

THE RUNSORT PERMUTON

NOGA ALON, COLIN DEFANT, AND NOAH KRAVITZ

ABSTRACT. Suppose we choose a permutation π uniformly at random from S_n . Let $\text{runsort}(\pi)$ be the permutation obtained by sorting the ascending runs of π into lexicographic order. Alexandersson and Nabawanda recently asked if the plot of $\text{runsort}(\pi)$, when scaled to the unit square $[0, 1]^2$, converges to a limit shape as $n \rightarrow \infty$. We answer their question by showing that the measures corresponding to the scaled plots of these permutations $\text{runsort}(\pi)$ converge with probability 1 to a permuton (limiting probability distribution) that we describe explicitly. In particular, the support of this permuton is $\{(x, y) \in [0, 1]^2 : x \leq ye^{1-y}\}$.

1. INTRODUCTION

Let S_n denote the set of permutations of the set $[n] = \{1, \dots, n\}$. The *scaled plot* of a permutation $\pi = \pi_1 \cdots \pi_n \in S_n$ is the diagram showing the points $(i/n, \pi_i/n)$ for $i \in [n]$. The scaled plot of π is closely related to the probability measure γ_π on the unit square $[0, 1]^2$ defined as follows. Consider a point $(x, y) \in [0, 1]^2$. If $(i-1)/n \leq x < i/n$ and $(\pi_i-1)/n \leq y < \pi_i/n$ for some $i \in [n]$, then γ_π has density n at (x, y) ; otherwise, γ_π has density 0 at (x, y) . In other words, we divide $[0, 1]^2$ into an $n \times n$ grid of squares, each with side length $1/n$, and we assign each square a constant density of either n or 0, according to whether or not the upper right corner of the square is a point in the scaled plot of π .

A *permuton* is a probability measure γ on the unit square $[0, 1]^2$ that has uniform marginals, in the sense that $\gamma([a, b] \times [0, 1]) = \gamma([0, 1] \times [a, b]) = b - a$ for all $0 \leq a \leq b \leq 1$. Note that if $\pi \in S_n$, then γ_π is a permuton. (This is why we chose to scale with density n .) There is a natural topology on the space of permutons obtained by restricting the weak topology on probability measures. This




FIGURE 1. On the left is the scaled plot of $\text{runsort}(\pi)$, where π is a permutation chosen uniformly at random from S_{50000} . On the right is the runsort permuton \mathbf{R} . The dark curve \mathcal{C} is the support of the singular continuous part of \mathbf{R} . Shading within the region \mathcal{C}^+ indicates the density of the absolutely continuous part of \mathbf{R} .

coincides with the topology induced by the metric d_{\square} defined by

$$d_{\square}(\gamma_1, \gamma_2) = \sup_R |\gamma_1(R) - \gamma_2(R)|,$$

where we take the supremum over all axis-parallel rectangles in the unit square. There has been a great deal of recent interest in describing permutons that arise as limits of large random permutations. Several results in this area concern the interplay between random permutations and permutation patterns [3, 8–11, 16]. In a different direction, Dauvergne [6] recently proved a beautiful result about permutons arising from random sorting networks (see also [2, 7]). Our goal in this article is to describe the permuton that emerges when we apply an operator called `runsort` to a large random permutation.

An *ascending run* of a permutation (henceforth called a *run* for simplicity) is a maximal consecutive increasing subsequence. For instance, the runs of 351476298 are 35, 147, 6, 29, and 8. Given a permutation $\pi \in S_n$, let `runsort`(π) be the permutation obtained by sorting the runs of π into lexicographic order. Equivalently, `runsort` sorts the runs so that their smallest entries appear in increasing order. For example, `runsort`(351476298) = 147293568. Note that `runsort` is an idempotent operator.

Motivated by the study of flattened partitions (see [5, 12, 13, 17]), Alexandersson and Nabawanda [4] proved several interesting combinatorial properties of `runsort`. When they chose $\pi \in S_n$ (for large n) uniformly at random and plotted the permutation `runsort`(π), they observed that it tended to have a very distinctive shape (see the left side of Figure 1). Furthermore, they noticed that the scaled plot of `runsort`(π) appeared to be bounded by a certain enveloping curve, and they asked if this curve approaches some limit curve as $n \rightarrow \infty$. In this paper, we answer their question (in a strong form) by showing that $\gamma_{\text{runsort}(\pi)}$ converges with high probability to a specific permuton.

Consider the curve

$$\mathcal{C} = \{(x, y) \in [0, 1]^2 : x = ye^{1-y}\},$$

and let

$$\mathcal{C}^+ = \{(x, y) \in [0, 1]^2 : x < ye^{1-y}\}$$

be the region in the unit square above \mathcal{C} . We define the *runsort permuton* \mathbf{R} to be the permuton given by

$$\mathbf{R}(B) = \int_{B \cap \mathcal{C}^+} e^{y-1} dx dy + \int_{B \cap \mathcal{C}} (1-y) dy$$

for every measurable set $B \subseteq [0, 1]^2$. Thus, \mathbf{R} is the sum of its absolutely continuous part, which has support \mathcal{C}^+ , and its singular continuous part, which has support \mathcal{C} . One can easily confirm by direct computation that \mathbf{R} is in fact a permuton (and we encourage the reader to do this).

Our main result is that \mathbf{R} is the limiting distribution for the image under `runsort` of a large random permutation.

Theorem 1.1 (Main Theorem). *Fix any $\varepsilon > 0$, and choose $\pi^{(n)}$ uniformly at random from S_n . Then $d_{\square}(\gamma_{\text{runsort}(\pi^{(n)})}, \mathbf{R}) < \varepsilon$ with probability tending to 1 as $n \rightarrow \infty$.*

In other words, if we randomly choose permutations $\pi^{(n)}$, then the measures $\gamma_{\text{runsort}(\pi^{(n)})}$ converge to \mathbf{R} with probability 1.

We remark that a combination of absolutely continuous and singular continuous parts, such as what \mathbf{R} exhibits, seems to be rare in previous work on permutons.

