in various orders.

Although our focus is on finite permutations, for clarity, we start by discussing the case of infinite one-sided permutations $\pi : \mathbb{N} \rightarrow \mathbb{N}$ in which the construction is simpler. In this case, as explained in Section 1.1,

$$\begin{aligned} \mathbb{P}[\pi_s = t | \pi_1, \dots, \pi_{s-1}] \\ = (1 - q) \cdot q^{|\{t\} \setminus \{\pi_1, \dots, \pi_{s-1}\}| - 1} \quad \text{for } t \notin \{\pi_1, \dots, \pi_{s-1}\}. \end{aligned}$$

In other words, the value of π_s , conditioned on the values of π_1, \dots, π_{s-1} , has the geometric distribution with success probability $1 - q$ on the values in $\mathbb{N} \setminus \{\pi_1, \dots, \pi_{s-1}\}$. This gives rise to the following sampling method: starting with a two-dimensional infinite array $(a_{s,t})_{s,t \geq 1}$ of independent Bernoulli random variables, each satisfying

$$\mathbb{P}[a_{s,t} = 1] = 1 - \mathbb{P}[a_{s,t} = 0] = 1 - q,$$

we may generate the permutation π by setting

$$\pi_s := \min\{t \geq 1 | t \notin \{\pi_1, \dots, \pi_{s-1}\}, a_{s,t} = 1\}.$$

Examination of this formula shows that the rule for deciding whether the point (s, t) belongs to the graph Γ_π depends only on the value of the bit $a_{s,t}$ and the portions of the graph Γ_π , which lie strictly below (s, t) or strictly to the left of (s, t) ,

$$(11) \quad \Gamma_\pi \cap \{x < s, y = t\} \quad \text{and} \quad \Gamma_\pi \cap \{x = s, y < t\}.$$

Precisely, $(s, t) \in \Gamma_\pi$ if and only if $a_{s,t} = 1$ and the two sets in (11) are empty. This viewpoint allows for iterative generation of the graph Γ_π in many different manners. In the sequel, we shall focus on diagonal generation, in which we expose the portion of the graph intersecting the square $\{x < t, y < t\}$ for increasing values of t .

Our next lemma gives an analogous generating method for the graph of a finite Mallows permutation, $\pi \in \mathcal{S}_n$, showing that many of the essential features of the above construction are preserved.

LEMMA 3.1. *Let $\pi \sim \mu_{n,q}$, let $s, t \in [n]$ and set $U := \{x < s \text{ or } y < t\}$. Then*

$$(12) \quad \mathbb{P}[\pi_s = t | \Gamma_\pi \cap U] = \frac{1 - q}{1 - q^{|\Gamma_\pi \cap U^c|}} \cdot \mathbb{1} \left\{ \begin{array}{l} \Gamma_\pi \cap \{x < s, y = t\} = \emptyset \\ \Gamma_\pi \cap \{x = s, y < t\} = \emptyset \end{array} \right\}.$$

We point out that the right-hand side of (12) does not depend on the full information in $\Gamma_\pi \cap U$. Indeed, to evaluate the right-hand side one only needs to know whether the sets in (11) are empty and the size of the set $|\Gamma_\pi \cap U^c|$, which may be computed, for instance, via

$$(13) \quad |\Gamma_\pi \cap U^c| = n - s - t + 2 + |\Gamma_\pi \cap \{x < s, y < t\}|.$$

As one application, one may use the equality (12) iteratively to compute the probability distribution of the portion of the graph $\Gamma_\pi \cap \{x \geq i, y < j\}$ conditioned on the portion of the graph $\Gamma_\pi \cap \{x < i, y < j\}$. The equality (12) shows that this probability distribution remains the same if we condition additionally on $\Gamma_\pi \cap \{x < i, y \geq j\}$. Therefore, we obtain the following conditional independence statement: for each $i, j \in [n]$,

$$(14) \quad \begin{aligned} &\text{conditionally on } \Gamma_\pi \cap \{x < i, y < j\}, \\ &\Gamma_\pi \cap \{x \geq i, y < j\} \text{ and } \Gamma_\pi \cap \{x < i, y \geq j\} \text{ are independent.} \end{aligned}$$

PROOF OF LEMMA 3.1. Our proof relies upon the formula (6). In the notation used there,

$$(15) \quad \mathbb{P}[\pi_s = j_k | \pi_1, \dots, \pi_{s-1}, \pi_s \geq j_k] = \frac{1 - q}{1 - q^{n-s-k+2}}.$$

We first claim that

$$(16) \quad \mathbb{P}[\pi_s = t | \Gamma_\pi \cap U \cap \{x \leq s\}] = \frac{1 - q}{1 - q^{|\Gamma_\pi \cap U^c|}} \cdot \mathbb{1}_A,$$

where $A := \{\Gamma_\pi \cap \{x < s, y = t\} = \emptyset\} \cap \{\Gamma_\pi \cap \{x = s, y < t\} = \emptyset\}$. The equality (16) certainly holds on A^c as both sides are zero. Now set $k := |[t - 1] \setminus \{\pi_1, \dots, \pi_{s-1}\}|$ and note that, by (13), $|\Gamma_\pi \cap U^c| = n - s - k + 2$. Observe that, in the notation of (15), we have that $t = j_k$ on the event A , as the set $\{\pi_1, \dots, \pi_{s-1}\}$ misses exactly $k - 1$ elements out of $[t - 1]$. It then follows from formula (15) that

$$\mathbb{P}[\pi_s = t | \Gamma_\pi \cap U \cap \{x \leq s\}] = \frac{1 - q}{1 - q^{n-s-k+2}} = \frac{1 - q}{1 - q^{|\Gamma_\pi \cap U^c|}} \quad \text{on } A,$$

completing the proof of (16).

Next, we observe that the argument used above to derive (14) from (12) may also be used to derive (14) from (16). Applying (14) with $i = s + 1$ and $j = t$ shows that π_s is conditionally independent of $\Gamma_\pi \cap U \cap \{x > s\}$ conditioned on $\Gamma_\pi \cap U \cap \{x \leq s\}$. Thus the formula (12) is a consequence of (16). \square

3.2. *Diagonal exposure and the arc chain process.* The main lemma of the previous section, Lemma 3.1, provides a procedure for calculating the distribution of certain portions of the graph Γ_π , of a Mallows permutation π , given others. This gives rise to several iterative algorithms for exposing the full graph. The proofs of our main theorems rely on a particular method of exposing Γ_π which will turn out to be particularly convenient. The total portion of Γ_π that will be revealed by time t will consist of $\Gamma_\pi \cap \{x \leq t, y \leq t\}$. Equivalently, as we pass from time t to $t + 1$ we reveal

$$(17) \quad \begin{aligned} &\Gamma_\pi \cap \{x = t + 1, y \leq t\}, \quad \Gamma_\pi \cap \{x \leq t, y = t + 1\} \quad \text{and} \\ &\Gamma_\pi \cap \{x = t + 1, y = t + 1\}. \end{aligned}$$

Formally, we define a finite filtration consisting of the sigma-algebras

$$\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}_{n-1} \subseteq \mathcal{F}_n$$

defined by

$$(18) \quad \mathcal{F}_t := \sigma(\Gamma_\pi \cap \{x \leq t, y \leq t\}).$$

Thus, \mathcal{F}_0 is the trivial σ -algebra and \mathcal{F}_n is the σ -algebra generated by π . We call this exposure procedure the *diagonal exposure* of π as the procedure exposes the graph in the diagonal direction. Corresponding to this filtration, we introduce the notation:

$$\mathbb{P}_t[A] := \mathbb{P}[A|\mathcal{F}_t] \quad \text{and} \quad \mathbb{E}_t[X] := \mathbb{E}[X|\mathcal{F}_t],$$

for an event A and a random variable X .

An important quantity to keep track of during the diagonal exposure process is the number of elements of π in the revealed portion of the graph Γ_π at each time t , that is, $|\Gamma_\pi \cap \{x \leq t, y \leq t\}|$. Our next definition introduces the counting process of an equivalent quantity, $|\Gamma_\pi \cap \{x \leq t, y > t\}|$, which will appear more frequently in our analysis. This quantity, as we elaborate upon in Section 4, counts the number of open ‘‘arcs’’, that is, portions of cycles that are yet to be closed, which are known using the information in \mathcal{F}_t .

DEFINITION 3.2 (Arc chain). The *arc chain* (κ_t) , $0 \leq t \leq n$, of a permutation $\pi \in \mathbb{S}_n$ is defined by

$$(19) \quad \kappa_t = \kappa_t(\pi) := |\{i \in [t] | \pi_i > t\}| = t - |\Gamma_\pi \cap \{x \leq t, y \leq t\}|,$$

that is, κ_t counts the number of π_1, \dots, π_t that are greater than t .

The arc chain is adapted to the filtration (\mathcal{F}_t) , that is, κ_t is determined by \mathcal{F}_t . One should note that π and π^{-1} share the same arc chain, that is,

$$(20) \quad \kappa_t = |\{t < i \leq n | \pi_i \leq t\}|.$$

The next proposition formalizes the fact that (κ_t) is a time-inhomogeneous Markov chain.

PROPOSITION 3.3. Let $\pi \sim \mu_{n,q}$. The arc chain κ of π is a time-inhomogeneous Markov chain, with respect to the filtration (\mathcal{F}_t) , satisfying $\kappa_0 = 0$ and $|\kappa_{t+1} - \kappa_t| \leq 1$, with transition probabilities given by

$$(21) \quad \begin{aligned} \mathbb{P}_t[\kappa_{t+1} = \kappa_t - 1] &= \left(\frac{1 - q^{\kappa_t}}{1 - q^{n-t}} \right)^2, \\ \mathbb{P}_t[\kappa_{t+1} = \kappa_t] &= \frac{q^{\kappa_t} - q^{n-t}}{1 - q^{n-t}} \cdot \frac{2 - q^{\kappa_t} - q^{\kappa_t+1}}{1 - q^{n-t}}, \\ \mathbb{P}_t[\kappa_{t+1} = \kappa_t + 1] &= \frac{q^{\kappa_t} - q^{n-t}}{1 - q^{n-t}} \cdot \frac{q^{\kappa_t+1} - q^{n-t}}{1 - q^{n-t}}. \end{aligned}$$

As an illustration of the usefulness of the arc chain, we note that the probability that π has a fixed point at position $t + 1$, given the information in \mathcal{F}_t , has a simple expression in terms of κ_t .

LEMMA 3.4. *Let $\pi \sim \mu_{n,q}$ and κ be its arc chain. Then*

$$\mathbb{P}_t[\pi_{t+1} = t + 1] = \frac{q^{\kappa_t} - q^{\kappa_t+1}}{1 - q^{n-t}} \cdot \frac{q^{\kappa_t} - q^{n-t}}{1 - q^{n-t}}, \quad 0 \leq t < n.$$

We prove the proposition and lemma together.

PROOF OF PROPOSITION 3.3 AND LEMMA 3.4. The newly revealed portions of Γ_π at time $t + 1$ were described in (17). Denote the first two of these portions by

$$X := \Gamma_\pi \cap \{x = t + 1, y \leq t\} \quad \text{and} \quad Y := \Gamma_\pi \cap \{x \leq t, y = t + 1\}.$$

We claim that

$$(22) \quad \mathbb{P}_t[X \neq \emptyset] = \mathbb{P}_t[Y \neq \emptyset] = \frac{1 - q^{\kappa_t}}{1 - q^{n-t}}.$$

It is convenient to derive this directly from (6). Write $j_1 < j_2 < \dots < j_{n-t}$ for the values in $[n] \setminus \{\pi_1, \dots, \pi_t\}$ and observe that $j_k \leq t$ if and only if $k \leq \kappa_t$, see (20). Therefore, it follows from (6) that

$$\mathbb{P}_t[X \neq \emptyset] = \mathbb{E}_t[\mathbb{P}[X \neq \emptyset | \pi_1, \dots, \pi_t]] = \mathbb{E}_t \left[\sum_{k=1}^{\kappa_t} \frac{1 - q}{1 - q^{n-t}} \cdot q^{k-1} \right] = \frac{1 - q^{\kappa_t}}{1 - q^{n-t}}.$$

The equality $\mathbb{P}_t[X \neq \emptyset] = \mathbb{P}_t[Y \neq \emptyset]$ follows from the symmetry (9), as π and π^{-1} share the same arc chain.

Now note that


$$\mathbb{P}_t[\pi_{t+1} = t + 1] = \mathbb{P}_t[X = Y = \emptyset] \cdot \mathbb{P}_t[\pi_{t+1} = t + 1 | X = Y = \emptyset].$$

The second factor can be computed directly from Lemma 3.1,

$$\mathbb{P}_t[\pi_{t+1} = t + 1 | X = Y = \emptyset] = \frac{1 - q}{1 - q^{n-t-\kappa_t}}.$$

In addition, observe that X and Y are conditionally independent given \mathcal{F}_t , as follows from (14). Thus the value of the first factor may be calculated from (22), which completes the proof of Lemma 3.4.

Observe that $\kappa_{t+1} - \kappa_t = 1 - |Z|$, where $Z := X \cup Y \cup (\Gamma_\pi \cap \{(t + 1, t + 1)\})$, as follows from the definition (19) of (κ_t) . As $|Z| \leq 2$, it follows that $|\kappa_{t+1} - \kappa_t| \leq 1$.



(a) The blue line depicts a single sample of an arc chain. The red dotted lines depict (linearly interpolated) percentiles of the distribution of the arc chain: 1st, 50th and 99th.
 (b) The red line is the graph of the stationary distribution of an $(\infty, 0.99)$ -arc chain. The blue dotted line is the graph of a normal distribution.

FIG. 4. Distributions associated with the arc chain of a Mallows permutation with parameters $n = 1000$ and $q = 0.99$.

The equations in (21) can be verified by computations similar to the ones used to prove Lemma 3.4, as

$$\begin{aligned} \mathbb{P}_t[\kappa_{t+1} = \kappa_t - 1] &= \mathbb{P}_t[|Z| = 2] = \mathbb{P}_t[X \neq \emptyset \text{ and } Y \neq \emptyset], \\ \mathbb{P}_t[\kappa_{t+1} = \kappa_t + 1] &= \mathbb{P}_t[|Z| = 0] \\ &= \mathbb{P}_t[X = Y = \emptyset] \cdot \mathbb{P}_t[\pi_{t+1} \neq t + 1 | X = Y = \emptyset], \end{aligned}$$

and the value of $\mathbb{P}_t[\kappa_{t+1} = \kappa_t]$ is derived from these using that $|\kappa_{t+1} - \kappa_t| \leq 1$. \square

3.3. *The distribution of the arc chain.* In this section, we study the distribution of the arc chain of a Mallows permutation at a fixed time t . Our main result, Theorem 3.6 below, states that the value of the chain is unlikely to be much larger than the value of ξ given in (8). Figure 4(a) depicts the percentiles of the distribution of the arc chain for certain values of n and q and all times t . These suggest that the typical values of the arc chain are close to ξ when t is bounded away from 1 and n . For such times, we establish in Proposition 3.8 below a formula for the limiting distribution of the arc chain when n tends to infinity with q fixed.

It is convenient to refer to the arc chain as an abstract Markov chain, without reference to an underlying Mallows permutation, as facilitated by the following definition.

DEFINITION 3.5. A random sequence (κ_t) , $0 \leq t \leq n$, is an (n, q) -arc chain, denoted $\kappa \sim \mathcal{AC}_{n,q}$, if κ is a time-inhomogeneous Markov chain with transition probabilities as in Proposition 3.3 and some initial distribution κ_0 supported in $\{0, \dots, n\}$.

We point out that the formulas in Proposition 3.3 constitute valid transition probabilities (i.e., they are nonnegative and sum to 1) when $0 \leq \kappa_t \leq n - t$. Using induction on t , one checks that an (n, q) -arc chain satisfies this condition for all t almost surely.

THEOREM 3.6. *Let $\kappa \sim \mathcal{AC}_{n,q}$ with $\kappa_0 = 0$. Then for all $0 \leq t \leq n$ and $d \in \mathbb{N}$ we have*

$$(23) \quad \mathbb{P}[\kappa_t > \xi + d] \leq \frac{q^{d^2+d}}{1 - q^{2d}}.$$

The idea of proof involves the definition of a limiting time-homogeneous Markov chain, corresponding formally to the case that $n = \infty$, and bounding the distribution of the (n, q) -arc chain by the stationary distribution of the limiting chain. Some of the tools that we develop here will be used later in the paper as well.

We recall that a time-homogeneous Markov chain (α_t) which takes values in the nonnegative integers and satisfies $|\alpha_t - \alpha_{t+1}| \leq 1$ is called a *birth-and-death chain*.

DEFINITION 3.7. A random sequence $(\hat{\kappa}_t)_{t \geq 0}$ is an (∞, q) -arc chain, denoted $\hat{\kappa} \sim \mathcal{AC}_{\infty,q}$, if κ is a birth-and-death chain with transition probabilities given by

$$(24) \quad \begin{aligned} \mathbb{P}[\hat{\kappa}_{t+1} = \hat{\kappa}_t - 1 | \hat{\kappa}_t] &= (1 - q^{\hat{\kappa}_t})^2, \\ \mathbb{P}[\hat{\kappa}_{t+1} = \hat{\kappa}_t | \hat{\kappa}_t] &= 2q^{\hat{\kappa}_t} - q^{2\hat{\kappa}_t} - q^{2\hat{\kappa}_t+1}, \\ \mathbb{P}[\hat{\kappa}_{t+1} = \hat{\kappa}_t + 1 | \hat{\kappa}_t] &= q^{2\hat{\kappa}_t+1}. \end{aligned}$$

There is a formula for the stationary measure of a birth-and-death chain. If (α_t) is a birth-and-death chain taking values in $\{0, \dots, m\}$, where m may be finite or infinite, having positive transition probabilities between consecutive integers in $\{0, \dots, m\}$, then α has a stationary measure z defined by²

$$(25) \quad z_0 := 1 \quad \text{and} \quad z_s := \prod_{i=1}^s \frac{\mathbb{P}[\alpha_{t+1} = i | \alpha_t = i - 1]}{\mathbb{P}[\alpha_{t+1} = i - 1 | \alpha_t = i]} \quad \text{for } 1 \leq s \leq m.$$

This is straightforward to verify directly. It is also not difficult to check that the stationary measure is unique up to scaling though we shall not use this fact (see also [19], Section 2.5).

²Here and later in the paper, to avoid introducing extra notation, we denote the transition probability of the chain from $i - 1$ to i at time t by $\mathbb{P}[\alpha_{t+1} = i | \alpha_t = i - 1]$, even if $\mathbb{P}[\alpha_t = i - 1] = 0$, and use similar notation for other transition probabilities.