In our setting, however, it is quite natural. As will become clear later, the singular part comes from the first entries of the runs, and the continuous part comes from the remaining entries. A random permutation in S_n has $(n+1)/2$ runs on average, and the reader can check that indeed half of the total mass of \mathbf{R} lies on the curve \mathcal{C} . More surprising is that the pointwise density of \mathbf{R} in \mathcal{C}^+ depends only on the vertical coordinate y . This “horizontal uniformity” is not combinatorially

obvious, and we do not know how to establish it directly (that is, without recourse to the explicit characterization of \mathbf{R}).

In Section 2, we use martingale concentration inequalities to show that it suffices to understand the expectation of $\gamma_{\text{runsort}(\pi^{(n)})}$. In Section 3, we treat the enveloping curve \mathcal{C} and determine the distribution of mass on it. Section 4 is devoted to the density in the region \mathcal{C}^+ . Section 5 concludes the proof of Theorem 1.1. Finally, in Section 6, we briefly mention a generalization of `runsort` in the context of pattern avoidance, and we outline how our method could be adapted to treat these new problems.

1.1. Notation. We write π_i for the i -th entry of a permutation π . We write $\mathbb{P}[A]$ for the probability of an event A and write $\mathbb{E}[X]$ for the expected value of a random variable X .

2. CONCENTRATION OF DISTRIBUTION

Our task is to show that if π is chosen randomly from S_n , then $\gamma_{\text{runsort}(\pi)}$ is suitably concentrated everywhere. The key observation is that transposing two entries of π has only a small effect on $\gamma_{\text{runsort}(\pi)}$, as long as π does not have any unusually long runs. In order to make this notion precise, we define the following variant of `runsort`: Let $\pi \in S_n$ be a permutation, and let A_1, \dots, A_t be the runs of π . We can break each run A_k into shorter subsequences $A_{k,1}, \dots, A_{k,t(k)}$, which we call *segments*, as follows: If A_k has length smaller than $\log n$, then set $t(k) = 1$ and $A_{k,1} = A_k$; if A_k has length greater than $\log n$, then break it directly before every position (of π) that is an integer multiple of $\lfloor \log n \rfloor$, so that $A_{k,\ell}$ has length exactly $\lfloor \log n \rfloor$ for $1 < \ell < t(k)$, and $A_{k,1}$ and $A_{k,t(k)}$ have length at most $\lfloor \log n \rfloor$. We then define $\overline{\text{runsort}}(\pi)$ to be the permutation obtained by sorting the runs $A_{k,\ell}$ into lexicographic order. Note that $\text{runsort}(\pi) = \overline{\text{runsort}}(\pi)$ if all of the runs of π have length at most $\lfloor \log n \rfloor$; in particular, the following lemma tells us that if π is chosen randomly from S_n , then $\text{runsort}(\pi) = \overline{\text{runsort}}(\pi)$ with probability tending very quickly to 1 as $n \rightarrow \infty$.

Lemma 2.1. *Suppose π is chosen uniformly at random from S_n . If n is sufficiently large, then the probability that every run of π has length at most $\log n$ is at least $1 - n^{-(\log \log n)/2}$.*

Proof. Let $T = \lfloor \log n \rfloor + 1$. For each $i \in [n - T + 1]$, the probability that the entries of π in positions $i, i + 1, \dots, i + T - 1$ appear in increasing order is $1/T!$. Therefore, the probability that there is a run of π of length at least T is at most $(n - T + 1)/T!$, which, by Stirling's Formula, is at most $n^{-(\log \log n)/2}$ for n sufficiently large. \square

Lemma 2.2. *Let $B = [x_1, x_2] \times [y_1, y_2] \subset [0, 1]^2$ be a rectangle. Let $\pi \in S_n$ be a permutation, and let $\pi' = \pi \circ (i_1 i_2)$ be the permutation obtained from π by swapping the entries in positions i_1 and i_2 (that is, applying the transposition $(i_1 i_2)$). Then $|\gamma_{\text{runsort}(\pi)}(B) - \gamma_{\overline{\text{runsort}}(\pi)}(B)| \leq 20 \log n/n$.*

Proof. We note that $|\gamma_{\overline{\text{runsort}}(\pi)}(B) - \gamma_{\text{runsort}(\pi')} (B)|$, viewed as a function of x_1, x_2, y_1, y_2 , clearly attains its maximum value when x_1, x_2, y_1, y_2 are all integer multiples of $1/n$, so it suffices to prove the lemma when B has this special form. In this case, we discretize the problem by writing

$$(1) \quad \gamma_{\overline{\text{runsort}}(\pi)}(B) = \frac{\#\{i : x_1 n < i \leq x_2 n \text{ and } y_1 n < \overline{\text{runsort}}(\pi)_i \leq y_2 n\}}{n}$$

(since each point $(i, \overline{\text{runsort}}(\pi)_i)$ contributes mass $1/n$ after rescaling), and likewise for $\gamma_{\text{runsort}(\pi')} (B)$. When we apply the transposition $(i_1 i_2)$ to π , all but at most four of the segments $A_{k,\ell}$ remain the same. Since each segment has length at most $\log n$, we see that at most $4 \log n$ entries are in these ‘‘affected’’ segments. It follows that for each entry j not in one of these affected runs, the horizontal positions of j in $\overline{\text{runsort}}(\pi)$ and $\overline{\text{runsort}}(\pi')$ differ by at most $4 \log n$. In particular, these horizontal positions are either both in the interval $(x_1 n, x_2 n]$ or neither in this interval, unless the position in $\overline{\text{runsort}}(\pi)$ was within $4 \log n$ of either $x_1 n$ or $x_2 n$. So the change in the numerator on the right-hand side of (1) coming from these entries is at most $16 \log n$, and we must absorb an additional error of $4 \log n$ because we have no control over the positions of the entries of the affected segments. \square

(It is possible to replace the constant 20 by 4 in this lemma, but we do not optimize this constant because it is irrelevant in what follows.)

We require the following martingale inequality from [14] (see also page 35 of the book [15]).

Proposition 2.3 ([14, 15]). *Let $f : S_n \rightarrow \mathbb{R}$ be a function on permutations such that if $\pi, \pi' \in S_n$ differ by a transposition, then $|f(\pi) - f(\pi')| \leq z$. Then for $\pi \in S_n$ chosen uniformly at random, we have*

$$\mathbb{P}[|f(\pi) - \mathbb{E}[f]| \geq cz] \leq 2e^{-c^2/(4n)}.$$

We note that by following the standard proof of Azuma's Inequality (see, e.g., [1, Theorem 7.2.1]) with the obvious modification needed to deal with permutations, it is possible to improve the above exponent from $-c^2/(4n)$ to $-c^2/(2n)$, but this is not necessary for our applications.

Before applying this proposition, we fix a bit of notation. For $1 \leq i, j \leq n$, let $p_n(i, j)$ denote the probability that the entry j is in the i -th position of $\text{runsort}(\pi)$ when $\pi \in S_n$ is chosen uniformly at random. For $B = [x_1, x_2] \times [y_1, y_2] \subseteq [0, 1]^2$, let

$$E_n(B) = \frac{1}{n} \sum_{\substack{x_1 n < i \leq x_2 n, \\ y_1 n < j \leq y_2 n}} p_n(i, j).$$

Define the analogous quantities $\bar{p}_n(i, j)$ and $\bar{E}_n(B)$ with runsort replaced by $\overline{\text{runsort}}$.

Lemma 2.4. *Fix any small $\varepsilon > 0$, and suppose that $\pi \in S_n$ is chosen uniformly at random. If n is sufficiently large (as a function of ε), then with probability at least $1 - \varepsilon$, the concentration inequality*

$$|\gamma_{\text{runsort}(\pi)}(B) - E_n(B)| < \varepsilon$$

holds simultaneously for all axis-parallel rectangles $B \subseteq [0, 1]^2$.

Proof. We make two reductions. First, write $B = [x_1, x_2] \times [y_1, y_2]$ and define $B' = [x'_1, x'_2] \times [y'_1, y'_2]$ where $x'_1 = \lfloor nx_1 \rfloor / n$, $x'_2 = \lfloor nx_2 \rfloor / n$, $y'_1 = \lfloor ny_1 \rfloor / n$, $y'_2 = \lfloor ny_2 \rfloor / n$. Then $E_n(B) = E_n(B')$ and

$$|\gamma_{\text{runsort}(\pi)}(B) - \gamma_{\text{runsort}(\pi)}(B')| \leq 2/n.$$

Since this can be made smaller than $\varepsilon/3$, it suffices to show that

$$|\gamma_{\text{runsort}(\pi)}(B) - E_n(B)| < 2\varepsilon/3$$

for all B whose vertices have coordinates that are integer multiples of $1/n$. Note that in this case, the quantities $E_n(B)$ and $\bar{E}_n(B)$ are precisely the expected values of $\gamma_{\text{runsort}(\pi)}(B)$ and $\gamma_{\overline{\text{runsort}}(\pi)}(B)$, respectively.

Second, we know from Lemma 2.1 that $\text{runsort}(\pi) = \overline{\text{runsort}}(\pi)$ with probability at least $1 - n^{-(\log \log n)/2}$. In particular, $|p_n(i, j) - \bar{p}_n(i, j)| \leq n^{-(\log \log n)/2}$ for all i, j , so $|E_n(B) - \bar{E}_n(B)| \leq n \cdot n^{-(\log \log n)/2}$. We can make this last quantity smaller than $\varepsilon/3$ by choosing n large enough, so it suffices to show that with probability at least $1 - 2\varepsilon/3$, the inequality

$$|\gamma_{\overline{\text{runsort}}(\pi)}(B) - \bar{E}_n(B)| < \varepsilon/3$$

holds simultaneously for all axis-parallel rectangles B whose coordinates are integer multiples of $1/n$.

For each such B with coordinates that are integer multiples of $1/n$, combining Lemma 2.2 and Proposition 2.3 (with $z = 20 \log n/n$ and $c = \varepsilon n/60 \log n$) gives that

$$|\gamma_{\overline{\text{runsort}}(\pi)}(B) - \bar{E}_n(B)| < \varepsilon/3$$

with probability at least $1 - e^{-\frac{\varepsilon^2 n}{14400(\log n)^2}}$. A union bound gives that with probability at least

$$1 - n^4 e^{-\frac{\varepsilon^2 n}{14400(\log n)^2}},$$

the above inequality holds simultaneously for all B with coordinates that are integer multiples of $1/n$, and taking n large guarantees that this probability is at least $1 - 2\varepsilon/3$, as needed. \square

This lemma tells us that it will suffice to work with the expectations $E_n(B)$. In particular, we have reduced Theorem 1.1 to the following more manageable-looking statement. For each n , define the probability measure \mathbf{R}_n via $\mathbf{R}_n(B) = E_n(B)$ for all axis-parallel rectangles $B \subseteq [0, 1]^2$.

Theorem 2.5. *The measures \mathbf{R}_n converge to \mathbf{R} as n goes to infinity.*

3. THE ENVELOPING CURVE

In this section, we address the ‘‘singular’’ behavior that comes from the first entries of runs; as mentioned in the Introduction, this corresponds to the mass in \mathbf{R} that lies on the curve \mathcal{C} .

For $y \in [0, 1]$ and $\pi \in S_n$, we let $L_\pi(y)$ denote the largest position $i \in [n]$ such that $\pi_i \leq yn$ if $y \geq 1/n$, and we define $L_\pi(y) = 0$ if $y < 1/n$. In other words, $L_\pi(y)$ is the smallest i such that all of the entries up to yn appear in the first i positions. Note that the curve $x = L_\pi(y)$ is the lower envelope of the scaled plot of π . We start by computing the expected value of $L_{\text{runsort}(\pi)}(y)$ and $L_{\overline{\text{runsort}(\pi)}}(y)$.