This fact allows us to find the stationary distribution of an (∞, q) -arc chain $\hat{\kappa}$. Put

$$(26) \quad \begin{aligned} u_i &:= \mathbb{P}[\hat{\kappa}_{t+1} = i | \hat{\kappa}_t = i - 1] = q^{2^{i-1}} > 0, & i \geq 1, \\ v_i &:= \mathbb{P}[\hat{\kappa}_{t+1} = i - 1 | \hat{\kappa}_t = i] = (1 - q^i)^2 > 0, & i \geq 1. \end{aligned}$$

Then the sequence (ν_s) , $s \geq 0$, defined by

$$(27) \quad \nu_s := \frac{\prod_{i=1}^s \frac{u_i}{v_i}}{\sum_{j \geq 0} \prod_{i=1}^j \frac{u_i}{v_i}}$$

defines a stationary distribution for $\hat{\kappa}$, where, as usual, an empty product is interpreted as 1. The denominator in (27) is finite since $u_i \rightarrow 0$ and $v_i \rightarrow 1$ as $i \rightarrow \infty$. See Figure 4(b) for a graph of ν .

We study further the relation between the distributions of the (n, q) -arc chain κ and the (∞, q) -arc chain $\hat{\kappa}$. Our next proposition shows that in a suitable limit, in which q is fixed, the distribution of κ_t converges to the stationary distribution of $\hat{\kappa}$. This proposition will be of use in Section 3.5.

PROPOSITION 3.8. *Let $\kappa \sim \mathcal{AC}_{n,q}$ with $\kappa_0 = 0$ and set $t = t_n$. If both $t \rightarrow \infty$ and $n - t \rightarrow \infty$, then the law of κ_t converges to the stationary distribution of $\mathcal{AC}_{\infty,q}$, as n tends to infinity with q fixed.*

Our main tool for proving the above theorem and proposition is a coupling in which the (∞, q) -arc chain bounds the (n, q) -arc chain at all times. We first introduce a general method for performing such couplings.

Let α be a Markov chain, possibly time-inhomogeneous, taking values in the nonnegative integers and satisfying $|\alpha_t - \alpha_{t+1}| \leq 1$. Let (U_t) be a sequence of independent random variables, each uniformly distributed on $[0, 1]$. One may couple the Markov chain α with the sequence U as follows. The initial distribution α_0 is taken independent of U . Then, for each $t \geq 0$, $\alpha_{t+1} := \alpha_t + F_t^\alpha(\alpha_t, U_t)$ where

$$F_t^\alpha(a, u) := \mathbb{1}_{\{u > 1 - \mathbb{P}[\alpha_{t+1} = a + 1 | \alpha_t = a]\}} - \mathbb{1}_{\{u \leq \mathbb{P}[\alpha_{t+1} = a - 1 | \alpha_t = a]\}}.$$

This can be understood as “ α_{t+1} is a monotone function of U_t for a given α_t ”. We say that a set of Markov chains of the above type is *monotonically coupled* if they are all coupled to the same sequence U via the above method.

PROPOSITION 3.9. *Let $q, \hat{q} \in (0, 1)$ and $n, \hat{n} \in \mathbb{N} \cup \{\infty\}$ satisfy $q \leq \hat{q}$ and $n \leq \hat{n}$. Let an (n, q) -arc chain κ be monotonically coupled with an (\hat{n}, \hat{q}) -arc chain $\hat{\kappa}$ that satisfies $\kappa_0 \leq \hat{\kappa}_0$ almost surely. Then, almost surely, $\kappa_t \leq \hat{\kappa}_t$ for all $0 \leq t \leq n$.*

DEFINITION 3.10 (Bounding chain). For an (n, q) -arc chain κ , a *bounding chain* is any (∞, q) -arc chain $\hat{\kappa}$ that is monotonically coupled with κ and satisfies $\kappa_0 \leq \hat{\kappa}_0$ almost surely. Proposition 3.9 implies that any bounding chain satisfies $\kappa_t \leq \hat{\kappa}_t$, almost surely, for all $0 \leq t \leq n$.

PROOF OF PROPOSITION 3.9. The proof relies on the facts that $|\kappa_{t+1} - \kappa_t| \leq 1$, $|\hat{\kappa}_{t+1} - \hat{\kappa}_t| \leq 1$ and the following three inequalities:

$$\begin{aligned}
 & \mathbb{P}[\kappa_{t+1} = k + 1 | \kappa_t = k] \\
 & \leq 1 - \mathbb{P}[\hat{\kappa}_{t+1} = k | \hat{\kappa}_t = k + 1], \quad 0 \leq k \leq \min\{n - t, \hat{n} - t - 1\}, \\
 (28) \quad & \mathbb{P}[\kappa_{t+1} = k - 1 | \kappa_t = k] \geq \mathbb{P}[\hat{\kappa}_{t+1} = k - 1 | \hat{\kappa}_t = k], \quad 0 \leq k \leq n - t, \\
 & \mathbb{P}[\kappa_{t+1} = k + 1 | \kappa_t = k] \leq \mathbb{P}[\hat{\kappa}_{t+1} = k + 1 | \hat{\kappa}_t = k], \quad 0 \leq k \leq n - t.
 \end{aligned}$$

To prove these inequalities, observe that the transition probabilities of arc chains are given by

$$\begin{aligned}
 (29) \quad & \mathbb{P}[\kappa_{t+1} = k - 1 | \kappa_t = k] = \left(\frac{1 - q^k}{1 - q^{n-t}} \right)^2, \\
 & \mathbb{P}[\kappa_{t+1} = k + 1 | \kappa_t = k] = \left(1 - \frac{1 - q^k}{1 - q^{n-t}} \right) \cdot \left(1 - \frac{1 - q^{k+1}}{1 - q^{n-t}} \right).
 \end{aligned}$$

Thus the last two inequalities in (28) follow from the fact that $\frac{1 - q^a}{1 - q^b} = \frac{1 + q + \dots + q^{a-1}}{1 + q + \dots + q^{b-1}}$, with $1 \leq a \leq b$, decreases with both b and q . The first inequality in (28) follows from the third inequality there and the fact that $\mathbb{P}[\hat{\kappa}_{t+1} = k + 1 | \hat{\kappa}_t = k] \leq 1 - \mathbb{P}[\hat{\kappa}_{t+1} = k | \hat{\kappa}_t = k + 1]$. This last fact follows by substituting the formulas in (29), using that $1 - x^2 = (1 - x)(1 + x)$ and taking out the nonnegative common factor $1 - \frac{1 - q^{k+1}}{1 - q^{n-t}}$.

We proceed to prove the proposition. Suppose $t < n$ is such that $\kappa_t \leq \hat{\kappa}_t$ almost surely and let us show that $\kappa_{t+1} \leq \hat{\kappa}_{t+1}$, almost surely. Recall that $0 \leq \kappa_t \leq n - t$ and $0 \leq \hat{\kappa}_t \leq \hat{n} - t$ almost surely, and let us consider separately the following three cases:

- The inequality is clear if $\hat{\kappa}_t - \kappa_t \geq 2$.
- If $\kappa_t = \hat{\kappa}_t - 1$ then $\kappa_{t+1} \leq \hat{\kappa}_{t+1}$ follows from the first inequality in (28).
- Lastly, if $\kappa_t = \hat{\kappa}_t$ then $\kappa_{t+1} \leq \hat{\kappa}_{t+1}$ is a consequence of the second and third inequality in (28). \square

As a corollary of this proposition, we deduce that an (n, q) -arc chain κ with $\kappa_0 = 0$ satisfies

$$(30) \quad \mathbb{P}[\kappa_t \geq d] \leq \nu[d, \infty] = \sum_{i \geq d} \nu_i \quad \text{for all } t \leq n \text{ and } d \geq 0,$$

where ν is the stationary distribution of an (∞, q) -arc chain, as given by (27). This follows by letting $\hat{\kappa}$ be the bounding chain of κ having $\hat{\kappa}_0 \sim \nu$. Then (30) is a consequence of the facts that $\hat{\kappa}_t \sim \nu$ and $\kappa_t \leq \hat{\kappa}_t$.

We are now ready to prove Theorem 3.6 and Proposition 3.8.

PROOF OF THEOREM 3.6. By (30), it suffices to prove that the stationary distribution ν of an (∞, q) -arc chain satisfies

$$(31) \quad \nu[\xi + d + 1, \infty] \leq \frac{q^{d^2+d}}{1 - q^{2d}}, \quad d \geq 1.$$

Let u_i and v_i be as in (26) and set $w_i := u_i/v_i$. Formula (27) implies that

$$(32) \quad \nu_{\xi+s+1} = \nu_{\xi+1} \cdot \prod_{j=1}^s w_{\xi+j+1} \leq \prod_{j=1}^s w_{\xi+j+1}, \quad s \geq 1.$$

Observe that $u_{\xi+1} \leq \frac{1}{4}$ and $v_{\xi+1} \geq \frac{1}{4}$, by the definition (8) of ξ , yielding $w_{\xi+1} \leq 1$. One may verify that $w_{i+1} \leq q^2 \cdot w_i$, which yields that $w_{\xi+j+1} \leq q^{2j} \cdot w_{\xi+1} \leq q^{2j}$. By substituting this in (32), we conclude that

$$\nu_{\xi+s+1} \leq q^{s^2+s}.$$

Summing this inequality over $s \geq d$ yields (31), by bounding the sum with a geometric progression with quotient q^{2d} . \square

PROOF OF PROPOSITION 3.8. Recall the *convergence theorem for finite-state Markov chains*: if a finite-state time-homogeneous Markov chain (a_t) is aperiodic and irreducible, then it has a stationary distribution and the distribution of a_t converges to this stationary distribution as $t \rightarrow \infty$.

Equation (30) states that ν dominates κ_t . We will construct distributions $\tilde{\nu}_h$ that are asymptotically dominated by κ_t as $t \rightarrow \infty$. Then, with the limit of κ_t sandwiched between $\tilde{\nu}_h$ and ν , we will show that $\tilde{\nu}_h$ approaches ν as $h \rightarrow \infty$. Let h be some fixed positive integer and assume without loss of generality that $n - t \geq h$. Let $\tilde{\kappa}$ be the birth-and-death chain having $\kappa_0 = 0$ and transition probabilities determined by

$$\begin{aligned} \mathbb{P}[\tilde{\kappa}_{t+1} = \tilde{\kappa}_t - 1 | \tilde{\kappa}_t] &= \left(\frac{1 - q^{\tilde{\kappa}_t}}{1 - q^h} \right)^2 \quad \text{and} \\ \mathbb{P}[\tilde{\kappa}_{t+1} = \tilde{\kappa}_t + 1 | \tilde{\kappa}_t] &= \frac{q^{\tilde{\kappa}_t} - q^h}{1 - q^h} \frac{q^{\tilde{\kappa}_t+1} - q^h}{1 - q^h}. \end{aligned}$$

Observe that $0 \leq \tilde{\kappa}_t \leq h$ for all t , almost surely. Let $\tilde{\kappa}$ be monotonically coupled with κ . It is not hard to check that the pair $(\tilde{\kappa}, \kappa)$ satisfies the analogous inequalities of (28) for $t \leq n - h$, which implies, by following the proof of Proposition 3.9, that $\tilde{\kappa}_t \leq \kappa_t$ for all $t \leq n - h$, almost surely.

By applying the convergence theorem for finite-state Markov chains, we obtain that $\tilde{\kappa}$ has stationary distribution $\tilde{\nu} = \tilde{\nu}_h$ and that $\tilde{\kappa}_t$ converges to $\tilde{\nu}$ in distribution. Since κ_t dominates $\tilde{\kappa}_t$ for all $t \leq n - h$, we obtain

$$\tilde{\nu}[d, \infty] \leq \liminf \mathbb{P}[\kappa_t \geq d] \leq \limsup \mathbb{P}[\kappa_t \geq d] \leq \nu[d, \infty] \quad \text{for } d \geq 0,$$

where the limits are taken for $n, t \rightarrow \infty$ with the restriction $n - t \geq h$. It remains to verify, using (25) and the fact that the ratios $\frac{\mathbb{P}[\tilde{\kappa}_{t+1}=i|\tilde{\kappa}_t=i-1]}{\mathbb{P}[\tilde{\kappa}_{t+1}=i-1|\tilde{\kappa}_t=i]}$ increase with h , that $\tilde{\nu}_s \rightarrow \nu_s$ as $h \rightarrow \infty$, for all s . Thus $\tilde{\nu}[d, \infty]$ converges to $\nu[d, \infty]$ as $h \rightarrow \infty$, completing the proof of the proposition. \square

3.4. *The hitting time of zero.* The times in which the arc chain is at zero can be thought of as cut points for the graph of the permutation in the sense that if $\kappa_t = 0$ then $\Gamma_\pi \subseteq \{x \leq t, y \leq t\} \cup \{x > t, y > t\}$. This leads one to consider the evolution of the arc chain as performing a sequence of excursions away from zero; a point of view which will be useful for us in the regime that q is bounded away from 1 since, as we now prove, the excursions tend to be relatively short in this regime.

The recent work of Basu and Bhatnagar [2] uses a similar viewpoint in their analysis of the longest monotone subsequences in a random Mallows permutation. There, a Markov chain related to our (∞, q) -arc chain is considered. While the two chains differ, they share the same visit times to zero and the work [2] contains an analysis of the distribution of the return times to zero, related to our discussion here.

The following theorem will be instrumental in the proofs of the upper bounds of Theorem 1.1, Theorem 1.3 and Theorem 1.4 in this regime.

THEOREM 3.11. *Let κ be an (n, q) -arc chain with $\kappa_0 = 0$ and let $0 \leq s \leq n$. For any $\varepsilon > 0$, there exists a constant $c_\varepsilon > 0$ such that*

$$T_s := \min\{t \geq s | \kappa_t = 0\} \quad \text{satisfies}$$

$$\mathbb{E}[(T_s - s)^2] \leq c_\varepsilon \cdot q \quad \text{for all } q \in (0, 1 - \varepsilon).$$

This theorem is a consequence of the following two statements.

PROPOSITION 3.12. *There exists a monotone nondecreasing function $f: (0, 1) \rightarrow [0, \infty)$ such that the following holds. Let $k \geq 0$ and let $\hat{\kappa}$ be an (∞, q) -arc chain with $\hat{\kappa}_0 = k$. Then*

$$T := \min\{t \geq 0 | \hat{\kappa}_t = 0\} \quad \text{satisfies} \quad \mathbb{E}[T^2] \leq f(q) \cdot k^2 \quad \text{for all } k \geq 0.$$

Recall the definition of ξ from (8).

LEMMA 3.13. *Let ν be the stationary distribution of the (∞, q) -arc chain. One has*

$$(33) \quad \nu_x \leq 2^{2\xi - x} \quad \text{for } q \in (0, 1),$$

$$(34) \quad \nu_x \leq \frac{4q}{2^x} \quad \text{for } x > 0 \text{ and } q \in \left(0, \frac{1}{4}\right).$$

PROOF OF THEOREM 3.11. Let $\hat{\kappa}$ be a bounding chain of κ with $\hat{\kappa}_0$ having the stationary distribution of the (∞, q) -arc chain. Define \hat{T}_s by

$$\hat{T}_s := \min\{t \geq s \mid \hat{\kappa}_t = 0\}.$$

Proposition 3.9 implies that $T_s \leq \hat{T}_s$. Hence it suffices to prove that there exists $c_\varepsilon > 0$ such that

$$\begin{aligned} \mathbb{E}[(\hat{T}_s - s)^2] &\leq c_\varepsilon && \text{for } q \in (0, 1 - \varepsilon), \text{ and} \\ \mathbb{E}[(\hat{T}_s - s)^2] &\lesssim q && \text{for } q \in \left(0, \frac{1}{4}\right). \end{aligned}$$

As $\hat{\kappa}$ is a time-homogeneous Markov chain, Proposition 3.12 implies that

$$\mathbb{E}[(\hat{T}_s - s)^2 \mid \hat{\kappa}_s] \leq f(q) \cdot \hat{\kappa}_s^2.$$

By taking expectations, we obtain

$$(35) \quad \mathbb{E}[(\hat{T}_s - s)^2] \leq f(q) \cdot \sum_{k>0} k^2 \cdot \nu_k.$$

Using the geometric bounds on ν_s provided in (33) and the fact that f is a monotone nondecreasing function, it follows that the right-hand side in (35) is uniformly bounded for $q \in (0, 1 - \varepsilon)$.

Now consider the case $q \leq \frac{1}{4}$. By applying (34) to (35), we conclude that

$$\mathbb{E}[(\hat{T}_s - s)^2] \leq f\left(\frac{1}{4}\right) \cdot \sum_{k>0} \nu_k \cdot k^2 \leq 4q \cdot f\left(\frac{1}{4}\right) \cdot \sum_{k>0} \frac{k^2}{2^k} \lesssim q. \quad \square$$

Now we need only prove Lemma 3.13 and Proposition 3.12.

PROOF OF LEMMA 3.13. Let u_i and v_i be as in (26). As u_i is monotone decreasing in i and v_i is monotone increasing in i , it follows that for $i > 2\xi$ one has

$$\frac{u_i}{v_i} \leq \frac{u_{2\xi+1}}{v_\xi} \leq q^\xi \cdot \frac{u_{\xi+1}}{v_\xi} \leq q^\xi \leq \frac{1}{2},$$

where we have used that $u_{\xi+1} \leq \frac{1}{4}$ and $v_\xi \geq \frac{1}{4}$. Thus, using (25), we obtain

$$(36) \quad \nu_x = \nu_{2\xi} \cdot \prod_{i=2\xi+1}^x \frac{u_i}{v_i} \leq \nu_{2\xi} \cdot 2^{2\xi-x} \leq 2^{2\xi-x} \quad \text{for } x > 2\xi.$$

This completes the proof of (33). We proceed with the proof of (34) and assume $q \leq \frac{1}{4}$. In this case, one may verify that $\frac{u_s}{v_s} \leq \frac{1}{2}$ for $s > 1$, while $\frac{u_1}{v_1} \leq 2q$. Similarly, as we obtained (36), we conclude

$$\nu_x \leq \frac{q}{2^{x-2}} \nu_0 \leq \frac{q}{2^{x-2}} \quad \text{for } x > 0. \quad \square$$

We proceed to prove Proposition 3.12. Let \hat{k} be an (∞, q) -arc chain. We shall write \mathbb{E}^k to denote the expectation under the measure where $\hat{k}_0 = k$. We also define the stopping times:

$$\tau_i = \min\{t \geq 0 \mid \hat{k}_t \leq i\}, \quad i \geq 0.$$

CLAIM 3.14. *If $k > 4\xi$, then $\mathbb{E}^k[3^{\tau_{4\xi}}] \leq 9^{k-4\xi}$ and if $1 \leq k \leq 4\xi$ then $\mathbb{E}^k[\lambda_{k,q}^{\tau_{k-1}}] < \infty$ for some $\lambda_{k,q} > 1$.*

PROOF. Suppose first that $k > 4\xi$. Consider the random sequence $X_t := 3^{2\hat{k}_t - 8\xi + t}$ and note that $X_0 = 9^{k-4\xi}$. Let u_i and v_i be as in (26) with $v_0 := 0$. The sequence (X_t) satisfies

$$\mathbb{E}^k[X_{t+1} \mid \hat{k}_t = i] = \left(27 \cdot u_{i+1} + 3 \cdot (1 - u_{i+1} - v_i) + \frac{1}{3} v_i \right) \cdot X_t.$$

One may verify that for all $i > 4\xi$ one has $27 \cdot u_{i+1} + 3 \cdot (1 - u_{i+1} - v_i) + \frac{1}{3} v_i \leq \frac{3}{4}$. This implies that

$$\mathbb{E}^k[X_{t+1} \mid \hat{k}_t = i] \leq \frac{3}{4} X_t \quad \text{for all } i > 4\xi.$$

Denoting $t \wedge \tau_{4\xi} = \min\{t, \tau_{4\xi}\}$ and using the facts that $\tau_{4\xi}$ is a stopping time for (X_t) and $\hat{k}_t > 4\xi$ when $\tau_{4\xi} > t$, it follows that $\mathbb{E}^k[X_{t \wedge \tau_{4\xi}}]$ is nonincreasing in t . As $3^{t \wedge \tau_{4\xi}} \leq X_{t \wedge \tau_{4\xi}}$ we conclude by the monotone convergence theorem that

$$\mathbb{E}^k[3^{\tau_{4\xi}}] = \lim_{t \rightarrow \infty} \mathbb{E}^k[3^{t \wedge \tau_{4\xi}}] \leq \lim_{t \rightarrow \infty} \mathbb{E}^k[X_{t \wedge \tau_{4\xi}}] \leq X_0 = 9^{k-4\xi}.$$

Now suppose that $1 \leq k \leq 4\xi$. By induction and using the previous case we may assume that there exists some $\lambda_{k+1,q} > 1$ for which

$$(37) \quad \mathbb{E}^{k+1}[\lambda_{k+1,q}^{\tau_k}] < \infty.$$

Let $t \geq 1$ and $1 < \lambda < \lambda_{k+1,q}$. By conditioning on the first step of the Markov chain, we have

$$\begin{aligned} & \mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t}] \\ (38) \quad &= v_k \lambda + (1 - u_{k+1} - v_k) \mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t} \mid \hat{k}_1 = k] + u_{k+1} \mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t} \mid \hat{k}_1 = k + 1] \\ &= v_k \lambda + (1 - u_{k+1} - v_k) \mathbb{E}^k[\lambda^{(1+\tau_{k-1}) \wedge t}] + u_{k+1} \mathbb{E}^{k+1}[\lambda^{(1+\tau_{k-1}) \wedge t}]. \end{aligned}$$

Now observe that $\mathbb{E}^k[\lambda^{(1+\tau_{k-1}) \wedge t}] \leq \lambda \cdot \mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t}]$. In addition,

$$\mathbb{E}^{k+1}[\lambda^{(1+\tau_{k-1}) \wedge t}] \leq \lambda \cdot \mathbb{E}^{k+1}[\lambda^{\tau_k + (\tau_{k-1} - \tau_k) \wedge t}] = \lambda \cdot \mathbb{E}^{k+1}[\lambda^{\tau_k}] \cdot \mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t}],$$

where in the last equality we used the strong Markov property and the fact that τ_k is almost surely finite under the measure where $\hat{k}_0 = k + 1$ by (37). Substituting these two bounds into (38) and rearranging the terms, we conclude that

$$(1 - (1 - u_{k+1} - v_k) \cdot \lambda - u_{k+1} \cdot \lambda \cdot \mathbb{E}^{k+1}[\lambda^{\tau_k}]) \cdot \mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t}] \leq v_k \cdot \lambda.$$

Thus, using (37), we may pick $\lambda > 1$ sufficiently small to make the coefficient of $\mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t}]$ positive. With this choice, we conclude that $\mathbb{E}^k[\lambda^{\tau_{k-1} \wedge t}]$ is bounded uniformly in t . Taking the limit $t \rightarrow \infty$ completes the proof. \square

PROOF OF PROPOSITION 3.12. The proposition is trivial for $k = 0$ so we assume that $k \geq 1$. Observe that

$$T = T_{4\xi} + \sum_{i=0}^{4\xi-1} (\tau_i - \tau_{i+1}).$$

The Cauchy–Schwarz inequality then implies that

$$T^2 \leq (4\xi + 1) \left(T_{4\xi}^2 + \sum_{i=0}^{4\xi-1} (\tau_i - \tau_{i+1})^2 \right).$$

If $i \geq k$, then $\tau_i = \tau_{i+1}$ whereas if $i < k$ then the strong Markov property and Claim 3.14 imply that

$$\mathbb{E}^k[(\tau_i - \tau_{i+1})^2] = \mathbb{E}^{i+1}[\tau_i^2] < \infty.$$

Similarly, if $4\xi \geq k$ then $T_{4\xi} = 0$ while if $4\xi < k$ it follows from Claim 3.14 and the fact that $\log_3^2 x$ is concave for $x \geq e$ that

$$\mathbb{E}^k[\tau_{4\xi}^2] = \mathbb{E}^k[\log_3^2(3^{\tau_{4\xi}})] \leq \log_3^2 \mathbb{E}^k[3^{\tau_{4\xi}}] \leq g(q)k^2$$

for some $g(q) > 0$. Combining all of the above facts, we conclude that

$$\mathbb{E}^k[T^2] \leq g(q)(4\xi + 1)k^2 + (4\xi + 1) \cdot 4\xi \cdot \max_{0 \leq i \leq 4\xi-1} \mathbb{E}^k[(\tau_i - \tau_{i+1})^2] \leq h(q)k^2$$

for some $h(q) > 0$. This bound implies a similar bound in which h is replaced by a monotone nondecreasing function f as $\mathbb{E}^k[T^2]$ is a nondecreasing function of q by Proposition 3.9. \square

3.5. Induced Mallows permutations and a stitching process. Our discussion so far was based on the results of Section 3.1, describing the distribution of a portion of the graph Γ_π of a Mallows permutation π conditioned on the parts of the graph “to the left and below this portion”. In our proof of Theorem 1.2, pertaining to the Poisson–Dirichlet limit law, we will need to understand the distribution of portions of Γ_π under more general conditioning events. Our first result in this section discusses the distribution of Γ_π restricted to a rectangle, given the complementary part of Γ_π . As it turns out, in this case the relative ordering of the points of Γ_π is itself distributed via a Mallows distribution. This is formulated precisely below.

Given a finite set of points $\Gamma \subseteq \mathbb{R}^2$, no two of which have equal x or equal y coordinate, we define the *relative order of Γ* as a permutation λ characterized by the following properties:

$$(39) \quad \begin{aligned} &\text{if } \Gamma = \{(x_1, y_1), \dots, (x_k, y_k)\} \text{ with } x_1 < \dots < x_k \\ &\text{then } \lambda \in \mathbb{S}_k \text{ and } \lambda_i := |\Gamma \cap \{y \leq y_i\}|. \end{aligned}$$

The name relative order stems from the fact that for each pair $i, j \in [k]$, one has $\lambda_i < \lambda_j$ if and only if $y_i < y_j$.

LEMMA 3.15. *Let $R := \{x_1 \leq x \leq x_2, y_1 \leq y \leq y_2\} \subseteq [n]^2$ be a discrete rectangle. Let $\pi \sim \mu_{n,q}$ and Γ_π be the graph of π . When $\Gamma_\pi \cap R$ is nonempty the relative order (39) of $\Gamma_\pi \cap R$, conditioned on $\Gamma_\pi \setminus R$, has the Mallows distribution $\mu_{m,q}$ with $m := |\Gamma_\pi \cap R|$.*

The lemma generalizes the more familiar special case where R is a vertical rectangle $\{x_1 \leq x \leq x_2\}$ (see, for instance, [6], Corollary 2.7, for a proof of this special case). The original paper of Mallows [20] contained a discussion of related facts. In fact, the above lemma can be deduced from the special case, though we will provide a direct proof below.

So far we have discussed methods for sampling the graph of a Mallows permutation iteratively “from beginning to end”. The above lemma gives rise to a method for sampling the graph “from a mid-point”. Indeed, one can consider, say, the top and bottom parts of the graph,

$$(40) \quad A := \Gamma_\pi \cap \{y \leq s\} \quad \text{and} \quad B := \Gamma_\pi \cap \{y > s\}$$

for a given $0 \leq s \leq n$. Then, due to the lemma, the relative orders of A and B are independent and have distributions $\mu_{s,q}$ and $\mu_{n-s,q}$, respectively. We now provide a “stitching” procedure for determining the full permutation π from the relative orders. Indeed, given the relative orders, one may determine the full permutation from the projections of A and B on the x -axis. Defining the process (χ_t) , $0 \leq t \leq n$, by

$$(41) \quad \chi_t := |\Gamma_\pi \cap \{x > t, y \leq s\}|$$

we note that the projection of A on the x -axis is exactly the set of descents of χ , $\{t \in [n] | \chi_t = \chi_{t-1} - 1\}$, and the projection of B on the x -axis is the complementary set. Thus the following lemma provides a computational procedure for determining these projections.

LEMMA 3.16. *The process χ defined in (41) is a time-inhomogeneous Markov chain with transition probabilities given by*

$$(42) \quad \begin{aligned} \mathbb{P}[\chi_{t+1} = \chi_t - 1 | \pi_1, \dots, \pi_t] &= \frac{1 - q^{\chi_t}}{1 - q^{n-t}} \quad \text{and} \\ \mathbb{P}[\chi_{t+1} = \chi_t | \pi_1, \dots, \pi_t] &= \frac{q^{\chi_t} - q^{n-t}}{1 - q^{n-t}}. \end{aligned}$$

We remark that the formulas (42) provide more than the transition probabilities of χ ; namely, that these probabilities remain the same even when conditioning on π_1, \dots, π_t .

Putting together the above two lemmas, we obtain the following method for sampling a Mallows permutation π . Let $0 \leq s \leq n$ and let A and B be as in (40). Observe that, according to Lemma 3.15, the relative order of A is independent of χ , as χ is determined by B . Similarly, the relative order of B is independent of both χ and A . Thus we may sample π by independently sampling χ , the relative order of A and the relative order of B . The sampling of χ can be performed using the transition probabilities given in Lemma 3.16. The usual sampling algorithm (6) may then be used to sample the relative orders of A and B .

The infinite case. Gnedin and Olshanski defined an infinite two-sided Mallows distribution as the unique q -exchangeable measure on one-to-one and onto $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$; see [14] for details. A method for sampling from this distribution was given in [14]. Although the infinite two-sided Mallows distribution is not used in our work, we note here that the above sampling algorithm may be extended to produce another sampling method for it.

We do not define the infinite two-sided Mallows distribution and shall rely only on the fact, proved in [14], Proposition 7.6, that this distribution is the limit of the distributions of finite Mallows permutations in a suitable sense. To give precise meaning to this, let us extend the definition of the finite Mallows permutation to arbitrary finite, nonempty, intervals $I \subset \mathbb{Z}$ by saying that $\pi \sim \mu_{I,q}$ if $\pi : I \rightarrow I$ is a bijection and $\mathbb{P}[\pi]$ is proportional to $q^{\text{inv}(\pi)}$ as in (1). This is the same as saying that $P_I^{-1} \circ \pi \circ P_I \sim \mu_{|I|,q}$ where P_I is the unique increasing bijection from $[|I|]$ to I . We view bijections $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$ as elements of $\mathbb{Z}^{\mathbb{Z}}$ with the product topology and identify each bijection $\pi : I \rightarrow I$ with a bijection $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$ by setting $\pi(i) = i$ for $i \notin I$.

FACT 3.17 ([14], Proposition 7.6). Let (I_n) be an arbitrary sequence of finite, nonempty, intervals increasing to \mathbb{Z} and let $\pi^{(n)} \sim \mu_{I_n,q}$. Then $\pi^{(n)}$ converges in distribution to the infinite two-sided Mallows distribution.

We augment this with the following useful tightness property.

CLAIM 3.18. Let (I_n) be an arbitrary sequence of finite, nonempty, intervals increasing to \mathbb{Z} and let $\pi^{(n)} \sim \mu_{I_n,q}$. Then

$$\limsup_{t \rightarrow \infty} \sup_n \mathbb{P}[\Gamma_{\pi^{(n)}} \cap \{x > t, y \leq 0\} \neq \emptyset] = 0.$$

Here and later, similar to before, given $\pi : I \rightarrow I$ with $I \subseteq \mathbb{Z}$, we write

$$\Gamma_\pi := \{(s, \pi_s) | s \in I\}.$$

PROOF OF CLAIM 3.18. The claim follows either from Theorem 3.11 or from tail bounds on the displacement of elements as in, say, [6], Theorem 1.1. Let us

argue from Theorem 3.11. Write $I_n = \{-a, -a + 1, \dots, b\}$ and assume that $a, b \geq 0$ as otherwise, deterministically, $\Gamma_{\pi^{(n)}}$ does not intersect the quadrant $\{x > 0, y \leq 0\}$. Let $\tilde{\pi}^{(n)} \sim \mu_{b+a+1, q}$. By the definition of $\mu_{I_n, q}$, we see that

$$(43) \quad \mathbb{P}[\Gamma_{\pi^{(n)}} \cap \{x > t, y \leq 0\} \neq \emptyset] = \mathbb{P}[\Gamma_{\tilde{\pi}^{(n)}} \cap \{x > t + a, y \leq a\} \neq \emptyset].$$

Let κ be the arc chain of $\tilde{\pi}^{(n)}$ and define $T := \min\{s \geq a \mid \kappa_s = 0\}$. By (20), assuming also $t \geq 0$,

$$\{\Gamma_{\tilde{\pi}^{(n)}} \cap \{x > t + a, y \leq a\} \neq \emptyset\} \subseteq \{T > t + a\}.$$

Thus Theorem 3.11 and Markov’s inequality imply that the probabilities in (43) are at most c_q/t^2 for some $c_q > 0$ depending only on q , from which the claim follows. \square

We require the notion of an infinite *one-sided* Mallows distribution. We recall that one may define an infinite one-sided Mallows permutation $\pi : \mathbb{N} \rightarrow \mathbb{N}$ with parameter $0 < q < 1$ via the formula (6) with the formal substitution $n = \infty$ and this yields a convenient sampling algorithm. This construction may be slightly generalized: For a countably infinite subset $I \subset \mathbb{Z}$ with either a minimal or maximal element, let P_I be the unique monotone bijection $P_I : \mathbb{N} \rightarrow I$ (increasing if I has a minimal element and decreasing if it has a maximal element). Given two countably infinite $I, J \subset \mathbb{Z}$, each with a minimal or maximal element, one defines the infinite one-sided Mallows distribution from I to J with parameter q as the measure on bijections $\pi : I \rightarrow J$ satisfying that $P_J^{-1} \circ \pi \circ P_I$ has the Mallows distribution from \mathbb{N} to \mathbb{N} with parameter q .

We now describe the sampling algorithm for the infinite two-sided case. Let $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$ have the infinite two-sided Mallows distribution with parameter q . Let

$$A := \Gamma_\pi \cap \{y \leq 0\} \quad \text{and} \quad B := \Gamma_\pi \cap \{y > 0\},$$

be the “bottom” and “top” portions of the graph of π , similar to the finite case. Let A_x and B_x be the projections onto the x -axis of A and B , respectively. It is known that, almost surely, A_x has a maximal element and B_x has a minimal element, and also that

$$(44) \quad \begin{aligned} &\text{conditioned on } A_x \text{ and } B_x, \text{ the restrictions } \pi|_{A_x} \text{ and } \pi|_{B_x} \\ &\text{are independent and have the infinite one-sided Mallows} \\ &\text{distributions with parameter } q. \end{aligned}$$

These facts were also noted and used in one of the sampling algorithms presented in [14]. With a bit of work, they also follow from Fact 3.17: the facts on the maximal and minimal elements follow using Claim 3.18, together with the reversal symmetry (10), and the fact (44) follows from Lemma 3.15. Thus, to complete the description of our sampling method for π it suffices to give an algorithm for sampling the projections A_x and B_x , that is, a method to “stitch” the one-sided infinite bijections to a two-sided infinite bijection.

Define the process χ by

$$\chi_t := |\Gamma_\pi \cap \{x > t, y \leq 0\}|, \quad t \in \mathbb{Z}$$

and note that A_x is exactly the set of descents of χ , that is, $A_x = \{t \in \mathbb{Z} | \chi_t = \chi_{t-1} - 1\}$. The distribution of $(\chi_t)_{t \geq 0}$ is described by the following two facts, whose proof we postpone:

(45) χ_0 is distributed as the stationary distribution of the (∞, q) -arc chain; see (27).

(46) χ is a time-homogeneous Markov chain with transition probabilities given by

$$\begin{aligned} \mathbb{P}[\chi_{t+1} = \chi_t - 1 | \pi_i \text{ for } i \leq t] &= 1 - q^{\chi_t} \quad \text{and} \\ \mathbb{P}[\chi_{t+1} = \chi_t | \pi_i \text{ for } i \leq t] &= q^{\chi_t}. \end{aligned}$$

Thus we may easily sample $A_x \cap \mathbb{N}$. To finish, we need only sample $B_x \setminus \mathbb{N}$ conditioned on $A_x \cap \mathbb{N}$ and χ_0 , as $B_x \setminus \mathbb{N}$ together with $A_x \cap \mathbb{N}$ determine both A_x and B_x . To this end, we rely on the following facts, whose proof is again postponed:

(47) Given χ_0 , $A_x \cap \mathbb{N}$ and $B_x \setminus \mathbb{N}$ are conditionally independent and have the same distribution up to reflection. Precisely, given χ_0 ,

$$B_x \setminus \mathbb{N} \stackrel{d}{=} -(A_x \cap \mathbb{N}) + 1.$$

In conclusion, one may sample A_x and B_x as follows: First, sample χ_0 from the distribution (27). Make two independent samples of $(\chi_t)_{t \geq 0}$, with the same given χ_0 , via the Markov chain transition probabilities in (46). Then take $A_x \cap \mathbb{N}$ to be the set of descents of the first copy of $(\chi_t)_{t \geq 0}$ and reconstruct B_x by taking $1 - (B_x \setminus \mathbb{N})$ to be the set of descents of the second copy of $(\chi_t)_{t \geq 0}$. The sets A_x and B_x are determined from $A_x \cap \mathbb{N}$ and $B_x \setminus \mathbb{N}$. The full permutation π may now be reconstructed using the property (44) and the sampling algorithm for infinite one-sided Mallows permutations.

We now return to prove (45), (46) and (47). Define the discrete intervals

$$I_n = \{-n + 1, -n, \dots, n\}$$

and let $\pi^{(n)} \sim \mu_{I_n, q}$. Define the processes $\chi^{(n)}$ by

$$\chi_t^{(n)} := |\Gamma_{\pi^{(n)}} \cap \{x > t, y \leq 0\}|, \quad t \in \mathbb{Z}.$$

Then Fact 3.17 together with Claim 3.18 imply that

(48) $\chi^{(n)}$ converges in distribution to χ .