Proposition 3.1. *Fix $y \in [0, 1]$, and suppose that $\pi \in S_n$ is chosen uniformly at random. Then*

$$\mathbb{E}[L_{\text{runsort}(\pi)}(y)] = nye^{1-y} + O(\log n)$$

and

$$\mathbb{E}[L_{\overline{\text{runsort}(\pi)}}(y)] = nye^{1-y} + O(\log n).$$

Proof. We start with the first statement. For $j \in [n]$, define the random variable X_j by

$$X_j(\pi) = \begin{cases} 0, & \text{if } j \text{ is not the beginning of a run of } \pi; \\ k, & \text{if } j \text{ is the beginning of a run of length } k \text{ of } \pi. \end{cases}$$

It follows from the definition of `runsort` that

$$\sum_{j \leq yn} X_j(\pi) - m(\pi) < L_{\text{runsort}(\pi)}(y) \leq \sum_{j \leq yn} X_j(\pi),$$

where $m(\pi)$ is the maximum length of a run in π . We know from Lemma 2.1 that $m(\pi) \leq \log n$ with probability at least $1 - n^{-(\log \log n)/2}$. Since $m(\pi) \leq n$ for all $\pi \in S_n$, we have

$$\mathbb{E}[L_{\text{runsort}(\pi)}(y)] \geq \sum_{j \leq yn} \mathbb{E}[X_j(\pi)] - (1 - n^{-(\log \log n)/2}) \log n - n \cdot n^{-(\log \log n)/2}.$$

Hence,

$$\mathbb{E}[L_{\text{runsort}(\pi)}(y)] = \sum_{j \leq yn} \mathbb{E}[X_j(\pi)] + O(\log n).$$

Consider $j \in [n]$ with $j \leq yn$, and let r be such that $\pi_r = j$. We wish to estimate the probability that $X_j(\pi) \geq k+1$ for each $k \geq 0$. If $k > \log n$, then $\mathbb{P}[X_j(\pi) \geq k+1] < n^{-(\log \log n)/2}$ by Lemma 2.1 (and this contribution will turn out to be negligible). To handle the case $k \leq \log n$, note that we have $X_j(\pi) \geq k+1$ if and only if the following both hold:

- Either $r = 1$, or $2 \leq r \leq n - k$ and $\pi_{r-1} > j$;
- $j < \pi_{r+1} < \pi_{r+2} < \dots < \pi_{r+k}$.

Therefore,

$$\mathbb{P}[X_j(\pi) \geq k+1] = \frac{1}{n \cdot k!} \cdot \frac{\binom{n-j}{k}}{\binom{n-1}{k}} + \frac{n-k-1}{n \cdot k!} \cdot \frac{\binom{n-j}{k+1}}{\binom{n-1}{k+1}}.$$

Note that the first term is at most $1/n$. To estimate the second term, we write

$$\begin{aligned} \frac{\binom{n-j}{k+1}}{\binom{n-1}{k+1}} &= \prod_{t=0}^k \frac{n-j-t}{n-1-t} \\ &= \prod_{t=0}^k \left(1 - \frac{j}{n} + \frac{n-j-jt}{n(n-1-t)} \right) \\ &= \left(1 - \frac{j}{n} + O\left(\frac{t+1}{n}\right) \right)^{k+1} \\ &= (1-j/n)^{k+1} + O((k+1)^2/n), \end{aligned}$$

where the last bound uses the fact that $k \leq \log n$. Combining these estimates yields

$$\begin{aligned} \mathbb{P}[X_j(\pi) \geq k+1] &= O(1/n) + \frac{1}{k!} (1 - (k+1)/n) \left((1-j/n)^{k+1} + O((k+1)^2/n) \right) \\ &= \frac{(1-j/n)^{k+1}}{k!} + O(1/n), \end{aligned}$$

again using $k \leq \log n$. So

$$\begin{aligned} \mathbb{E}[X_j(\pi)] &= \sum_{k=0}^{n-1} \mathbb{P}[X_j(\pi) \geq k+1] \\ &= \sum_{k=0}^{\lfloor \log n \rfloor} \left(\frac{(1-j/n)^{k+1}}{k!} + O(1/n) \right) + O\left(n \cdot n^{-(\log \log n)/2}\right) \\ &= (1-j/n)(e^{1-j/n} + O(1/\lfloor \log n \rfloor!)) + O(\log n/n) + O\left(n \cdot n^{-(\log \log n)/2}\right) \\ &= (1-j/n)e^{1-j/n} + O(\log n/n). \end{aligned}$$

Summing over $j \leq yn$ gives

$$\begin{aligned} \sum_{j \leq yn} \mathbb{E}[X_j(\pi)] &= \sum_{j \leq yn} \left((1-j/n)e^{1-j/n} + O(\log n/n) \right) \\ &= n \int_0^y (1-t)e^{1-t} dt + O(1) + O(\log n) \\ &= ny e^{1-y} + O(\log n), \end{aligned}$$

and we conclude that

$$\mathbb{E}[L_{\text{runsort}(\pi)}(y)] = ny e^{1-y} + O(\log n).$$

For the statement about $\mathbb{E}[L_{\overline{\text{runsort}}(\pi)}(y)]$, note that $L_{\overline{\text{runsort}}(\pi)}(y) = L_{\text{runsort}(\pi)}(y)$ with probability at least $1 - n^{-(\log \log n)/2}$ by Lemma 2.1; when these quantities do differ, they differ by at most n , and $n \cdot n^{-(\log \log n)/2}$ can be absorbed into the error term. \square

In order to apply Proposition 2.3, we need an analogue of Lemma 2.2. As in the previous section, it is more convenient to work with $\overline{\text{runsort}}$.