Let us elaborate on the proof of this fact. Observe that

$$\begin{aligned} \chi_t^{(n)} - \chi_s^{(n)} &= |\{i | t < i \leq s, \pi_i^{(n)} \leq 0\}| \quad \text{and} \\ \chi_t - \chi_s &= |\{i | t < i \leq s, \pi_i \leq 0\}|, \quad t \leq s, \end{aligned}$$

so that these differences depend only on the value of the permutations at finitely many indices. Thus Fact 3.17 implies that $(\chi_t^{(n)} - \chi_s^{(n)})_{t \leq s}$ converges in distribution to $(\chi_t - \chi_s)_{t \leq s}$. This may be upgraded to (48) by using the fact that $\mathbb{P}[\chi_s \neq 0] \xrightarrow{s \rightarrow \infty} 0$ (since A_x has a maximal element, almost surely) and $\sup_n \mathbb{P}[\chi_s^{(n)} \neq 0] \xrightarrow{s \rightarrow \infty} 0$ by Claim 3.18.

Observe that $\chi_0^{(n)}$ has the distribution of the arc chain of $\pi^{(n)}$ at 0. By using (48), property (45) follows from Proposition 3.8 and property (46) follows from Lemma 3.16. Finally, property (47) is a consequence of the reversal symmetry (10) and Lemma 3.16 applied to $\pi^{(n)}$.

Proofs of Lemma 3.15 and Lemma 3.16. To complete this section, we need only prove these two lemmas.

PROOF OF LEMMA 3.15. Throughout the proof, we condition on $\Gamma_\pi \setminus R$ and assume that $|\Gamma_\pi \cap R| > 0$. Let ρ be the relative order of $\Gamma_\pi \cap R$. Observe that the permutation ρ uniquely determines π and ρ may assume, with positive probability, any value in $\mathbb{S}_{|\Gamma_\pi \cap R|}$. Hence, the distribution of ρ is proportional to $q^{\text{inv}(\pi)}$ by the definition of the Mallows distribution (1); while we need to prove that the distribution of ρ is proportional to $q^{\text{inv}(\rho)}$. Therefore, it suffices to verify that

$$(49) \quad \text{inv}(\pi) - \text{inv}(\rho) \quad \text{is determined by } \Gamma_\pi \setminus R.$$

We say that two points $(v_x, v_y), (w_x, w_y) \in \mathbb{R}^2$ form an inversion if $(v_x - w_x) \cdot (v_y - w_y) < 0$. For two finite subsets $V, W \subset \mathbb{R}^2$ we define

$$\begin{aligned} \text{inv}(V, W) &:= |\{(v, w) \in V \times W | (v, w) \text{ forms an inversion}\}| \quad \text{and} \\ \text{inv}(V) &:= \frac{1}{2} \text{inv}(V, V). \end{aligned}$$

The definitions are chosen so that $\text{inv}(\sigma) = \text{inv}(\Gamma_\sigma)$ for any permutation σ . Consider the following equality:

$$\text{inv}(\pi) = \text{inv}(\Gamma_\pi \setminus R) + \text{inv}(\Gamma_\pi \setminus R, \Gamma_\pi \cap R) + \text{inv}(\Gamma_\pi \cap R).$$

Observe that $\text{inv}(\Gamma_\pi \setminus R)$ is determined by $\Gamma_\pi \setminus R$ and that $\text{inv}(\rho) = \text{inv}(\Gamma_\pi \cap R)$. Thus we need only prove that

$$(50) \quad \begin{aligned} &\text{inv}(\Gamma_\pi \setminus R, \Gamma_\pi \cap R) \\ &= \sum_{u \in \Gamma_\pi \setminus R} \text{inv}(\{u\}, \Gamma_\pi \cap R) \quad \text{is determined by } \Gamma_\pi \setminus R. \end{aligned}$$

Let $u = (u_x, u_y) \in \Gamma_\pi \setminus R$. Since $u \in R^c$, we know that at least one of four inequalities occur: $u_x < x_1, u_x > x_2, u_y < y_1, u_y > y_2$. Assume $u_x > x_2$. We have

$$\begin{aligned} \text{inv}(\{u\}, \Gamma_\pi \cap R) &= |\{(x, y) \in \Gamma_\pi \cap R \mid y > u_y\}| \\ &= n - u_y - |\{(x, y) \in \Gamma_\pi \setminus R \mid y > u_y\}|. \end{aligned}$$

Thus $\text{inv}(\{u\}, \Gamma_\pi \cap R)$ is determined by $\Gamma_\pi \setminus R$ and u . Applying similar reasoning in the other three cases shows that (50), and hence (49) holds. \square

PROOF OF LEMMA 3.16. As $0 \leq \chi_t - \chi_{t+1} \leq 1$, it suffices to establish the formula for $\mathbb{P}[\chi_{t+1} = \chi_t - 1 \mid \pi_1, \dots, \pi_t]$. It is convenient to use the formula (6). Let $j_1 < \dots < j_{n-t}$ be the elements of $[n] \setminus \{\pi_1, \dots, \pi_t\}$. Observe that $j_k \leq s$ if and only if $k \leq \chi_t$. Hence by (6) and the definition of χ_t we have

$$\mathbb{P}[\chi_{t+1} = \chi_t - 1 \mid \pi_1, \dots, \pi_t] = \sum_{k=1}^{\chi_t} \mathbb{P}[\pi_{t+1} = j_k \mid \pi_1, \dots, \pi_t] = \frac{1 - q^{\chi_t}}{1 - q^{n-t}}. \quad \square$$

4. Main theorems. We start by introducing several definitions which we will need for proving our main theorems. A nonempty subset \mathfrak{a} of $[n]$ is called an *arc* of the permutation $\pi \in \mathbb{S}_n$ if its elements can be ordered so that $\mathfrak{a} = \{a_1, \dots, a_{|\mathfrak{a}|}\}$ with $\pi(a_i) = a_{i+1}$ for $1 \leq i < |\mathfrak{a}|$. We say the arc has *length* $|\mathfrak{a}|$ and say the arc is *closed* if it forms a cycle, that is, if also $\pi(a_{|\mathfrak{a}|}) = a_1$. A nonclosed arc is called *open*. When the arc is open, the above ordering is unique, in which case we call a_1 and $a_{|\mathfrak{a}|}$ the *tail* and *head* of the arc \mathfrak{a} , respectively, and denote them by $\text{tail}(\mathfrak{a})$ and $\text{head}(\mathfrak{a})$.

In our proofs of the main theorems, we will rely upon the diagonal exposure process introduced in Section 3.2. We recall that by time t of this process, we expose the portion of the graph Γ_π contained in $\{x \leq t, y \leq t\}$. This information allows us to determine all arcs which are contained in $[t]$ and, moreover, to tell whether each such arc is open or closed. This motivates the following definitions.

Let $\pi \in \mathbb{S}_n$ and $0 \leq t \leq n$. We say that an arc \mathfrak{a} of π is $[t]$ -maximal (with respect to inclusion) if $\mathfrak{a} \subseteq [t]$ and if every arc $\mathfrak{b} \subseteq [t]$ which contains \mathfrak{a} is in fact equal to \mathfrak{a} . Denote

$$\begin{aligned} \mathcal{A}_t(\pi) &:= \{\mathfrak{a} \mid \mathfrak{a} \text{ is } [t]\text{-maximal}\}, \\ \mathcal{O}_t(\pi) &:= \{\mathfrak{a} \mid \mathfrak{a} \text{ is } [t]\text{-maximal and open}\}. \end{aligned}$$

Recalling the definition of the arc chain κ from (19) we observe that

$$(51) \quad \kappa_t(\pi) = |\mathcal{O}_t(\pi)|,$$

which is the origin of the name ‘‘arc chain’’. We note further that for each $1 \leq s \leq t$ there exists a unique $[t]$ -maximal arc containing s and we denote this arc by $\text{arc}_s^t(\pi)$. When the permutation π is clear from the context, we shall abbreviate $\text{arc}_s^t(\pi), \mathcal{A}_t(\pi)$ and $\mathcal{O}_t(\pi)$ to $\text{arc}_s^t, \mathcal{A}_t$ and \mathcal{O}_t .

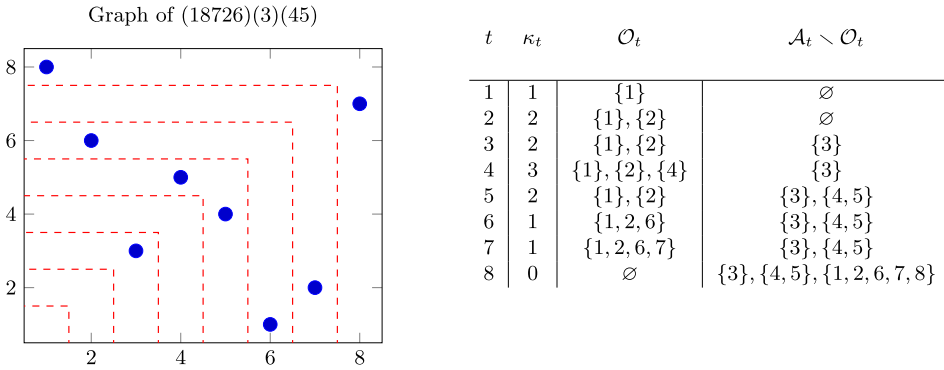


FIG. 5. Graph of the permutation (18726)(3)(45) alongside its arc chain process (κ_t) , its maximal open arcs process (\mathcal{O}_t) and its maximal closed arcs process $(\mathcal{A}_t \setminus \mathcal{O}_t)$.

Let us describe how \mathcal{A}_t evolves during the diagonal exposure process, that is, the relationship between \mathcal{A}_t and \mathcal{A}_{t+1} ; see Figure 5 for an example. The newly exposed portions of the graph Γ_π at time $t + 1$ were described in (17). Thus the set of $[t + 1]$ -maximal arcs is formed from the set of $[t]$ -maximal arcs by having the element $t + 1$ either: (i) form an, open or closed, arc by itself; (ii) extend an open arc to a new, open or closed, longer arc; or (iii) merge two open arcs into a longer open arc. These three possibilities are considered below according to their effect on the number of open arcs:

- If $\kappa_{t+1} = \kappa_t + 1$, then necessarily $\{t + 1\}$ is a $[t + 1]$ -maximal open arc and \mathcal{A}_{t+1} equals \mathcal{A}_t with $\{t + 1\}$ added.
- If $\kappa_{t+1} = \kappa_t$, then either $t + 1$ is a fixed point of π or $t + 1$ extends an open arc in \mathcal{A}_t to a longer open arc in \mathcal{A}_{t+1} , either as the head or as the tail of the arc.
- If $\kappa_{t+1} = \kappa_t - 1$, then either two open arcs were merged via $t + 1$ or an open arc was extended by $t + 1$ to a closed arc.

We now consider the probabilities for the process (\mathcal{A}_t) to evolve according to the above possibilities when π is a Mallows random permutation. As mentioned above, we note that \mathcal{A}_t and \mathcal{O}_t are measurable with respect to \mathcal{F}_t , that is, \mathcal{A} and \mathcal{O} are adapted to the diagonal exposure filtration.

The probability, conditioned on \mathcal{F}_t , that the element $t + 1$ forms a $[t + 1]$ -maximal arc by itself was already calculated in Proposition 3.3 (for the case that it forms an open arc, or equivalently that $\kappa_{t+1} = \kappa_t + 1$) and Lemma 3.4 (for the case that it is a fixed point).

The other options, in which the element $t + 1$ either extends an existing arc or merges two arcs, are determined from the basic events $\{\pi_{t+1}^{-1} = \text{head}(\mathbf{a})\}$ and $\{\pi_{t+1} = \text{tail}(\mathbf{a})\}$ for an arc $\mathbf{a} \in \mathcal{O}_t$. Indeed, for the merging event one checks in a straightforward manner that if $\mathbf{a}, \mathbf{b} \in \mathcal{O}_t$ are distinct arcs then

$$\{\mathbf{a} \cup \{t + 1\} \cup \mathbf{b} \in \mathcal{O}_{t+1} \text{ and } \pi_{t+1} \in \mathbf{b}\} \iff$$

$$\{\pi_{t+1}^{-1} = \text{head}(\mathbf{a}) \text{ and } \pi_{t+1} = \text{tail}(\mathbf{b})\}$$

and for the event that $t + 1$ extends the open arc $\mathbf{a} \in \mathcal{O}_t$ to a closed arc we have

$$\{\mathbf{a} \cup \{t + 1\} \in \mathcal{A}_{t+1} \setminus \mathcal{O}_{t+1}\} \iff \{\pi_{t+1}^{-1} = \text{head}(\mathbf{a}) \text{ and } \pi_{t+1} = \text{tail}(\mathbf{a})\}.$$

The event that $t + 1$ extends an open arc to a longer open arc is the complement of the other possibilities. Thus the probabilities of these events may be derived from the following lemma.

LEMMA 4.1. *Let $\pi \sim \mu_{n,q}$, let $0 \leq t < n$ and let $\mathbf{a}, \mathbf{b} \in \mathcal{O}_t$ be two, not necessarily distinct, arcs. The events $\{\pi_{t+1}^{-1} = \text{head}(\mathbf{a})\}$ and $\{\pi_{t+1} = \text{tail}(\mathbf{b})\}$ are conditionally independent given \mathcal{F}_t . Furthermore,*

$$(52) \quad \begin{aligned} \mathbb{P}_t[\pi_{t+1}^{-1} = \text{head}(\mathbf{a})] &= q^{i_{\mathbf{a}}^h} \cdot \frac{1 - q}{1 - q^{n-t}} \quad \text{and} \\ \mathbb{P}_t[\pi_{t+1} = \text{tail}(\mathbf{b})] &= q^{i_{\mathbf{b}}^t} \cdot \frac{1 - q}{1 - q^{n-t}}, \end{aligned}$$

where

$$i_{\mathbf{a}}^h = |\{\mathbf{c} \in \mathcal{O}_t \mid \text{head}(\mathbf{c}) < \text{head}(\mathbf{a})\}| \quad \text{and} \quad i_{\mathbf{b}}^t = |\{\mathbf{c} \in \mathcal{O}_t \mid \text{tail}(\mathbf{c}) < \text{tail}(\mathbf{b})\}|.$$

A consequence of the lemma is that for $q \geq \frac{1}{2}$ and any two open arcs $\mathbf{a}, \mathbf{b} \in \mathcal{O}_t$ one has

$$\mathbb{P}_t[\pi_{t+1}^{-1} = \text{head}(\mathbf{a}) \text{ and } \pi_{t+1} = \text{tail}(\mathbf{b})] \approx \left(\frac{1 - q}{1 - q^{n-t}}\right)^2 \quad \text{on } \{\kappa_t \leq 2\xi\},$$

which eventually leads to the appearance of the length scale $\frac{1}{(1-q)^2}$ in our theorems.

PROOF OF LEMMA 4.1. The fact that the events $\{\pi_{t+1}^{-1} = \text{head}(\mathbf{a})\}$ and $\{\pi_{t+1} = \text{tail}(\mathbf{b})\}$ are conditionally independent given \mathcal{F}_t is a direct consequence of (14).

It is not difficult to verify that

$$i_{\mathbf{b}}^t = |\{1 \leq j < \text{tail}(\mathbf{b}) \mid j \notin \{\pi_1, \dots, \pi_t\}\}|.$$

Thus the sampling formula (6) yields that

$$\begin{aligned} \mathbb{P}_t[\pi_{t+1} = \text{tail}(\mathbf{b})] &= \mathbb{E}_t \mathbb{P}[\pi_{t+1} = \text{tail}(\mathbf{b}) \mid \pi_1, \dots, \pi_t] \\ &= \mathbb{E}_t \left[q^{i_{\mathbf{b}}^t} \cdot \frac{1 - q}{1 - q^{n-t}} \right] = q^{i_{\mathbf{b}}^t} \cdot \frac{1 - q}{1 - q^{n-t}} \end{aligned}$$

as we wanted to show. The formula for $\mathbb{P}_t[\pi_{t+1}^{-1} = \text{head}(\mathbf{a})]$ follows from the formula for $\mathbb{P}_t[\pi_{t+1} = \text{tail}(\mathbf{b})]$ applied to the inverse permutation π^{-1} , as we have the inverse symmetry (9) and the fact that \mathcal{F}_t and \mathcal{O}_t are invariant under this symmetry. \square

4.1. *Expected number of cycles.* In this section, we prove Theorem 1.5. Let $\pi \sim \mu_{n,q}$ and κ be the arc chain of π . Recall from the Introduction that \mathcal{C}_s is the set of points in the cycle of π which contains s .

Our proof is based on the fact that the number of cycles in π equals the number of points $s \in [n]$ that satisfy $s = \max(\mathcal{C}_s)$, that is,

$$(53) \quad \text{number of cycles in } \pi = |\{s \in [n] | s = \max(\mathcal{C}_s)\}| = \sum_{s \in [n]} \mathbb{1}_{\{s = \max(\mathcal{C}_s)\}}.$$

The following proposition provides an estimate for the conditional probability of the event $\{s = \max(\mathcal{C}_s)\}$ given \mathcal{F}_{s-1} .

PROPOSITION 4.2. *For $1 \leq s \leq n$, one has*

$$q^{\kappa_{s-1}} \cdot \frac{1 - q}{1 - q^{n-s+1}} \leq \mathbb{P}_{s-1}[s = \max(\mathcal{C}_s)] \leq \frac{1 - q}{1 - q^{n-s+1}}.$$

PROOF. Given π_1, \dots, π_{s-1} there is a unique element $t \in [n]$, distinct from π_1, \dots, π_{s-1} , such that $s = \max(\mathcal{C}_s)$ if and only if $\pi_s = t$ (indeed, there is a unique $k \geq 1$ and unique distinct $t_1, \dots, t_k \in [s]$ satisfying $t_k = s, \pi_{t_i} = t_{i+1}$ for $1 \leq i \leq k - 1$ and $t_1 \notin \{\pi_1, \dots, \pi_{s-1}\}$ whence we set $t = t_1$). We shall derive the proposition from this fact and formula (6). Write $j_1 < \dots < j_{n-s+1}$ for the elements of $[n] \setminus \{\pi_1, \dots, \pi_{s-1}\}$ and let k be such that $t = j_k$. Then

$$\mathbb{P}_{s-1}[s = \max(\mathcal{C}_s)] = \mathbb{E}_{s-1} \mathbb{P}[\pi_s = t | \pi_1, \dots, \pi_{s-1}] = \mathbb{E}_{s-1} \left[q^{k-1} \frac{1 - q}{1 - q^{n-s+1}} \right].$$

The proposition follows as $0 \leq k - 1 \leq \kappa_{s-1}$. \square

We augment this proposition with the following simple estimate on $q^{\kappa_{s-1}}$.