Lemma 3.2. *Fix $y \in [0, 1]$. Let $\pi \in S_n$ be a permutation, and let $\pi' = \pi \circ (i_1 i_2)$ be the permutation obtained from π by swapping the entries in positions i_1 and i_2 . Then*

$$|L_{\overline{\text{runsort}}(\pi)}(y) - L_{\overline{\text{runsort}}(\pi')}(y)| \leq 9 \log n.$$

Proof. Write $\{A_{k,\ell}\}$ and $\{A'_{k,\ell}\}$ for the sets of segments of π and π' , respectively. As in the proof of Lemma 2.2, we note that multiplying π on the right by the transposition $(i_1 i_2)$ affects at most four of the segments $A_{k,\ell}$, which together contain at most $4 \log n$ entries. In particular, for each entry j not in one of these affected segments, the horizontal positions of j in $\overline{\text{runsort}}(\pi)$ and $\overline{\text{runsort}}(\pi')$ differ by at most $4 \log n$. Hence, the first $L_{\overline{\text{runsort}}(\pi)}(y) + 4 \log n$ entries of $\overline{\text{runsort}}(\pi')$ certainly contain all of the entries up to yn except for possibly some of the $4 \log n$ entries in the affected segments. These missing small entries (if any exist) are contained in at most four segments $A'_{k,\ell}$ of π' , and these segments must appear in $\overline{\text{runsort}}(\pi')$ directly after the last run that includes all smaller entries. Since all runs have length at most $\log n$, we see that all of the entries up to yn are among the first

$$L_{\overline{\text{runsort}}(\pi)}(y) + 4 \log n + \log n + 4 \log n = L_{\overline{\text{runsort}}(\pi)}(y) + 9 \log n$$

entries of $\overline{\text{runsort}}(\pi')$; that is, $L_{\overline{\text{runsort}}(\pi')}(y) \leq L_{\overline{\text{runsort}}(\pi)}(y) + 9 \log n$. By the same argument, we have $L_{\overline{\text{runsort}}(\pi)}(y) \leq L_{\overline{\text{runsort}}(\pi')}(y) + 9 \log n$. \square

Following the same strategy as in the previous section, we obtain a concentration inequality for $L_{\text{runsort}(\pi)}(y)$.

Lemma 3.3. *Fix $y \in [0, 1]$, and suppose that $\pi \in S_n$ is chosen uniformly at random. Then with probability at least $1 - 2n^{-(\log \log n)/2}$, we have the concentration inequality*

$$|L_{\text{runsort}(\pi)}(y) - nye^{1-y}| = O(\sqrt{n}(\log n)^2).$$

Proof. Applying Proposition 2.3 with $z = 9 \log n$ and $c = 2\sqrt{n} \log n$ gives that with probability at least $1 - 2e^{-(\log n)^2}$, the quantity $L_{\overline{\text{runsort}}(\pi)}(y)$ differs from its expectation by at most $18\sqrt{n}(\log n)^2$. Plugging in the estimate from Proposition 3.1, we have that

$$|L_{\overline{\text{runsort}}(\pi)} - nye^{1-y}| = O(\sqrt{n}(\log n)^2)$$

with probability at least $1 - 2e^{-(\log n)^2}$. Since $\text{runsort}(\pi) = \overline{\text{runsort}}(\pi)$ with probability at least $1 - n^{-(\log \log n)/2}$ (by Lemma 2.1), we see that in fact $|L_{\overline{\text{runsort}}(\pi)} - nye^{1-y}| = O(\sqrt{n}(\log n)^2)$ holds with probability at least

$$1 - 2e^{-(\log n)^2} - n^{-(\log \log n)/2} > 1 - 2n^{-(\log \log n)/2}. \quad \square$$

When the entry $j = yn$ is the beginning of a run of π , the position of j in $\text{runsort}(\pi)$ is precisely $L_{\text{runsort}(\pi)}(y)$. Thus, the previous lemma tells us that in the scaled plot of $\text{runsort}(\pi)$, the beginnings of runs cluster around the curve \mathcal{C} . To make this precise, let $q_n(i, j)$ denote the probability that the entry j is the beginning of a run of π and is in the i -th position of $\text{runsort}(\pi)$ when $\pi \in S_n$ is chosen uniformly at random. (Here, $q_n(i, j)$ differs from $p_n(i, j)$ in that the former looks at only the first entry of each run and the latter looks at all entries.) We obtain an estimate on the distribution of $q_n(i, j)$ for fixed j .

Lemma 3.4. *There exists a constant $C > 0$ such that the following holds for all $y \in [0, 1]$:*

$$\left| \sum_{|i - nye^{1-y}| \leq C\sqrt{n}(\log n)^2} q_n(i, yn) - (1 - y + 1/n) \right| \leq 2n^{-(\log \log n)/2}$$

and

$$\sum_{|i - nye^{1-y}| > C\sqrt{n}(\log n)^2} q_n(i, yn) \leq 2n^{-(\log \log n)/2}.$$

Proof. Recall that the entry yn is the beginning of a run of π with probability $(n - yn + 1)/n$. This implies that $\sum_{i \in [n]} q_n(i, yn) = 1 - y + 1/n$. Both statements now follow from Lemma 3.3. \square

Let us summarize in words what this lemma tells us about the contribution to the measures \mathbf{R}_n (and eventually also \mathbf{R}) from the beginnings of runs: This contribution is concentrated close to the curve \mathcal{C} , and the “weighting” in the y -direction is $(1 - y) dy$. We will make these observations precise in Section 5.

Finally, we record a version of this result that will be convenient in the next section.

Lemma 3.5. *There exists a constant $C > 0$ such that the following holds for all $y \in [0, 1]$: If $\ell > C\sqrt{n}(\log n)^2$ and $k \geq yn$ are integers, then $q_n(nye^{1-y} - \ell, k) \leq 2n^{-(\log \log n)/2}$.*

Proof. This follows from Lemma 3.3 and the observation that $y \leq y'$ implies that $L_{\text{runsort}(\pi)}(y) \leq L_{\text{runsort}(\pi)}(y')$ for every π . \square

4. THE INTERIOR DENSITY

We now compute the limiting density for the non-singular part of the permuton. Suppose that π is a random permutation of length n . Recall that $p_n(i, j)$ denotes the probability that the entry j is in the i -th position of $\text{runsort}(\pi)$. In the previous section, we analyzed the case where the entry j is the beginning of a run of π , so we focus on the remaining case: Let $p'_n(i, j) = p_n(i, j) - q_n(i, j)$ denote the probability that the entry j is in the i -th position of $\text{runsort}(\pi)$ and is not the beginning of a run of π . We will be interested in the situation where $i = xn$ and $j = yn$ for fixed $x, y \in (0, 1)$, where the point (x, y) lies in \mathcal{C}^+ (i.e., strictly above \mathcal{C}) and n tends to infinity. By the results of the previous section, we know that $p'_n(i, j)$ is very close to $p_n(i, j)$ in this regime. Whenever we use xn or yn as an input for a function that takes integer values (such as p_n , p'_n , or q_n), we really mean $\lceil xn \rceil$ and $\lceil yn \rceil$; we simply omit the ceiling symbols to avoid a typhoon of ceiling symbols.