CLAIM 4.3. *For $1 \leq s \leq n$, one has*

$$\mathbb{E}[q^{\kappa_{s-1}}] \approx 1.$$

PROOF. Clearly, $q^{\kappa_{s-1}} \leq 1$ so we only need to prove the lower bound. We consider two cases. If $q \geq \frac{1}{8}$ then, by Theorem 3.6,

$$\mathbb{E}[q^{\kappa_{s-1}}] \geq q^{2\xi} \mathbb{P}[\kappa_{s-1} \leq 2\xi] \geq \frac{1}{4} \left(1 - \frac{q^{\xi^2 + \xi}}{1 - q^{2\xi}} \right) \gtrsim 1.$$

Whereas if $q \leq \frac{1}{8}$ then, by (30) and Lemma 3.13,

$$\mathbb{E}[q^{\kappa_{s-1}}] \geq \mathbb{P}[\kappa_{s-1} = 0] \gtrsim 1. \quad \square$$

Putting together (53) and the previous two estimates shows that

$$\mathbb{E}[\text{number of cycles in } \pi] \approx \sum_{s=1}^n \frac{1 - q}{1 - q^{n-s+1}}.$$

Theorem 1.5 is a direct consequence of this fact, together with the observation that

$$(54) \quad \frac{1 - q}{1 - q^k} \approx 1 - q + \frac{1}{k}, \quad k \geq 1.$$

To verify (54), we consider two cases. Recall ξ from (8).

If $k \geq \xi$, then $q^k \leq \frac{1}{2}$ and $\frac{1}{k} \lesssim 1 - q$; therefore $\frac{1 - q}{1 - q^k} \approx 1 - q \approx 1 - q + \frac{1}{k}$.

If $k < \xi$, then $q^k > \frac{1}{2}$ and $\frac{1}{k} \gtrsim 1 - q$; thus $\frac{1 - q}{1 - q^k} = \frac{1 - q}{1 + q + \dots + q^{k-1}} \approx \frac{1}{k} \approx 1 - q + \frac{1}{k}$.

4.2. *Expected cycle diameter.* In this section, we prove Theorem 1.4. Some of the tools developed here will also be used in proving the rest of our main theorems.

We prove the lower and upper bounds on the quantity $\mathbb{E}[\max(\mathcal{C}_s) - s]$ as given in (4). The other bounds in the theorem follow: The bounds on $\mathbb{E}[s - \min(\mathcal{C}_s)]$ given in (5) are equivalent to those of (4) via the reversal symmetry (10). Put together, the bounds in (4) and (5) yield the bounds on $\mathbb{E}[\max(\mathcal{C}_s) - \min(\mathcal{C}_s)]$ given in (3).

Throughout this section, we let $\pi \sim \mu_{n,q}$ and κ be the arc chain of π .

Lower bound. Let us begin with the proof of the lower bound of (4), that is,

$$(55) \quad \mathbb{E}[\max(\mathcal{C}_s) - s] \gtrsim \min\{q \cdot \xi^2, n - s\},$$

where we recall the definition of ξ from (8). First, we are going to provide upper bounds for closing the cycle of s at time $s \leq t < n$. Observe that for $s \leq t < n$, conditioned on the event $\{\text{arc}_s^t \text{ is open}\}$, one has

$$\text{arc}_s^{t+1} \text{ is closed} \iff \pi_{t+1}^{-1} = \text{head}(\text{arc}_s^t) \quad \text{and} \quad \pi_{t+1} = \text{tail}(\text{arc}_s^t).$$

Hence the following equation is an immediate corollary of Lemma 4.1:

$$(56) \quad \text{for } s \leq t < n, \quad \mathbb{P}_t[\text{arc}_s^{t+1} \text{ is closed}] \leq \left(\frac{1 - q}{1 - q^{n-t}}\right)^2 \quad \text{on } \{\text{arc}_s^t \in \mathcal{O}_t\}.$$

It is worth noting that the probability of closing the cycle of s at time $t = s$ has a larger estimate, as given in Proposition 4.2, as the event $s = \max(\mathcal{C}_s)$ occurs either when s is a fixed point or when any one of the open arcs closes at time $t = s$.

Combining Proposition 4.2 with (56), we derive by induction the following corollary.

COROLLARY 4.4. *For $0 < s \leq t \leq n$, one has³*

$$\mathbb{P}_{s-1}[\text{arc}_s^t \text{ is open}] \geq \frac{q - q^{n-s+1}}{1 - q^{n-s+1}} \prod_{s \leq i < t} \left(1 - \left(\frac{1 - q}{1 - q^{n-i}}\right)^2\right).$$

Now we state the main proposition.

³In the case when $s = t$, the empty product is assumed to be 1.

PROPOSITION 4.5. *Let $1 \leq s < n$ and set $r := \min\{s + \xi^2, n - 1\}$. Then for $q \geq \frac{1}{2}$ we have*

$$\mathbb{P}_{s-1}[\text{arc}_s^r \text{ is open}] \gtrsim 1.$$

Observe that equation (55) in the case $q \geq \frac{1}{2}$ is an immediate corollary of this proposition. One may also verify that Proposition 4.2 yields (55) in the case $q \leq \frac{1}{2}$. Hence it suffices to prove Proposition 4.5.

We use the following calculus fact: for a sequence (x_i) with $x_i \in [0, 1)$, one has

$$(57) \quad \text{if } \sum_i x_i \lesssim 1 \text{ and } 1 - x_i \gtrsim 1 \text{ for all } i, \text{ then } \prod_i (1 - x_i) \gtrsim 1.$$

PROOF OF PROPOSITION 4.5. Due to Corollary 4.4 and the fact that $\frac{q - q^{n-s+1}}{1 - q^{n-s+1}} \geq \frac{q - q^2}{1 - q^2} \geq \frac{1}{3}$, it suffices to verify that

$$\prod_{s \leq i < r} \left(1 - \left(\frac{1 - q}{1 - q^{n-i}} \right)^2 \right) \gtrsim 1.$$

Due to (57), we need only verify that one has

$$1 - \left(\frac{1 - q}{1 - q^{n-i}} \right)^2 \gtrsim 1, \quad \text{for } s \leq i < r, \quad \text{and} \quad \sum_{s \leq i < r} \left(\frac{1 - q}{1 - q^{n-i}} \right)^2 \lesssim 1.$$

The first inequality follows from the fact that $q \geq \frac{1}{2}$ and $i < n - 1$. For the second inequality, observe that

$$\sum_{s \leq i < r} \left(\frac{1 - q}{1 - q^{n-i}} \right)^2 \leq \sum_{t=1}^{r-s} \left(\frac{1 - q}{1 - q^t} \right)^2 \leq \sum_{t=1}^{\xi^2} \left(\frac{1 - q}{1 - q^t} \right)^2 \lesssim \sum_{t=1}^{\xi^2} \left((1 - q)^2 + \frac{1}{t^2} \right) \lesssim 1,$$

where we applied the estimate $\left(\frac{1-q}{1-q^t}\right)^2 \lesssim (1 - q)^2 + \frac{1}{t^2}$ which follows from (54). □

Upper bound. We proceed to establish the upper bound

$$(58) \quad \mathbb{E}[\max(\mathcal{C}_s) - s] \lesssim \min\{q \cdot \xi^2, n - s\}.$$

We are going to state analogues to (56) and Corollary 4.4 for the upper bound. Lemma 4.1 implies that

$$(59) \quad \text{for } s \leq t < n, \quad \mathbb{P}_t[\text{arc}_s^{t+1} \text{ is closed}] \geq \left(\frac{q^{\kappa_t} - q^{\kappa_t+1}}{1 - q^{n-t}} \right)^2.$$

PROPOSITION 4.6. *For $1 \leq s \leq t \leq n$ and $d \geq 0$, we have*

$$(60) \quad \mathbb{P}_s[\text{arc}_s^t \text{ is open and } \kappa_i \leq d \text{ for all } i \in [s, t]] \leq \exp(-(t - s) \cdot (1 - q)^2 \cdot q^{2d}).$$

PROOF. For $s < r \leq n$, we set

$$E_r := \{\text{arc}_s^r \text{ is open and } \kappa_i \leq d \text{ for all } i \in [s, r]\}.$$

Via (59) one may verify that for such r ,

$$\begin{aligned} \mathbb{P}_s[E_r] &\leq \mathbb{P}_s[E_{r-1} \cap \{\text{arc}_s^r \text{ is open}\}] \leq \mathbb{P}_s[E_{r-1}] \cdot \left(1 - \left(\frac{q^d - q^{d+1}}{1 - q^{n-r+1}}\right)^2\right) \\ &\leq \mathbb{P}_s[E_{r-1}] \cdot \exp(-(1 - q)^2 \cdot q^{2d}). \end{aligned}$$

By applying this inequality for all $s < r \leq t$, we obtain inequality (60). \square

PROPOSITION 4.7. For $1 \leq s \leq n$, one has

$$\mathbb{E}[(\max(C_s) - s)^2] \lesssim \min\{q \cdot \xi^4, (n - s)^2\}.$$

This proposition implies (58) for $q \geq \frac{1}{2}$ using the fact that $\mathbb{E}[X^2] \geq \mathbb{E}[X]^2$ for any integrable random variable X . For $q \leq \frac{1}{2}$ we use the fact that $\max(C_s) - s \leq (\max(C_s) - s)^2$ to again deduce (58) from the same proposition.

Proposition 4.7 is stronger than what we need here as it bounds the second moment of $\max(C_s) - s$, but this extra strength will be used in the proof of Theorem 1.3.

PROOF. The inequality $\max(C_s) - s \leq n - s$ holds true by definition. So we need only verify $\mathbb{E}[(\max(C_s) - s)^2] \lesssim q \cdot \xi^4$. In the regime when q is bounded away from 1, the inequality is a direct consequence of Theorem 3.11 due to

$$\kappa_t = 0 \implies \max(C_i) \leq t \quad \text{for all } i \leq t,$$

as κ_t counts the number of open arcs at time t , see (51). Thus to complete the proof of proposition it suffices to verify that

$$(61) \quad \mathbb{E}[(\max(C_s) - s)^2] \lesssim \xi^4, \quad \text{for } q \text{ sufficiently close to 1.}$$

We assume without loss of generality that $n > s + 2\xi^2$. Our starting point is the inequality

$$(62) \quad \mathbb{E}[(\max(C_s) - s)^2] \leq 4\xi^4 + 3 \sum_{t=s+2\xi^2}^{n-1} (t - s) \cdot \mathbb{P}[\max(C_s) > t].$$

Define

$$d_t := \left\lceil \frac{1}{4} \log_{1/q} \left(\frac{t - s}{\xi^2} \right) \right\rceil \quad \text{and} \quad E_t := \{\kappa_i \leq d_t + \xi \text{ for all } i \in [s, t]\}.$$

Now write

$$\begin{aligned} & \sum_{t=s+2\xi^2}^{n-1} (t-s)\mathbb{P}[\max(\mathcal{C}_s) > t] \\ & \leq \sum_{t=s+2\xi^2}^{n-1} (t-s)\mathbb{P}[E_t^c] + \sum_{t=s+2\xi^2}^{n-1} (t-s)\mathbb{P}[\{\max(\mathcal{C}_s) > t\} \cap E_t]. \end{aligned}$$

We estimate the two sums separately. For the first sum, using Theorem 3.6 and the observation that $d_t \geq 12$ for $t \geq s + 2\xi^2$ as we assumed that q is sufficiently close to 1, we obtain

$$\sum_{t=s+2\xi^2}^{n-1} (t-s)\mathbb{P}[E_t^c] \lesssim \sum_{t=s+2\xi^2}^{n-1} (t-s)q^{12d_t} \leq \sum_{t=s+2\xi^2}^{n-1} (t-s)\left(\frac{\xi^2}{t-s}\right)^3 \lesssim \xi^4.$$

For the second sum, Proposition 4.6 implies that

$$\begin{aligned} & \sum_{t=s+2\xi^2}^{n-1} (t-s)\mathbb{P}[\{\max(\mathcal{C}_s) > t\} \cap E_t] \\ & \leq \sum_{t=s+2\xi^2}^{n-1} (t-s) \exp(-(t-s)(1-q)^2 q^{2d_t+2\xi}) \\ & \leq \sum_{t \geq s+2\xi^2} (t-s) \exp\left(- (t-s)(1-q)^2 \frac{q^{2+2\xi} \xi}{\sqrt{t-s}}\right) \\ & \leq \sum_{t \geq s+2\xi^2} (t-s) \exp(-c\sqrt{t-s}(1-q)) \lesssim \xi^4, \end{aligned}$$

for a positive absolute constant $c > 0$, where the last estimate is not difficult to check directly. \square

4.3. *Expected cycle length.* In this section, we prove Theorem 1.1. As $|\mathcal{C}_s| \leq \text{diam}(\mathcal{C}_s) + 1$ the upper bound for $\mathbb{E}|\mathcal{C}_s|$ follows immediately from the upper bound on the diameter of \mathcal{C}_s proved in Theorem 1.4. Thus we need only prove the lower bound, namely that

$$(63) \quad \mathbb{E}|\mathcal{C}_s| \gtrsim \min\{\xi^2, n\}.$$

Since $|\mathcal{C}_s| \geq 1$, we may and will restrict, for the rest of the section, to the regime where q is sufficiently close to 1 and n is sufficiently large.

Our starting point is the formula

$$\mathbb{E}|\mathcal{C}_s| = \sum_{t \in [n]} \mathbb{P}[t \in \mathcal{C}_s].$$

The lower bound (63) is an immediate consequence of this formula combined with the next lemma.

LEMMA 4.8. *Let $s, t \in [n]$ satisfy $|t - s| \leq \xi^2$ then*

$$\mathbb{P}[t \in \mathcal{C}_s] \gtrsim 1.$$

PROOF. It suffices to consider the case $s < t$. Define the events $(V_r), t \leq r \leq n$ by

$$\begin{aligned} V_t &:= \{\text{arc}_s^t = \text{arc}_t^t\} \quad \text{and} \\ V_r &:= \{\text{arc}_s^r = \text{arc}_t^r \text{ and } \text{arc}_s^{r-1} \neq \text{arc}_t^{r-1}\} \quad \text{for } r > t. \end{aligned}$$

By definition, V_r occurs if and only if the arc of s and the arc of t merge exactly at time r . Hence $\{t \in \mathcal{C}_s\}$ is the disjoint union $\bigsqcup_{t \leq r \leq n} V_r$, yielding that

$$\mathbb{P}[t \in \mathcal{C}_s] = \sum_{t \leq r \leq n} \mathbb{P}[V_r].$$

We shall prove the following estimates:

$$(64) \quad \begin{aligned} \mathbb{P}[V_t] &\gtrsim \frac{1 - q}{1 - q^{n-t+1}} \quad \text{and} \\ \mathbb{P}[V_r] &\gtrsim \left(\frac{1 - q}{1 - q^{n-r+1}}\right)^2 \quad \text{for } t < r < \min\{t + \xi^2, n - 2\}. \end{aligned}$$

The lemma follows easily from these, as if $t + \xi^2 < n - 2$ one may sum the estimates for $\mathbb{P}[V_r]$ and otherwise it suffices to consider only the estimate for $\mathbb{P}[V_t]$ or $\mathbb{P}[V_{n-3}]$. Let us prove these estimates. The following inequality is an immediate consequence of Lemma 4.1:

$$(65) \quad \mathbb{P}_{t-1}[V_t] \geq q^{2\xi} \cdot \frac{1 - q}{1 - q^{n-t+1}} \cdot \mathbb{1}_{A \cap \{\kappa_{t-1} \leq 2\xi\}} \quad \text{where } A := \{\text{arc}_s^{t-1} \text{ is open}\}.$$

Thus the estimate for $\mathbb{P}[V_t]$ follows by noting that $\mathbb{P}[A \cap \{\kappa_{t-1} \leq 2\xi\}] \gtrsim 1$ for q sufficiently close to 1. Indeed, $\mathbb{P}[A] \gtrsim 1$ by Proposition 4.5 and $\mathbb{P}[\kappa_{t-1} \leq 2\xi]$ tends to 1 as q tends to 1, uniformly in t and n , by Theorem 3.6.

We proceed to estimate $\mathbb{P}[V_r]$ for $t < r < \min\{t + \xi^2, n - 2\}$. Define

$$B_i := \{\text{arc}_s^i \text{ and } \text{arc}_t^i \text{ are open and distinct}\}, \quad t \leq i \leq n.$$

Similar to (65), we have

$$\mathbb{P}_{r-1}[V_r] \geq q^{4\xi} \cdot \left(\frac{1 - q}{1 - q^{n-t+1}}\right)^2 \cdot \mathbb{1}_{B_{r-1} \cap \{\kappa_{r-1} \leq 2\xi\}}$$

and, as before, it suffices to show that $\mathbb{P}[B_{r-1}] \gtrsim 1$. Lemma 4.1 implies that for $i > t$,

$$\begin{aligned} \mathbb{P}[B_i] &= \mathbb{E} \left[\mathbb{1}_{B_{i-1}} \cdot \left(1 - \mathbb{P}_{i-1} \left[\begin{array}{l} \pi_i \in \{\text{tail}(\text{arc}_s^{i-1}), \text{tail}(\text{arc}_t^{i-1})\}, \\ \pi_i^{-1} \in \{\text{head}(\text{arc}_s^{i-1}), \text{head}(\text{arc}_t^{i-1})\} \end{array} \right] \right) \right] \\ &\geq \mathbb{P}[B_{i-1}] \cdot \left(1 - 4 \left(\frac{1-q}{1-q^{n-i+1}} \right)^2 \right). \end{aligned}$$

Furthermore, for $i = t$ we obtain using Proposition 4.2 and Lemma 4.1 that

$$\begin{aligned} \mathbb{P}[B_t] &= \mathbb{P}[A \cap \{\pi_t \neq \text{tail}(\text{arc}_s^{t-1}), \pi_t^{-1} \neq \text{head}(\text{arc}_s^{t-1}), t \neq \max(\mathcal{C}_t)\}] \\ &\geq \mathbb{P}[A] \cdot \left(1 - 3 \cdot \frac{1-q}{1-q^{n-t+1}} \right). \end{aligned}$$

Combining these inequalities with the estimate $\mathbb{P}[A] \gtrsim 1$ proved previously shows that

$$\mathbb{P}[B_{r-1}] \gtrsim \left(1 - 3 \cdot \frac{1-q}{1-q^{n-t+1}} \right) \prod_{t < i < r} \left(1 - 4 \left(\frac{1-q}{1-q^{n-i+1}} \right)^2 \right).$$

By estimating the product as in the proof of Proposition 4.5, we conclude that $\mathbb{P}[B_{r-1}] \gtrsim 1$, as we wanted to show. \square

4.4. *Variance of cycle length.* In this section, we prove Theorem 1.3. Throughout we let $\pi \sim \mu_{n,q}$. We need to show that for every $s \in [n]$ we have

$$\text{Var} |\mathcal{C}_s| \approx \min\{q \cdot \xi^4, (n-1)^2\}.$$

We divide the proof into 3 cases:

1. The lower bound for $q \geq \frac{1}{8}$ follows from a general nonconcentration argument together with Theorem 1.1.
2. The lower bound for $q \leq \frac{1}{8}$ is a corollary of Lemma 3.4.
3. The upper bound is a direct corollary of Proposition 4.7.

Case 1: Assume that $q \geq \frac{1}{8}$. Assume that $n > 2$ as the case $n = 1$ is trivial. Further, using the reversal symmetry (10), assume that $s < n$.

We consider the following equivalence relation on \mathbb{S}_n : we say that $\sigma_1 \sim \sigma_2$ if

$$\sigma_1 \in \{\sigma_2, \tau \circ \sigma_2, \sigma_2 \circ \tau, \tau \circ \sigma_2 \circ \tau\} \quad \text{where } \tau := (s, s+1).$$

Let X be the random equivalence class of π in this equivalence relation. We shall prove that

$$(66) \quad \text{Var}(|\mathcal{C}_s| | X) \gtrsim \mathbb{E}[|\mathcal{C}_s|^2 | X].$$

The proof of the lower bound in this case follows from the inequality together with Theorem 1.1 as

$$\text{Var}(|\mathcal{C}_s|) \geq \mathbb{E}[\text{Var}(|\mathcal{C}_s| | X)] \gtrsim \mathbb{E}[|\mathcal{C}_s|^2] \geq \mathbb{E}[|\mathcal{C}_s|]^2.$$

Let us proceed with the proof of (66). Composing a permutation with an adjacent transposition, like $\tau = (s, s + 1)$, changes the number of inversions in the permutation exactly by 1. It follows that any two permutations $\sigma_1 \sim \sigma_2$ satisfy $|\text{inv}(\sigma_1) - \text{inv}(\sigma_2)| \leq 2$, whence (1) implies that

$$\mathbb{P}[\pi = \sigma_1] \approx \mathbb{P}[\pi = \sigma_2].$$

As X is an equivalence class of \mathbb{S}_n of size at most 4, we conclude that

$$(67) \quad \mathbb{P}[\pi = \sigma | X] \approx 1, \quad \text{for all } \sigma \in X.$$

The equivalence class X necessarily contains a permutation ρ satisfying $\mathcal{C}_s(\rho) \neq \mathcal{C}_{s+1}(\rho)$. Choosing such a ρ , one checks that for each $\sigma \in X$, $|\mathcal{C}_s(\sigma)|$ is either $|\mathcal{C}_s(\rho)|$, $|\mathcal{C}_{s+1}(\rho)|$ or $|\mathcal{C}_s(\rho)| + |\mathcal{C}_{s+1}(\rho)|$ and each of these values occurs for some $\sigma \in X$. Thus (66) is a consequence of (67).

REMARK 4.9. The above argument is a general argument for showing non-concentration of cycle lengths, that is,

$$\text{Var}(|\mathcal{C}_s|) \gtrsim \mathbb{E}[|\mathcal{C}_s|^2].$$

It may be applied to other random permutation models satisfying the following assumption. There exists some $1 \leq t \leq n$, different from s , for which

$$\mathbb{P}[\sigma] \approx \mathbb{P}[\sigma \circ (s, t)] \approx \mathbb{P}[(s, t) \circ \sigma] \quad \text{for all } \sigma \in \mathbb{S}_n.$$

Case 2: Assume now that $q \leq \frac{1}{8}$. We need to prove that

$$\text{Var } |\mathcal{C}_s| \gtrsim q \quad \text{for } n > 1.$$

It suffices to verify that

$$1 - \mathbb{P}[|\mathcal{C}_s| = 1] \gtrsim q \quad \text{and} \quad \mathbb{P}[|\mathcal{C}_s| = 1] \gtrsim 1.$$

Let κ be the arc chain of π . Due to the reversal symmetry (10), we may assume that $n - s \geq 1$. Lemma 3.4 implies that

$$\begin{aligned} \mathbb{P}[|\mathcal{C}_s| = 1] &= \mathbb{P}[\pi_s = s] = \mathbb{E} \left[\frac{q^{\kappa_{s-1}} - q^{\kappa_{s-1}+1}}{1 - q^{n-s+1}} \cdot \frac{q^{\kappa_{s-1}} - q^{n-s+1}}{1 - q^{n-s+1}} \right] \\ &\leq \frac{1 - q}{1 - q^2} = \frac{1}{1 + q}. \end{aligned}$$

Hence $1 - \mathbb{P}[|\mathcal{C}_s| = 1] \geq \frac{q}{1+q} \gtrsim q$. Now we shall verify that $\mathbb{P}[|\mathcal{C}_s| = 1] \gtrsim 1$. Observe that

$$\begin{aligned} \mathbb{P}[\pi_s = s] &\geq \mathbb{P}[\pi_s = s | \kappa_{s-1} = 0] \cdot \mathbb{P}[\kappa_{s-1} = 0] \geq (1 - q) \cdot \mathbb{P}[\kappa_{s-1} = 0] \\ &\gtrsim \mathbb{P}[\kappa_{s-1} = 0] \end{aligned}$$

by Lemma 3.4 and our assumption that $q \leq \frac{1}{8}$. The fact that $\mathbb{P}[\kappa_{s-1} = 0] \gtrsim 1$ follows from (30) and Lemma 3.13.

Case 3: We prove the upper bound for the variance of cycle length. Since $|\mathcal{C}_s| \geq 1$, it follows that

$$(68) \quad \begin{aligned} \text{Var } |\mathcal{C}_s| &\leq \mathbb{E}[(|\mathcal{C}_s| - 1)^2] \leq \mathbb{E}[\text{diam}(\mathcal{C}_s)^2] \\ &\leq 2\mathbb{E}[(\max(\mathcal{C}_s) - s)^2 + (s - \min(\mathcal{C}_s))^2]. \end{aligned}$$

Proposition 4.7, via the reversal symmetry (10), states that

$$\begin{aligned} \mathbb{E}[(\max(\mathcal{C}_s) - s)^2] &\lesssim \min\{q\xi^4, (n - s)^2\} \quad \text{and} \\ \mathbb{E}[(s - \min(\mathcal{C}_s))^2] &\lesssim \min\{q\xi^4, (s - 1)^2\}, \end{aligned}$$

which yields the upper bound via (68).

4.5. Poisson–Dirichlet law. In this section, we prove Theorem 1.2. We need to prove two facts: that the normalized length $\frac{1}{n}|\mathcal{C}_{s_n}|$ converges in distribution to the uniform distribution on $[0, 1]$ for any sequence (s_n) with $s_n \in [n]$ and that the distribution of the sorted and normalized cycle lengths converges to the Poisson–Dirichlet law. The proofs of these two facts are similar and we shall focus on the proof for the Poisson–Dirichlet law. At the end of the section, we point out the needed modifications to obtain the limiting distribution of $\frac{1}{n}|\mathcal{C}_{s_n}|$.

The Poisson–Dirichlet law in a space of multisets of reals. Denote by \mathcal{D} the space of sorted sequences $(\alpha_i), i \geq 1, \alpha_1 \geq \alpha_2 \geq \dots$ of nonnegative reals with finite sum. The space \mathcal{D} is endowed with the product topology, the topology inherited from $\mathbb{R}^{\mathbb{N}}$. We shall consider also the ℓ^2 metric on sequences in \mathcal{D} and take note of the fact that convergence in the ℓ^2 metric implies convergence in the product topology.

The Poisson–Dirichlet law with parameter one, denoted by \mathcal{PD} , is a distribution on \mathcal{D} , supported on sequences with sum 1. We will not need the precise definition of \mathcal{PD} , instead relying only on its relation with uniform random permutations, and the reader is referred, for example, to [12] for further background. Specifically, we shall use that if σ is a uniformly random permutation in \mathbb{S}_k and $\ell_1 \geq \ell_2 \geq \dots$ are the sorted lengths of cycles in σ then

$$(69) \quad \frac{1}{k}(\ell_1, \ell_2, \dots) \xrightarrow{d} \mathcal{PD} \quad \text{as } k \rightarrow \infty,$$

where a finite sequence is viewed as an element of \mathcal{D} by adding to it a trailing sequence of zeros.

It is convenient to work with an alternative, equivalent, description of the space \mathcal{D} . A sequence in \mathcal{D} may be equivalently described by a multiset of nonnegative reals with finite sum (summing elements according to their multiplicities),

in which the multiplicity of each number is the number of its occurrences in the sequence. Note that 0 is the only number possibly having an infinite multiplicity in this representation. We denote the space of such multisets by \mathcal{M} and make the identification of \mathcal{D} and \mathcal{M} in the sequel, putting the induced topology on \mathcal{M} , that is, the push-forward of the product topology via the identification map. We denote by d the metric on \mathcal{M} obtained as the push-forward of the ℓ^2 metric on \mathcal{D} . The metric d has the following explicit expression: given two multisets $X, Y \in \mathcal{M}$,

$$(70) \quad d(X, Y) := \min_{\varphi} \sqrt{\sum_{x \in X} (\varphi(x) - x)^2},$$

where the minimum is taken over all bijections $\varphi: X \rightarrow Y$ and it is understood that bijections may assign different images to multiple occurrences of the same element. To see that d coincides with the push-forward of the ℓ^2 metric on \mathcal{D} , we note that the minimum in (70) is obtained by taking φ to be a monotone nondecreasing mapping.

We will also consider multisets of finite cardinality of nonnegative reals as members of \mathcal{M} by adding an infinite amount of zeros to the multiset.

Convergence criterion. Here, we state a convergence criterion to the Poisson–Dirichlet distribution which generalizes (69). We start with some definitions.

For a permutation $\sigma: S \rightarrow S$ and a weight function $w: S \rightarrow \mathbb{R}$, where S is a nonempty finite set, we define the weighted length function $\mathcal{L}(\sigma, w)$ to be the multiset

$$(71) \quad \mathcal{L}(\sigma, w) := \left\{ \sum_{i \in \mathcal{C}^1} w_i, \dots, \sum_{i \in \mathcal{C}^m} w_i \right\},$$

where $\mathcal{C}^1, \dots, \mathcal{C}^m$ is a decomposition of S into orbits of σ . For instance, $\mathcal{L}(\sigma, \mathbb{1})$, where $\mathbb{1}$ denotes the constant 1 function, is the multiset of cycle lengths of σ .

The convergence in (69) may be stated equivalently as follows: if σ is uniform in \mathbb{S}_k then $\mathcal{L}(\sigma, \frac{1}{k}\mathbb{1}) \rightarrow \mathcal{PD}$ as $k \rightarrow \infty$. The following proposition allows us to generalize this fact by “adding weights”.

PROPOSITION 4.10. *Let S be a nonempty finite set, let $w: S \rightarrow [0, \infty)$ satisfy $\sum w_s \leq 1$ and let $\sigma: S \rightarrow S$ be a uniformly random permutation. Then*

$$\mathbb{E} \left[d^2 \left(\mathcal{L} \left(\sigma, \frac{1}{|S|} \mathbb{1} \right), \mathcal{L}(\sigma, w) \right) \right] \leq \frac{1}{2} \left(\left(1 - \sum w_s \right)^2 + \sum \left(\frac{1}{|S|} - w_s \right)^2 \right).$$

In particular, if $w = w(k): [k] \rightarrow [0, \infty)$ is a sequence of tuples satisfying $\sum_{i=1}^k w_i \leq 1$ and the limit relations $\sum_{i=1}^k w_i \rightarrow 1$ and $\max_{i \in [k]} w_i \rightarrow 0$ as $k \rightarrow \infty$, then $\mathcal{L}(\sigma, \frac{1}{k}\mathbb{1})$ and $\mathcal{L}(\sigma, w)$ share the same limit distribution, which is the \mathcal{PD} distribution by (69).

PROOF OF PROPOSITION 4.10. Using the notation of (71), the following inequality follows from (70):

$$\begin{aligned}
 (72) \quad d^2\left(\mathcal{L}\left(\sigma, \frac{1}{|S|}\mathbb{1}\right), \mathcal{L}(\sigma, w)\right) &\leq \sum_r \left(\sum_{s \in \mathcal{C}^r} \left(\frac{1}{|S|} - w_s\right)\right)^2 \\
 &= \sum_{s \in S} \sum_{t \in S} \left(\frac{1}{|S|} - w_s\right) \left(\frac{1}{|S|} - w_t\right) \cdot \mathbb{1}\{t \in \mathcal{C}_s(\sigma)\}.
 \end{aligned}$$

For each $s \in S$, as $\mathbb{P}[t \in \mathcal{C}_s(\sigma)] = \frac{1}{2}\mathbb{1}_{s=t} + \frac{1}{2}$,

$$\begin{aligned}
 (73) \quad \mathbb{E}\left[\sum_{t \in S} \left(\frac{1}{|S|} - w_s\right) \left(\frac{1}{|S|} - w_t\right) \cdot \mathbb{1}\{t \in \mathcal{C}_s(\sigma)\}\right] \\
 = \frac{1}{2} \left(\left(\frac{1}{|S|} - w_s\right) \cdot \left(1 - \sum_{t \in S} w_t\right) + \left(\frac{1}{|S|} - w_s\right)^2 \right).
 \end{aligned}$$

The proposition follows by taking expectation in (72) and substituting (73). \square

Proof of the Poisson–Dirichlet law. Let $q = q_n$ satisfy

$$(74) \quad (1 - q)^2 \cdot n \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Let $m = m_n \in \mathbb{N}$ be a sequence that converges to infinity sufficiently slowly so that

$$(75) \quad \frac{m}{n} \rightarrow 0 \quad \text{and} \quad q^{m^2} \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

For instance, we may take $m = \lfloor n^{1/4} \rfloor$ or $m = \lceil \log(n + 1) \rceil$. Let $\pi \sim \mu_{n,q}$. We shall analyze π conditioned on $\Gamma_\pi \cap U$, where

$$U := \{x \leq n - m \text{ or } y > n - m\}.$$

The plan is to show that the lengths of the long cycles can be coupled closely with the lengths of cycles in a uniform permutation. To do it, we first show that, despite the fact that $\Gamma_\pi \cap U$ is almost the whole graph of π , with high probability, it provides little to no information on the lengths of the long cycles of π . The lengths of the long cycles, given $\Gamma_\pi \cap U$, are mostly decided by $\Gamma_\pi \setminus U$, the remaining portion of the graph. We shall then conclude by utilizing the fact that the relative order [see (39)] of $\Gamma_\pi \setminus U$ is very close to a uniformly distributed permutation.

Our proof requires us to define permutations over a finite random set \mathcal{O}_U and analyze the multiset $\mathcal{L}(\cdot, w)$ of such permutations for various weight tuples w . The set \mathcal{O}_U may be empty, though this case does not impact on our analysis as its probability tends to 0. To avoid treating it in a special manner, as much as possible, we define the set of permutations over the empty set to consist of a single element denoted id_\emptyset . This element has no cycles and, in particular, $\mathcal{L}(\text{id}_\emptyset, w) = \emptyset$ for all w .

Let us begin the proof. We introduce additional definitions to discuss arcs, which are determined by $\Gamma_\pi \cap U$. We say that an arc α of π belongs to U if one may order its elements $\alpha = \{a_1, \dots, a_{|\alpha|}\}$ so that $a_{i+1} = \pi(a_i)$ and $(a_i, a_{i+1}) \in U$ for all $1 \leq i < |\alpha|$. If, in addition, $a_1 = \pi(a_{|\alpha|})$ and $(a_{|\alpha|}, a_1) \in U$ we say that α is relatively closed and otherwise deem it relatively open. When the arc is relatively open, then the above ordering is unique, in which case we call the elements a_1 and $a_{|\alpha|}$ the tail and head of the arc α , respectively, and denote them by $\text{tail}(\alpha)$ and $\text{head}(\alpha)$. We say that an arc α is U -maximal if α belongs to U and there are no other arcs that belong to U and contain α . Let \mathcal{A}_U be the set of U -maximal arcs and let \mathcal{O}_U be the set of relatively open U -maximal arcs. One should note that \mathcal{A}_U and \mathcal{O}_U are determined by $\Gamma_\pi \cap U$ and that $\sum_{\alpha \in \mathcal{A}_U} |\alpha| = n$. The last equality follows from the fact that every element of $[n]$ belongs to exactly one of the arcs of \mathcal{A}_U (possibly to an arc containing only this element).

We proceed to discuss the way that the cycles of π are formed from the arcs in \mathcal{A}_U and the portion of the graph $\Gamma_\pi \setminus U$. Each point $(s, t) \in \Gamma_\pi \setminus U$ satisfies $s = \text{head}(\alpha)$ and $t = \text{tail}(\beta)$ for some $\alpha, \beta \in \mathcal{O}_U$. Conversely, for each $\alpha \in \mathcal{O}_U$ there exist points $(s, t), (s', t') \in \Gamma_\pi \setminus U$ satisfying $s = \text{head}(\alpha)$ and $t' = \text{tail}(\alpha)$. Thus we may define a permutation $\tau : \mathcal{O}_U \rightarrow \mathcal{O}_U$ by setting

$$\tau(\alpha) = \beta \quad \text{if and only if} \quad (\text{head}(\alpha), \text{tail}(\beta)) \in \Gamma_\pi.$$

It is straightforward to check that each cycle of π is then either a cycle in $\mathcal{A}_U \setminus \mathcal{O}_U$, or formed by merging the open arcs in \mathcal{O}_U which are in the same orbit of τ . In particular,

$$\mathcal{L}\left(\pi, \frac{1}{n}\mathbb{1}\right) = \left\{ \frac{1}{n}|\alpha| \mid \alpha \in \mathcal{A}_U \setminus \mathcal{O}_U \right\} \cup \left\{ \frac{1}{n} \sum_{\alpha \in \mathcal{C}} |\alpha| \mid \mathcal{C} \text{ is a cycle of } \tau \right\},$$

where the equality and union are in the sense of multisets. Write $|\cdot|$ for the length map on arcs, $\alpha \mapsto |\alpha|$. Recalling (70), we conclude that

$$\begin{aligned} (76) \quad & d\left(\mathcal{L}\left(\pi, \frac{1}{n}\mathbb{1}\right), \mathcal{L}\left(\tau, \frac{1}{n}|\cdot|\right)\right) \\ & \leq \sqrt{\sum_{\alpha \in \mathcal{A}_U \setminus \mathcal{O}_U} \frac{|\alpha|^2}{n^2}} \leq \sum_{\alpha \in \mathcal{A}_U \setminus \mathcal{O}_U} \frac{|\alpha|}{n} = 1 - \sum_{\alpha \in \mathcal{O}_U} \frac{|\alpha|}{n}. \end{aligned}$$

The Poisson–Dirichlet law is a consequence of this inequality and the following lemma, which states the properties of the Mallows model that we require.

LEMMA 4.11. *Let $\pi \sim \mu_{n,q}$ and suppose $n \rightarrow \infty$ with (74) and (75) holding. Then:*

$$(77) \quad (i) \quad \frac{1}{n} \max_{\alpha \in \mathcal{O}_U} |\alpha| \rightarrow 0 \quad \text{and} \quad \frac{1}{n} \sum_{\alpha \in \mathcal{O}_U} |\alpha| \rightarrow 1 \quad \text{in probability.}$$

(ii) *There exists a coupling of τ and a permutation $\sigma : \mathcal{O}_U \rightarrow \mathcal{O}_U$ such that*

$$\mathbb{P}[\sigma \neq \tau] \rightarrow 0 \quad \text{as } n \rightarrow \infty$$

and, conditioned on $\Gamma_\pi \cap U$, σ has the uniform distribution on permutations of \mathcal{O}_U .