Fix n and $i, j \in [n]$ with $i > 1$, and suppose that we obtain π by first picking a random permutation π' on $[n] \setminus \{j\}$ and then inserting the entry j in a random position. Let k denote the entry in the $(i - 1)$ -th position of $\text{runsort}(\pi)$. The following two descriptions define the same event:

- The entry j is in the i -th position of $\text{runsort}(\pi)$ and is not the beginning of a run of π .
- $k < j$, and the entry j was inserted directly after the entry k in π .

The probability of the first bullet point occurring is (by definition) $p'_n(i, j)$, and the probability of the second bullet point occurring is

$$\frac{1}{n} \sum_{1 \leq k < j} p_{n-1}(i-1, k) = \frac{1}{n} \left(1 - \sum_{j \leq k \leq n-1} p_{n-1}(i-1, k) \right).$$

Putting these together, we find that

$$(2) \quad p'_n(i, j) = \frac{1}{n} \left(1 - \sum_{j \leq k \leq n-1} p_{n-1}(i-1, k) \right).$$

Since $p_n(i, j)$ and $p'_n(i, j)$ are very close, we incur a very small error if we replace $p_{n-1}(i-1, k)$ with $p'_{n-1}(i-1, k)$ on the right-hand side; we will address this carefully below. So, up to this small error, $p'_n(i, j)$ is equal to

$$\frac{1}{n} \left(1 - \sum_{j \leq k \leq n-1} p'_{n-1}(i-1, k) \right).$$

Note that we have the boundary conditions $p'_n(1, j) = 0$ for all j and $p'_n(i, n) = 1/n$ for $1 < i < n$. We now define the quantities $\tilde{p}_n(i, j)$ recursively via

$$\tilde{p}_n(i, j) = \frac{1}{n} \left(1 - \sum_{j \leq k \leq n-1} \tilde{p}_{n-1}(i-1, k) \right),$$

with the same boundary conditions $\tilde{p}_n(1, j) = 0$ and $\tilde{p}_n(i, n) = 1/n$ for $1 < i < n$. We will see that $p'_n(xn, yn)$ and $\tilde{p}_n(xn, yn)$ are very close as long as n is sufficiently large (with respect to (x, y)).

By repeatedly applying the recurrence relation for \tilde{p} , we can express $\tilde{p}_n(i, j)$ as a sum of terms involving $\tilde{p}_{n-i+1}(1, k)$ (for $j \leq k \leq n-i+1$) and $\tilde{p}_{n-r}(i-r, n-r)$ (for $i-r > 1$ and $j \leq k \leq n-r$), together with some constant terms:

- The terms $\tilde{p}_{n-i+1}(1, k)$ vanish by our boundary conditions.
- Each term $\tilde{p}_{n-r}(i-r, n-r)$ is equal to $1/(n-r)$ (by our boundary conditions) and appears $\binom{n-j+r-1}{r-1}$ times (by stars and bars), always weighted by $\frac{1}{n(n-1)\cdots(n-r+1)}$ and carrying the sign $(-1)^r$. Here, r ranges from 1 to $\min\{n-j, i-2\}$; call the latter quantity M .
- Each constant term $\frac{(-1)^s}{n(n-1)\cdots(n-s)}$ appears $\binom{n-j}{s}$ times, where s ranges from 0 to M .

Putting everything together, we arrive at the explicit formula

$$\tilde{p}_n(i, j) = \sum_{r=1}^M \frac{(-1)^r}{n(n-1)\cdots(n-r)} \binom{n-j+r-1}{r-1} + \sum_{s=0}^M \frac{(-1)^s}{n(n-1)\cdots(n-s)} \binom{n-j}{s}.$$

We remark that the only dependence on i is contained in the value of M . In the regime $i \geq n-j+2$, we see that M (and hence also $\tilde{p}_n(i, j)$) is completely independent of i ; this is a hint of the horizontal uniformity alluded to in the Introduction. For $i = xn$ and $j = yn$ with n tending to infinity, the second sum (call it S_2) will contribute the main term, and the first sum (call it S_1) will contribute a negligible error. Note that in this setting, M is asymptotically a positive constant multiple of n (since $x, y \in (0, 1)$).

We begin with the first sum. Expanding the binomial coefficient gives

$$S_1 = \sum_{r=1}^M \frac{(-1)^r}{(r-1)!} \cdot \frac{1}{n(n-1)} \cdot \frac{(n-j+r-1)(n-j+r-2)\cdots(n-j+1)}{(n-2)(n-3)\cdots(n-r)}.$$

Writing $n-j = (1-y)n$, we estimate

$$\frac{n-j+t}{n-r+t-1} = 1 - \frac{j}{n} + \frac{j(t-r-1) + n(r-1)}{n(n-r+t-1)} = (1-y) + O(r/n)$$

uniformly in n, r, t (so the implied constant depends only on (x, y)). Then

$$\begin{aligned} S_1 &= \frac{1}{n(n-1)} \sum_{r=1}^M \frac{(-1)^r}{(r-1)!} ((1-y)^{r-1} + O(r2^{r-1}/n)) \\ &= \frac{-e^{y-1}}{n^2} + O(1/n^3), \end{aligned}$$

which is $o(1/n)$.