To obtain the Poisson–Dirichlet limit law, let σ be the random permutation given by part (ii) of Lemma 4.11. Observe that

$$\begin{aligned} & d\left(\mathcal{L}\left(\pi, \frac{1}{n}\mathbb{1}\right), \mathcal{L}\left(\sigma, \frac{1}{|\mathcal{O}_U|}\mathbb{1}\right)\right) \\ & \leq \underbrace{d\left(\mathcal{L}\left(\pi, \frac{1}{n}\mathbb{1}\right), \mathcal{L}\left(\tau, \frac{1}{n}|\cdot|\right)\right)}_{=:I} + \underbrace{d\left(\mathcal{L}\left(\tau, \frac{1}{n}|\cdot|\right), \mathcal{L}\left(\sigma, \frac{1}{n}|\cdot|\right)\right)}_{=:II} \\ & \quad + \underbrace{d\left(\mathcal{L}\left(\sigma, \frac{1}{n}|\cdot|\right), \mathcal{L}\left(\sigma, \frac{1}{|\mathcal{O}_U|}\mathbb{1}\right)\right)}_{=:III}. \end{aligned}$$

We estimate the expectation of each of the last three summands separately. By (76) and part (i) of Lemma 4.11,

$$\mathbb{E}[I] \leq \mathbb{E}\left[1 - \sum_{a \in \mathcal{O}_U} \frac{|a|}{n}\right] \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Note that as $d(\mathcal{L}(\alpha, w), \mathcal{L}(\beta, w)) \leq 2 \sum_i |w_i|$ for any two permutations α, β on a finite set S and weight tuple $w : S \rightarrow \mathbb{R}$, we have $II \leq 2 \cdot \mathbb{1}_{\sigma \neq \tau}$. Thus, by part (ii) of Lemma 4.11,

$$\mathbb{E}[II] \leq 2 \cdot \mathbb{P}[\sigma \neq \tau] \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Lastly, $III = 0$ when $\mathcal{O}_U = \emptyset$ and by Proposition 4.10,

$$\begin{aligned} & \mathbb{E}[(III)^2 | \Gamma_\pi \cap U] \\ & \leq \frac{1}{2} \left(\left(1 - \frac{1}{n} \sum_{a \in \mathcal{O}_U} |a|\right)^2 + \sum_{a \in \mathcal{O}_U} \left(\frac{1}{|\mathcal{O}_U|} - \frac{|a|}{n}\right)^2 \right) \quad \text{on } \{\mathcal{O}_U \neq \emptyset\}. \end{aligned}$$

Thus, applying part (i) of Lemma 4.11 and observing that it implies, in particular, that $|\mathcal{O}_U| \rightarrow \infty$ in probability as $n \rightarrow \infty$, we obtain

$$\mathbb{E}[III] \leq \sqrt{\mathbb{E}[(III)^2]} \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

We conclude that

$$\mathbb{E}\left[d\left(\mathcal{L}\left(\pi, \frac{1}{n}\mathbb{1}\right), \mathcal{L}\left(\sigma, \frac{1}{|\mathcal{O}_U|}\mathbb{1}\right)\right)\right] \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

The limiting distribution of $\mathcal{L}(\sigma, \frac{1}{|\mathcal{O}_U|}\mathbb{1})$ is the \mathcal{PD} distribution as, conditioned on $\Gamma_\pi \cap U$, σ is a uniform permutation on \mathcal{O}_U by part (ii) of Lemma 4.11, using (69) and relying again on the fact that $|\mathcal{O}_U| \rightarrow \infty$ in probability as $n \rightarrow \infty$. Thus $\mathcal{L}(\pi, \frac{1}{n}\mathbb{1})$ converges also to the \mathcal{PD} distribution, as we wanted to prove.

Proof of part (ii) of Lemma 4.11. This is a corollary of Lemma 3.15. On the event $\{\mathcal{O}_U = \emptyset\}$, we simply set $\sigma = \tau$. On the complementary event $\{\mathcal{O}_U \neq \emptyset\}$ we do as follows. Set $k := |\mathcal{O}_U|$. Let $\tilde{\tau}$ be the relative order [see (39)] of $\Gamma_\pi \setminus U$. By our construction,

$$\tau = \alpha \circ \tilde{\tau} \circ \beta$$

for two bijections, $\alpha : [k] \rightarrow \mathcal{O}_U$ and $\beta : \mathcal{O}_U \rightarrow [k]$, which are determined by $\Gamma_\pi \cap U$. Explicitly, this follows by viewing $\tau : \mathcal{O}_U \rightarrow \mathcal{O}_U$ as the composition of 5 maps:

$$\mathcal{O}_U \xrightarrow{\text{head}} \text{head}(\mathcal{O}_U) \xrightarrow{\text{monotone}} [k] \xrightarrow{\tilde{\tau}} [k] \xrightarrow{\text{monotone}} \text{tail}(\mathcal{O}_U) \xrightarrow{\text{tail}^{-1}} \mathcal{O}_U,$$

where $A \xrightarrow{\text{monotone}} B$ stands for the unique monotone increasing bijection from A to B , provided that A and B are subsets of \mathbb{N} of the same size.

It thus suffices to couple $\tilde{\tau}$ with a permutation $\tilde{\sigma} : [k] \rightarrow [k]$ in a way that

$$\mathbb{P}[\tilde{\tau} \neq \tilde{\sigma} | k \geq 1] \rightarrow 0 \quad \text{as } n \rightarrow \infty$$

and, conditioned on $\Gamma_\pi \cap U$, $\tilde{\sigma}$ has the uniform distribution on \mathbb{S}_k , as we may then take $\sigma := \alpha \circ \tilde{\sigma} \circ \beta$.

By Lemma 3.15, conditioned on $\Gamma_\pi \cap U$, we have that $\tilde{\tau} \sim \mu_{k,q}$. Hence the following claim suffices to finish the proof, using our assumption (75) and the fact that $k \leq m$.

CLAIM 4.12. *Let $k \in \mathbb{N}$, $0 < q \leq 1$ and $\rho \sim \mu_{k,q}$. Then ρ may be coupled with a uniform random permutation λ in \mathbb{S}_k such that*

$$\mathbb{P}[\rho \neq \lambda] \leq 1 - q^{k^2}.$$

PROOF. We recall (see, e.g., [19], Proposition 4.7) that the total variation distance of two probability distributions μ and ν on a finite set S is given by

$$\text{TV}(\mu, \nu) := \frac{1}{2} \sum_{s \in S} |\mu[s] - \nu[s]| = 1 - \sum_{s \in S} \min\{\mu[s], \nu[s]\},$$

and that there exists a coupling of the two distributions, that is, random variables X, Y with X distributed μ and Y distributed ν , so that $\mathbb{P}[X \neq Y] = \text{TV}(\mu, \nu)$ (and, moreover, there is no coupling achieving a smaller value for $\mathbb{P}[X \neq Y]$). Thus it suffices to show that the total variation distance of $\mu_{k,q}$ and the uniform distribution on \mathbb{S}_k is at most $1 - q^{k^2}$.

Let $Z_{k,q}$ be as in the definition (1) of the Mallows permutation and note that $Z_{k,q} \leq k!$. Since $\text{inv}(\sigma) \leq k^2$ for all $\sigma \in \mathbb{S}_k$, we obtain from (1) that $\mu_{k,q}[\sigma] \geq \frac{q^{k^2}}{k!}$ for all $\sigma \in \mathbb{S}_k$. Thus the required total variation distance is at most $1 - \sum_{\sigma \in \mathbb{S}_k} \min\{\frac{q^{k^2}}{k!}, \frac{1}{k!}\} = 1 - q^{k^2}$. \square

Proof of part (i) of Lemma 4.11. The claim is derived from the following proposition regarding diagonal arcs.

PROPOSITION 4.13. *For $1 \leq s \leq r \leq n$, one has*

$$(78) \quad \mathbb{E}|\{1 \leq i \leq r \mid \text{arc}_i^r \text{ is closed}\}| \lesssim 1 + (1 - q)^2 \cdot n^2 + \frac{1}{n - r + 1} \cdot n,$$

$$(79) \quad \mathbb{E}|\text{arc}_s^r| \lesssim 1 + (1 - q)^2 \cdot n^2 + \frac{1}{n - r + 1} \cdot n.$$

Since each open arc of \mathcal{A}_{n-m} extends to an arc of \mathcal{O}_U , inequality (78), applied with $r = n - m$, implies the second limit in (77) by our assumptions that m tends to infinity with n and that (74) and (75) hold.

To derive the first limit in (77), we use the following general claim.

CLAIM 4.14. *Let I be a nonempty finite set. Let A_1, \dots, A_k , where $k \in \mathbb{N}$ is random, be pairwise disjoint random subsets of I with union I . For $i \in I$, let ℓ_i be the size of the A_j to which i belongs. Then*

$$\left(\frac{1}{|I|} \mathbb{E} \left[\max_{i \in I} \ell_i \right] \right)^3 \lesssim \frac{1}{|I|} \max_{i \in I} \mathbb{E}[\ell_i].$$

Now, inequality (79) yields that $\frac{1}{n} \mathbb{E}|\text{arc}_s^{n-m}| \rightarrow 0$ uniformly over all $s \in [n - m]$ by our assumptions that m tends to infinity with n and that (74) and (75) hold. Claim 4.14, applied with $I = [n - m]$ and A_1, \dots, A_k being the diagonal arcs of \mathcal{A}_{n-m} , allows one to deduce that in fact

$$(80) \quad \frac{1}{n} \mathbb{E} \left[\max_{a \in \mathcal{A}_{n-m}} |a| \right] = \frac{1}{n} \mathbb{E} \left[\max_{s \in [n-m]} |\text{arc}_s^{n-m}| \right] \rightarrow 0.$$

Since restricting an arc of \mathcal{O}_U to $[n - m]$ yields an open arc of \mathcal{A}_{n-m} which is shorter by at most m elements, the first limit in (77) follows from (80) using our assumption that $\frac{m}{n} \rightarrow 0$.

PROOF OF CLAIM 4.14. Set $L = \max_{i \in I} \ell_i$ and $\alpha := \frac{1}{|I|} \mathbb{E}[L]$. By Markov's inequality applied to $|I| - L$, using that $0 \leq L \leq |I|$,

$$(81) \quad \mathbb{P} \left[L \geq \frac{1}{2} \alpha |I| \right] \gtrsim \alpha.$$

The sum of the ℓ_i satisfies

$$(82) \quad \sum_{i \in I} \ell_i \geq \mathbb{1} \left\{ L \geq \frac{1}{2} \alpha |I| \right\} \cdot \sum_{i \in I} \ell_i \geq \frac{\alpha^2 |I|^2}{4} \cdot \mathbb{1} \left\{ L \geq \frac{1}{2} \alpha |I| \right\}.$$

The second inequality is due to the fact that on the event $\{L \geq \frac{1}{2} \alpha |I|\}$ there are at least $\frac{1}{2} \alpha |I|$ values of $i \in I$ for which $\ell_i \geq \frac{1}{2} \alpha |I|$. By taking expectation in (82) and substituting (81), we obtain

$$|I| \cdot \max_{i \in I} \mathbb{E}[\ell_i] \geq \sum_{i \in I} \mathbb{E}[\ell_i] \geq \frac{\alpha^2 |I|^2}{4} \mathbb{P} \left[L \geq \frac{1}{2} \alpha |I| \right] \gtrsim \alpha^3 |I|^2. \quad \square$$

PROOF OF PROPOSITION 4.13. Fix s, r satisfying $1 \leq s \leq r \leq n$. Define

$$M_t := |\{1 \leq i \leq t \mid \text{arc}_i^t \text{ is closed}\}| \quad \text{and for } t \geq s \text{ let } N_t := |\text{arc}_s^t|.$$

Our intention is to provide bounds for $\mathbb{E}[N_s]$, $\mathbb{E}[N_{t+1} - N_t]$ and $\mathbb{E}[M_{t+1} - M_t]$, and to derive the proposition from these bounds. Observe that

$$N_s = 1 + \sum_{\mathbf{a} \in \mathcal{O}_{s-1}} |\mathbf{a}| \cdot \mathbb{1} \left\{ \begin{array}{l} \pi_s = \text{tail}(\mathbf{a}) \text{ or} \\ \pi_s^{-1} = \text{head}(\mathbf{a}) \end{array} \right\},$$

as arc_s^s consists of s and up to two arcs of \mathcal{O}_{s-1} that have merged with $\{s\}$ via either their head or their tail. In light of this equality, Lemma 4.1 and the fact that $\sum_{\mathbf{a} \in \mathcal{O}_{s-1}} |\mathbf{a}| \leq n$ imply the following bound:

$$(83) \quad \mathbb{E}[N_s] \lesssim 1 + n \cdot \frac{1 - q}{1 - q^{n-s+1}}.$$

To bound $N_{t+1} - N_t$ and $M_{t+1} - M_t$, we define the events

$$\mathcal{M}_{t+1}(\mathbf{a}, \mathbf{b}) := \left\{ \begin{array}{l} \pi_{t+1} = \text{tail}(\mathbf{a}) \\ \pi_{t+1}^{-1} = \text{head}(\mathbf{b}) \end{array} \right\} \cup \left\{ \begin{array}{l} \pi_{t+1} = \text{tail}(\mathbf{b}) \\ \pi_{t+1}^{-1} = \text{head}(\mathbf{a}) \end{array} \right\}, \quad \mathbf{a}, \mathbf{b} \in \mathcal{O}_t,$$

denoting the merging of \mathbf{a} and \mathbf{b} when the two arcs are distinct and the closure of \mathbf{a} when they are equal. Note that $\mathbb{P}_t[\mathcal{M}_{t+1}(\mathbf{a}, \mathbf{b})] \leq (\frac{1-q}{1-q^{n-t}})^2$ for all $\mathbf{a}, \mathbf{b} \in \mathcal{O}_t$, as Lemma 4.1 implies.

Observe that for $t \geq s$ one has $N_{t+1} - N_t = 0$ if arc_s^t is closed and otherwise one has

$$\begin{aligned} N_{t+1} - N_t &= \mathbb{1} \left\{ \begin{array}{l} \pi_{t+1} = \text{tail}(\text{arc}_s^t) \text{ or} \\ \pi_{t+1}^{-1} = \text{head}(\text{arc}_s^t) \end{array} \right\} \\ &\quad + \sum_{\mathbf{a} \in \mathcal{O}_t \setminus \{\text{arc}_s^t\}} |\mathbf{a}| \cdot \mathbb{1}_{\mathcal{M}_{t+1}(\mathbf{a}, \text{arc}_s^t)} \quad \text{on } \{\text{arc}_s^t \in \mathcal{O}_t\}, \end{aligned}$$

as $N_{t+1} - N_t > 0$ only when $\{t + 1\}$ has merged with arc_s^t , in which case $N_{t+1} - N_t = |\alpha| + 1$ if $\{t + 1\}$ has also merged with another arc $\alpha \in \mathcal{O}_t \setminus \{\text{arc}_s^t\}$, and otherwise $N_{t+1} - N_t = 1$.

For the difference $M_{t+1} - M_t$, observe that

$$\begin{aligned} M_{t+1} - M_t &= \mathbb{1}\{t + 1 = \pi_{t+1}\} + \sum_{\alpha \in \mathcal{O}_t} (|\alpha| + 1) \cdot \mathbb{1}_{\mathcal{M}_{t+1}(\alpha, \alpha)} \\ &= \mathbb{1}\{t + 1 = \max(\mathcal{C}_{t+1})\} + \sum_{\alpha \in \mathcal{O}_t} |\alpha| \cdot \mathbb{1}_{\mathcal{M}_{t+1}(\alpha, \alpha)}, \end{aligned}$$

as $M_{t+1} - M_t > 0$ only when $t + 1$ closes a cycle, in which case $M_{t+1} - M_t = |\alpha| + 1$ when $t + 1$ closes the arc $\alpha \in \mathcal{O}_t$ and $M_{t+1} - M_t = 1$ when $t + 1$ forms a fixed point.

Note that $\sum_{\alpha \in \mathcal{O}_t} |\alpha| \leq n$ and apply Lemma 4.1 and Proposition 4.2 to the above formulas for $N_{t+1} - N_t$ and $M_{t+1} - M_t$ to obtain

$$(84) \quad \mathbb{E}[M_{t+1} - M_t] \lesssim \frac{1 - q}{1 - q^{n-t}} + n \cdot \frac{(1 - q)^2}{(1 - q^{n-t})^2} \quad \text{for } t \geq 0,$$

$$(85) \quad \mathbb{E}[N_{t+1} - N_t] \lesssim \frac{1 - q}{1 - q^{n-t}} + n \cdot \frac{(1 - q)^2}{(1 - q^{n-t})^2} \quad \text{for } t \geq s.$$

By using the bounds (83), (84) and (85) one may show that both $\mathbb{E}[N_r]$ and $\mathbb{E}[M_r]$ are bounded, up to multiplication by a positive absolute constant, by

$$1 + n \cdot \frac{1 - q}{1 - q^{n-r+1}} + \sum_{t=0}^{r-1} \frac{1 - q}{1 - q^{n-t}} + n \cdot \sum_{t=0}^{r-1} \frac{(1 - q)^2}{(1 - q^{n-t})^2}.$$

The bounds (78) and (79) follow by using (54). \square

Limiting distribution of $\frac{1}{n}|\mathcal{C}_{s_n}|$. Our proof of the fact that $\frac{1}{n}|\mathcal{C}_{s_n}|$ converges in distribution to $U[0, 1]$, the uniform distribution on $[0, 1]$, is very similar to our proof of the Poisson–Dirichlet limit law. Therefore, let us only elaborate on the main differences.

The proof is based on the following simple fact: For $k \in \mathbb{N}$, let $\rho \in \mathbb{S}_k$ be a uniformly random permutation and let $i_k \in [k]$ be arbitrary. Then

$$(86) \quad \frac{1}{k}|\mathcal{C}_{i_k}(\rho)| \rightarrow U[0, 1] \quad \text{in distribution as } k \rightarrow \infty.$$

We use the same notation as in the proof of the Poisson–Dirichlet limit law, for example, σ is the random permutation over \mathcal{O}_U from part (ii) of Lemma 4.11. For $\alpha \in \mathcal{O}_U$, we denote by $\mathcal{C}_\alpha(\sigma)$ the orbit of σ , and similarly with τ .

Let α be the arc of \mathcal{A}_U that contains s_n . We first claim that

$$(87) \quad \mathbb{P}[\alpha \in \mathcal{O}_U] \rightarrow 1 \quad \text{as } n \rightarrow \infty.$$

Indeed, this is a consequence of Corollary 4.4, used with the reversal symmetry (10) if $s_n > n - m$, and our assumptions that m tends to infinity with n and that (74) and (75) hold.