Performing the analogous computation for the second sum, we find that

$$\begin{aligned} S_2 &= \sum_{s=0}^M \frac{(-1)^s}{s!} \cdot \frac{1}{n} \cdot \frac{(n-j)(n-j-1)\cdots(n-j-s+1)}{(n-1)(n-2)\cdots(n-s)} \\ &= \frac{1}{n} \sum_{s=0}^M \frac{(-1)^s}{s!} ((1-y)^s + O((s+1)2^{s+1}/n)) \\ &= \frac{e^{y-1}}{n} + O(1/n^2). \end{aligned}$$

In summary, we have established the following proposition.

Proposition 4.1. *Fix $x, y \in (0, 1)$ such that $(x, y) \in \mathcal{C}^+$. Then*

$$\tilde{p}_n(xn, yn) = \frac{e^{y-1}}{n} + O(1/n^2),$$

where the implicit constant depends only on (x, y) .

It remains to bound the difference between \tilde{p} and p' . This consists of keeping track of the error terms $q = p - p'$ when we iterate the recurrence (2). At the ℓ -th stage, the number of terms $p_{n-\ell}$ is $\binom{(1-y)n}{\ell}$ (by stars and bars), and each such term is scaled (up to a sign) by $\frac{1}{n(n-1)\cdots(n-\ell+1)}$. Let Q be the maximum of the error terms q (which are certainly nonnegative). By the triangle inequality, we may ignore the signs of the errors, and we find that

$$|p'_n(xn, yn) - \tilde{p}_n(xn, yn)| \leq Q \sum_{\ell=1}^{(1-y)n} \binom{(1-y)n}{\ell} \frac{1}{n(n-1)\cdots(n-\ell+1)} \leq Q \sum_{\ell=1}^{(1-y)n} \frac{1}{\ell!} = O(Q).$$

We now bound Q .

Lemma 4.2. *Fix $x, y \in (0, 1)$ such that $(x, y) \in \mathcal{C}^+$. If n is sufficiently large (depending on (x, y)), then the following holds: For all $1 \leq \ell \leq \min\{xn - 1, (1-y)n\}$ and all $yn \leq k \leq n - \ell$, we have*

$$q_{n-\ell}(xn - \ell, k) \leq 2(yn)^{-(\log \log(yn))/2}.$$

Proof. We wish to apply Lemma 3.5 with n replaced by $n - \ell$. Let $\varepsilon = ye^{1-y} - x$ (which is strictly positive). First, we check that

$$k \geq yn \geq y(n - \ell).$$

Second, we wish to show that $xn - \ell < (n - \ell)ye^{1-y} - C\sqrt{n - \ell}(\log(n - \ell))^2$ (where C is the constant from Lemma 3.5); this inequality rearranges to

$$\varepsilon n + \ell(1 - ye^{1-y}) > C\sqrt{n - \ell}(\log(n - \ell))^2.$$

The term $1 - ye^{1-y}$ is nonnegative, and we see that the inequality holds as long as n is sufficiently large (depending on y and ε). So we can apply Lemma 3.5, which tells us that

$$q_{n-\ell}(xn - \ell, k) \leq 2(n - \ell)^{-(\log \log(n - \ell))/2}.$$

The right-hand side is an increasing function of ℓ , so the bound $\ell \leq (1 - y)n$ gives the desired inequality. \square

The previous lemma implies that $Q \leq 2(yn)^{-(\log \log(yn))/2}$ (which is certainly $O(1/n^2)$) for n sufficiently large, so we can deduce the main result of this section. (For $y = 1$, recall from above that $p'_n(xn, n) = 1/n$.)

Lemma 4.3. *Fix $x, y \in (0, 1]$ such that $(x, y) \in \mathcal{C}^+$. Then*

$$p'_n(xn, yn) = \frac{e^{y-1}}{n} + O(1/n^2),$$

where again the implied constant depends only on (x, y) .

Let us summarize in words what this lemma tells us about the contribution to the measures \mathbf{R}_n (and eventually also \mathbf{R}) from the entries that are not the beginnings of runs: In \mathcal{C}^+ , this contribution gives a density $e^{1-y} dx dy$ at the point (x, y) ; note that this is independent of x . We have said nothing about the contribution below \mathcal{C} ; that this contribution is 0 follows quickly from Lemma 3.3, but in fact we will give an alternative argument in the next section.

5. PUTTING EVERYTHING TOGETHER

We finally prove Theorem 2.5, which implies Theorem 1.1. The main idea is that we have already accounted for 100% of the mass of \mathbf{R} in our discussions in the previous two sections; this means that there will not be any mass below the curve \mathcal{C} and that we do not need to worry about additional contributions very close to \mathcal{C} from entries that are not the beginnings of runs.

For each axis-parallel rectangle $B \subseteq [0, 1]^2$, it follows from Lemmas 3.4 and 4.3 that

$$\liminf_{n \rightarrow \infty} \mathbf{R}_n(B) \geq \int_{B \cap \mathcal{C}^+} e^{y-1} dx dy + \int_{B \cap \mathcal{C}} (1-y) dy = \mathbf{R}(B).$$

Now fix an axis-parallel rectangle $B_1 \subseteq [0, 1]^2$. Choose a tiling of $[0, 1]^2$ by axis-parallel rectangles B_1, \dots, B_5 (such a tiling certainly exists). We have

$$1 = \limsup_{n \rightarrow \infty} \mathbf{R}_n([0, 1]^2) = \sum_{a=1}^5 \limsup_{n \rightarrow \infty} \mathbf{R}_n(B_a) \geq \sum_{a=1}^5 \liminf_{n \rightarrow \infty} \mathbf{R}_n(B_a) \geq \sum_{a=1}^5 \mathbf{R}(B_a) = \mathbf{R}([0, 1]^2) = 1.$$

These inequalities must all be equalities, so $\lim_{n \rightarrow \infty} \mathbf{R}_n(B_1)$ exists and equals $\mathbf{R}(B_1)$. As B_1 was arbitrary, this completes the proof of Theorem 2.5.

6. A GENERALIZED SETTING

In this brief concluding section, we mention a setting in which the ideas presented earlier—especially those concerning the concentration inequalities derived in Section 3—still hold.