In addition, we recall that $|\mathcal{O}_U| \rightarrow \infty$ in probability as $n \rightarrow \infty$ as a consequence of (77). Now, the limits (86) and (87) and the fact that, given $|\mathcal{O}_U|$, the distribution of σ is uniform imply that

$$\frac{1}{|\mathcal{O}_U|} |\mathcal{C}_a(\sigma)| \rightarrow U[0, 1] \quad \text{in distribution as } n \rightarrow \infty,$$

where it is understood that in the case when $a \notin \mathcal{O}_U$ we set $\frac{1}{|\mathcal{O}_U|} |\mathcal{C}_a(\sigma)| := 0$. It thus suffices to show that

$$\left| \frac{1}{n} |\mathcal{C}_{s_n}(\pi)| - \frac{1}{|\mathcal{O}_U|} |\mathcal{C}_a(\sigma)| \right| \rightarrow 0 \quad \text{in probability as } n \rightarrow \infty.$$

To see this, we write, similarly as in the proof of the Poisson–Dirichlet limit law, interpreting $\mathcal{C}_a(\tau)$ and $\mathcal{C}_a(\sigma)$ as empty when $a \notin \mathcal{O}_U$,

$$\begin{aligned} & \left| \frac{1}{n} |\mathcal{C}_{s_n}(\pi)| - \frac{1}{|\mathcal{O}_U|} |\mathcal{C}_a(\sigma)| \right| \\ & \leq \left| \frac{1}{n} |\mathcal{C}_{s_n}(\pi)| - \sum_{b \in \mathcal{C}_a(\tau)} \frac{|b|}{n} \right| + \left| \sum_{b \in \mathcal{C}_a(\tau)} \frac{|b|}{n} - \sum_{b \in \mathcal{C}_a(\sigma)} \frac{|b|}{n} \right| \\ & \quad + \left| \sum_{b \in \mathcal{C}_a(\sigma)} \frac{|b|}{n} - \frac{1}{|\mathcal{O}_U|} |\mathcal{C}_a(\sigma)| \right|. \end{aligned}$$

The first of the terms on the right-hand side is small, in probability, due to (87). The second term is small by part (ii) of Lemma 4.11. The last term can be bounded in a similar manner as in the proof of Proposition 4.10 and shown to be small by part (i) of Lemma 4.11.

5. Discussion and open questions. In this work, we study the Mallows model for random permutations, providing estimates for the typical length and diameter of cycles. We observe that macroscopic cycles emerge in the parameter range $\frac{1}{(1-q)^2} \gg n$. In this regime, we prove further that the joint distribution of the lengths of long cycles in the permutation converges to the Poisson–Dirichlet distribution. In this section, we discuss several further questions on the Mallows model as well as questions pertaining to other related models of random permutations.

The limiting distributions of the cycle length and cycle diameter. Let π have the Mallows distribution with parameters n and q . Recall that \mathcal{C}_s stands for the cycle in π containing the point $1 \leq s \leq n$, so that $|\mathcal{C}_s|$ and $\max(\mathcal{C}_s) - \min(\mathcal{C}_s)$ are the length and diameter of \mathcal{C}_s , respectively. What can be said about the limiting distributions of these quantities when $q \rightarrow 1$ and $n \rightarrow \infty$? To avoid boundary effects, we restrict to the case that there is some $\alpha \in (0, 1)$ for which $s = s_n$ satisfies $\frac{s}{n} \rightarrow \alpha$. We consider three cases.




FIG. 6. Distribution of the length of the cycle containing a uniform random point. Obtained empirically with 1,000,000 samples.

Macroscopic cycles. Suppose that $n(1 - q)^2 \rightarrow 0$. In this regime, as shown in Theorem 1.2, the normalized cycle length $\frac{1}{n}|\mathcal{C}_s|$ converges in distribution to the uniform distribution on $[0, 1]$. Figures 2(b) and 6(b) suggest that, in fact, a stronger convergence takes place. For any sequence $1 \leq k_n \leq n$ bounded away from 1 and n in a suitable manner, one has $\mathbb{P}[|\mathcal{C}_s| = k_n] \cdot n \rightarrow 1$.

Using Corollary 4.4, it may additionally be shown that the cycle spans the full interval, in the sense of the following convergence in distribution:

$$\frac{\min(\mathcal{C}_s)}{n} \xrightarrow{d} 0 \quad \text{and} \quad \frac{\max(\mathcal{C}_s)}{n} \xrightarrow{d} 1.$$

Microscopic cycles. Suppose that $n(1 - q)^2 \rightarrow \infty$. It appears from simulations (see Figures 2 and 6) that the limiting distribution of the normalized cycle length $(1 - q)^2|\mathcal{C}_s|$ exists in this regime, but it is unclear what its form is. This limiting distribution, if it indeed exists, cannot be concentrated on a single point due to the lower bound on the variance of $(1 - q)^2|\mathcal{C}_s|$ given in Theorem 1.3, used together with the bounds in Proposition 4.6 and Theorem 3.6.

The normalized cycle diameter $(1 - q)^2(\max(\mathcal{C}_s) - \min(\mathcal{C}_s))$ seems simpler to analyze. Our results imply that there exist absolute constants $0 < c_1 \leq c_2 < \infty$ such that for any fixed $x \geq 0$,

$$\begin{aligned} \liminf \mathbb{P}[(1 - q)^2(\max(\mathcal{C}_s) - s) \geq x] &\geq e^{-c_2x}, \\ \limsup \mathbb{P}[(1 - q)^2(\max(\mathcal{C}_s) - s) \geq x] &\leq e^{-c_1x}. \end{aligned}$$

The lower bound follows from Corollary 4.4 and the upper bound follows by again considering Proposition 4.6 together with Theorem 3.6. We conjecture that, in fact, there exists a single absolute constant $c > 0$ for which

$$\lim \mathbb{P}[(1 - q)^2(\max(\mathcal{C}_s) - s) \geq x] = e^{-cx}, \quad x \geq 0.$$

That is, that the limiting distribution of $(1 - q)^2(\max(\mathcal{C}_s) - s)$ is exponential. By symmetry, the same is conjectured for $(1 - q)^2(s - \min(\mathcal{C}_s))$. Furthermore, we conjecture that the dependence between $\max(\mathcal{C}_s)$ and $\min(\mathcal{C}_s)$ disappears in this limit, so that the diameter $(1 - q)^2(\max(\mathcal{C}_s) - \min(\mathcal{C}_s))$ converges in distribution to the sum of two independent, identically distributed, exponential random variables.

Intermediate regime. Suppose that $n(1 - q)^2 \rightarrow \beta \in (0, \infty)$. We expect the limiting distributions to still exist in this regime and interpolate in a continuous manner the previous two cases. For the cycle diameter, this interpolation may possibly be achieved by truncation, as certainly $\max(\mathcal{C}_s) - s \leq n - s$. Recalling that $\frac{s}{n} \rightarrow \alpha$, we conjecture, for instance, that the limiting distribution of $(1 - q)^2(\max(\mathcal{C}_s) - s)$ is equal to that of $\min\{X, \beta(1 - \alpha)\}$ where X is the exponential random variable conjectured to give the limiting distribution in the previous regime.

Extensions of the parameter range. As discussed in the [Introduction](#), there exist extensions of the Mallows distribution to infinite permutations; one-to-one and onto functions $\pi : \mathbb{N} \rightarrow \mathbb{N}$ or $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$. We expect the analogues of our main theorems regarding the expected length, variance of the length and expected diameter of cycles, [Theorem 1.1](#), [Theorem 1.3](#) and equation (3) of [Theorem 1.4](#), to continue to hold for these models, plugging formally $n = \infty$ and taking $s \in \mathbb{N}$ or $s \in \mathbb{Z}$ according to the case. Such results may follow from our methods, using the sampling mechanism described in [Section 3.5](#) for the case $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$, but we do not develop this further. The approximation theorems of Gnedin and Olshanski [[14](#)], [Section 7.2](#), may also prove useful in this context.

The Mallows distribution with parameters n and q is defined for the case that $q > 1$ via the same formula (1). The distributions with parameters q and $\frac{1}{q}$ are related: If $\pi \sim \mu_{n,q}$ then $\pi \circ r$, with $r_s = n - s + 1$, is distributed $\mu_{n,1/q}$. This operation corresponds to reflecting the graph of the permutation π across the line $x = \frac{n+1}{2}$. Cycles are significantly affected by this operation as, when $\pi \sim \mu_{n,q}$ with $q > 1$, π_i typically lies around $n + 1 - i$, as follows from (2) and the above relation. Thus, for instance, the diameter of the cycle containing 1 in π may be close to n for all $q > 1$. Still, we would expect that for all $q > 1$ the expected cycle lengths are still of order $\min\{\frac{1}{(1-1/q)^2}, n\}$ as in [Theorem 1.1](#) and further expect a Poisson–Dirichlet limit law for the long cycles when $n \rightarrow \infty$, $q \rightarrow 1+$ and $n(q - 1)^2 \rightarrow 0$ as in [Theorem 1.2](#). This is suggested by the fact that composing π with itself leads to a permutation whose graph is qualitatively similar to that of the Mallows model with $q < 1$; see [Figure 7](#).

Our study of the cycle lengths of the Mallows permutation with parameter $q < 1$ was based on the diagonal exposure process. Possibly, a similar process may be used to study the case $q > 1$ by exposing the graph of the permutation “from the center outwards”. That is, exposing after t iterations the portion of the graph contained in a square of side length $2t$ around the mid-point $(\frac{n+1}{2}, \frac{n+1}{2})$. The ideas in [Section 3.5](#) may be useful in making such an approach rigorous.




FIG. 7. On the left is a graph of a sample of the Mallows distribution $\mu_{n,q}$ with $n = 1000$ and $q = 1.02$. On the right is the graph of the composition of the same permutation with itself. The red lines are at vertical distance $\frac{2}{1-q}$ (left) and $\frac{2\sqrt{2}}{1-q}$ (right) from the diagonal. They delimit a region containing most of the points of the permutation.

Band models. We expect analogues of our results to hold for other natural models of random permutations whose graph typically has a “band structure”. For instance, for the interchange model on the one-dimensional graph $\{1, \dots, n\}$ with nearest-neighbor edges. In this case, as briefly discussed in Section 1.2.4, the analogous result for the expected cycle length has been proved by Kozma and Sidoravicius. Another natural model is the band-Poisson model. Here, for an integer $n \geq 1$ and real $0 < w \leq n$, one considers a Poisson point process with intensity $\frac{1}{w}$ in the continuous band given by

$$\{(x, y) \in [0, n]^2 \mid |x - y| \leq w\},$$

where the parameter w controls the width of the band. With this definition, each vertical strip of width 1 in $[0, n]^2$ contains on average a constant number of points. Each realization of the process gives rise to a permutation by taking the relative order of the points as in (39). We expect the analogues of our results to hold for this model with w standing for $\frac{q}{1-q}$.

Higher dimensions and general graphs. As discussed in Section 1.2.4, the study of cycles of spatial random permutations, random permutations biased towards the identity in an underlying geometry, is of great interest. The special case in which the geometry is that of \mathbb{R}^d or \mathbb{Z}^d is particularly significant with relations to models of statistical physics. In this context, our work pertains to the case $d = 1$. With other geometries in mind, we note here that a Mallows model may be defined on any finite connected graph $G = (V, E)$ and parameter $0 < q \leq 1$ by letting the probability of a permutation $\pi : V \rightarrow V$ be given by

$$\mathbb{P}_{G,q}[\pi] = \frac{1}{Z_{G,q}} q^{d(\pi, \text{Id})},$$

where $Z_{G,q}$ is a normalization constant and $d(\pi, \text{Id})$ stands for the minimal number of adjacent transpositions required to change π to the identity permutation $\text{Id}: v \mapsto v$. Here, an adjacent transposition is a transposition of the endpoints of an edge of G . For instance, any transposition is allowed on the complete graph on n vertices K_n , whence the model coincides with the well-studied Ewens model [11], with parameter $\theta = \frac{1}{q}$. This follows from the fact that

$$\mathbb{P}_{K_n,q}[\pi] = \frac{q^n}{Z_{K_n,q}} q^{-\mathcal{N}(\pi)}$$

with $\mathcal{N}(\pi)$ denoting the number of cycles (including fixed points) in π . Our analysis of the Mallows model is based on the exact sampling algorithm given by (6). Unfortunately, we are not aware of corresponding algorithms for general graphs (though in the specific case of the Ewens model an algorithm is given by the so-called Chinese restaurant process). Nonetheless, it is of interest to obtain results on the length of long cycles for general graphs G , with the case that G is a box in Z^d , $d \geq 2$, having special significance.

Acknowledgments. We thank Nayantara Bhatnagar, Gady Kozma, Grigori Olshanski and Sasha Sodin for useful discussions. We are also grateful to an anonymous referee for a detailed reading of the paper and many excellent comments which contributed significantly to the presentation.

REFERENCES

- [1] ANGEL, O. (2003). Random infinite permutations and the cyclic time random walk. In *Discrete Random Walks (Paris, 2003)*. *Discrete Math. Theor. Comput. Sci. Proc.*, AC 9–16. Assoc. Discrete Math. Theor. Comput. Sci., Nancy. [MR2042369](#)
- [2] BASU, R. and BHATNAGAR, N. (2017). Limit theorems for longest monotone subsequences in random Mallows permutations. *Ann. Inst. Henri Poincaré Probab. Stat.* **53** 1934–1951. [MR3729641](#)
- [3] BENJAMINI, I., BERGER, N., HOFFMAN, C. and MOSSEL, E. (2005). Mixing times of the biased card shuffling and the asymmetric exclusion process. *Trans. Amer. Math. Soc.* **357** 3013–3029. [MR2135733](#)
- [4] BERESTYCKI, N. (2011). Emergence of giant cycles and slowdown transition in random transpositions and k -cycles. *Electron. J. Probab.* **16** 152–173. [MR2754801](#)
- [5] BETZ, V. and UELTSCHI, D. (2011). Spatial random permutations and Poisson–Dirichlet law of cycle lengths. *Electron. J. Probab.* **16** 1173–1192. [MR2820074](#)
- [6] BHATNAGAR, N. and PELED, R. (2015). Lengths of monotone subsequences in a Mallows permutation. *Probab. Theory Related Fields* **161** 719–780. [MR3334280](#)
- [7] BORODIN, A., DIACONIS, P. and FULMAN, J. (2010). On adding a list of numbers (and other one-dependent determinantal processes). *Bull. Amer. Math. Soc. (N.S.)* **47** 639–670. [MR2721041](#)
- [8] BRAVERMAN, M. and MOSSEL, E. (2009). Sorting from noisy information. *CoRR abs/0910.1191*.
- [9] DIACONIS, P. and RAM, A. (2000). Analysis of systematic scan Metropolis algorithms using Iwahori–Hecke algebra techniques. *Michigan Math. J.* **48** 157–190. Dedicated to William Fulton on the occasion of his 60th birthday. [MR1786485](#)

- [10] DIACONIS, P. and SHAHSHAHANI, M. (1981). Generating a random permutation with random transpositions. *Z. Wahrsch. Verw. Gebiete* **57** 159–179. [MR0626813](#)
- [11] EWENS, W. J. (1972). The sampling theory of selectively neutral alleles. *Theor. Popul. Biol.* **3** 87–112; erratum, *ibid.* **3** (1972), 240; erratum, *ibid.* **3** (1972), 376. [MR0325177](#)
- [12] FENG, S. (2010). *The Poisson–Dirichlet Distribution and Related Topics: Models and Asymptotic Behaviors*. Springer, Heidelberg. [MR2663265](#)
- [13] GNEDIN, A. and OLSHANSKI, G. (2010). q -exchangeability via quasi-invariance. *Ann. Probab.* **38** 2103–2135. [MR2683626](#)
- [14] GNEDIN, A. and OLSHANSKI, G. (2012). The two-sided infinite extension of the Mallows model for random permutations. *Adv. in Appl. Math.* **48** 615–639. [MR2920835](#)
- [15] HAMMOND, A. (2013). Infinite cycles in the random stirring model on trees. *Bull. Inst. Math. Acad. Sin. (N.S.)* **8** 85–104. [MR3097418](#)
- [16] HAMMOND, A. (2015). Sharp phase transition in the random stirring model on trees. *Probab. Theory Related Fields* **161** 429–448. [MR3334273](#)
- [17] KENYON, R., KRÁL', D., RADIN, C. and WINKLER, P. Permutations with fixed pattern densities. Available at [arXiv:1506.02340](#).
- [18] KOTECKÝ, R., MIŁOŚ, P. and UELTSCHI, D. The random interchange process on the hypercube. Available at [arXiv:1509.02067](#).
- [19] LEVIN, D. A., PERES, Y. and WILMER, E. L. (2009). *Markov Chains and Mixing Times*. Amer. Math. Soc., Providence, RI. [MR2466937](#)
- [20] MALLOWS, C. L. (1957). Non-null ranking models. I. *Biometrika* **44** 114–130. [MR0087267](#)
- [21] MUELLER, C. and STARR, S. (2013). The length of the longest increasing subsequence of a random Mallows permutation. *J. Theoret. Probab.* **26** 514–540. [MR3055815](#)
- [22] MUKHERJEE, S. (2016). Fixed points and cycle structure of random permutations. *Electron. J. Probab.* **21** Paper No. 40, 18. [MR3515570](#)
- [23] OLSHANSKI, G. (2011). Random permutations and related topics. In *The Oxford Handbook of Random Matrix Theory* 510–533. Oxford Univ. Press, Oxford. [MR2932645](#)
- [24] SCHRAMM, O. (2005). Compositions of random transpositions. *Israel J. Math.* **147** 221–243. [MR2166362](#)
- [25] SPENCER, T. (2011). Random banded and sparse matrices. In *The Oxford Handbook of Random Matrix Theory* 471–488. Oxford Univ. Press, Oxford. [MR2932643](#)
- [26] STANLEY, R. P. (2012). *Enumerative Combinatorics. Volume 1*, 2nd ed. *Cambridge Studies in Advanced Mathematics* **49**. Cambridge Univ. Press, Cambridge. [MR2868112](#)
- [27] STARR, S. (2009). Thermodynamic limit for the Mallows model on S_n . *J. Math. Phys.* **50** 095208, 15. [MR2566888](#)
- [28] STARR, S. and WALTERS, M. Phase uniqueness for the Mallows measure on permutations. Available at [arXiv:1502.03727](#).
- [29] SÜTÓ, A. (2002). Percolation transition in the Bose gas. II. *J. Phys. A* **35** 6995–7002. [MR1945163](#)
- [30] TÓTH, B. (1993). Improved lower bound on the thermodynamic pressure of the spin 1/2 Heisenberg ferromagnet. *Lett. Math. Phys.* **28** 75–84. [MR1224836](#)

SCHOOL OF MATHEMATICAL SCIENCES
TEL AVIV UNIVERSITY
TEL AVIV, 69978
ISRAEL
E-MAIL: alx112358@gmail.com
peledron@post.tau.ac.il