The *standardization* of a sequence of n distinct integers is the permutation in S_n that has the same relative order as the sequence. For example, the standardization of 4917 is 2413. Let $\mathcal{F} \subseteq \bigcup_{n \geq 0} S_n$ be a family of permutations, and let $\mathcal{F}_n = \mathcal{F} \cap S_n$. Assume that every permutation obtained by taking the standardization of a prefix of a permutation in \mathcal{F} is also in \mathcal{F} . Let us also assume that \mathcal{F} contains the permutation $1 \in S_1$ and that there is some constant $c > 1$ such that $|\mathcal{F}_n| < n!/c^n$ for all sufficiently large n .

We can use the family \mathcal{F} to split an arbitrary permutation $\pi \in S_n$ into subsequences as follows. Let $\pi_{[a,b]}$ denote the subsequence of π consisting of entries in positions $a, a+1, \dots, b$, and let $\bar{\pi}_{[a,b]}$ be the standardization of $\pi_{[a,b]}$. Set $k_0 = 1$. Then let k_1 be the smallest integer that is greater than k_0 and satisfies $\bar{\pi}_{[k_0, k_1]} \notin \mathcal{F}$; we make the convention that $k_1 = n+1$ if $\pi \in \mathcal{F}$. If $k_1 \neq n+1$, let k_2 be the smallest integer that is greater than k_1 and satisfies $\bar{\pi}_{[k_1, k_2]} \notin \mathcal{F}$; we make the convention that $k_2 = n+1$ if $\bar{\pi}_{[k_1, n]} \in \mathcal{F}$. Continue defining integers k_i in this greedy fashion until reaching a step at which $k_r = n+1$. Note that $\bar{\pi}_{[k_0, k_1-1]}, \bar{\pi}_{[k_1, k_2-1]}, \dots, \bar{\pi}_{[k_{r-1}, k_r-1]}$ all belong to the family \mathcal{F} . Let us call the subsequences $\pi_{[k_0, k_1-1]}, \pi_{[k_1, k_2-1]}, \dots, \pi_{[k_{r-1}, k_r-1]}$ the \mathcal{F} -runs of π . Define \mathcal{F} -sort(π) to be the permutation obtained by sorting the \mathcal{F} -runs of π so that their minimal entries appear in increasing order. Note that \mathcal{F} -sort is the same as runsort when \mathcal{F} is the family of increasing permutations (consisting of one permutation of each length).

Suppose n is large. If we choose $\pi \in S_n$ uniformly at random, then the minimal entries of the \mathcal{F} -runs of π should concentrate along a certain curve after we apply \mathcal{F} -sort to π . It should be possible to make this concentration statement precise using the ideas from Section 3; however, determining exactly what the curve is could be very difficult. Here, we simply provide some images illustrating how this phenomenon might look in specific examples.

A *double descent* of a permutation $\pi \in S_n$ is an index $i \in \{2, \dots, n-1\}$ such that $\pi_{i-1} > \pi_i > \pi_{i+1}$. A *valley* of π is an index $i \in \{2, \dots, n-1\}$ such that $\pi_{i-1} > \pi_i < \pi_{i+1}$. Let $\mathcal{F}^{\text{ddes}}$ be the family of permutations with no double descents, and let \mathcal{F}^{val} be the family of permutations with no valleys. Figure 2 shows the scaled plots of $\mathcal{F}^{\text{ddes}}$ -sort(π) and \mathcal{F}^{val} -sort(π), where π is a permutation chosen uniformly at random from S_{50000} .




FIGURE 2. The scaled plots of $\mathcal{F}^{\text{ddes}}\text{-sort}(\pi)$ (left) and $\mathcal{F}^{\text{val}}\text{-sort}(\pi)$ (right), where π is a permutation chosen uniformly at random from S_{50000} .

One possibility for future research is the characterization of the limiting behavior of $\mathcal{F}\text{-sort}(\pi)$ ($\pi \in S_n$ chosen uniformly at random) for various specific choices of \mathcal{F} . Perhaps even more interesting would be the determination of more general properties such as the existence of a limiting permutation and conditions on \mathcal{F} that guarantee some form of “horizontal uniformity.” One could also consider the limiting behavior of $\text{runsort}(\pi)$ when $\pi \in S_n$ is chosen non-uniformly, e.g., according to the Mallows distribution.

ACKNOWLEDGEMENTS

The first author is supported in part by NSF grant DMS–1855464, BSF grant 2018267, and the Simons Foundation. The second author is supported by an NSF Graduate Research Fellowship (grant DGE–1656466) and a Fannie and John Hertz Foundation Fellowship. The third author is supported by an NSF Graduate Research Fellowship (grant DGE–2039656). We are grateful to Ryan Alweiss and Peter Winkler for helpful conversations.

REFERENCES

- [1] N. Alon and J. H. Spencer, *The Probabilistic Method, Fourth Edition*. Wiley (2016).
- [2] O. Angel, A. Holroyd, D. Romik, and B. Virág, Random sorting networks. *Adv. Math.*, **215** (2007), 839–868.
- [3] M. Atapour and N. Madras, Large deviations and ratio limit theorems for pattern-avoiding permutations. *Combin. Probab. Comput.*, **23** (2014), 161–200.
- [4] P. Alexandersson and O. Nabawanda, Peaks are preserved under run-sorting. *Enumer. Combin. Appl.*, **2** (2022).
- [5] D. Callan, Pattern avoidance in “flattened” partitions. *Discrete Math.*, **309** (2009), 4187–4191.
- [6] D. Dauvergne, The Archimedean limit of random sorting networks. *Preprint* arXiv:1802.08934 (2018).
- [7] D. Dauvergne and B. Virág, Circular support in random sorting networks. *Trans. Amer. Math. Soc.*, **373** (2020), 1529–1553.
- [8] T. Dokos and I. Pak, The expected shape of random doubly alternating Baxter permutations. *Online J. Anal. Comb.*, **9** (2014).
- [9] R. Glebov, A. Grzesik, T. Klímašová, and D. Král’, Finitely forcible graphons and permutons. *J. Combin. Theory Ser. B*, **110** (2015), 112–135.
- [10] C. Hoppen, Y. Kohayakawa, C. G. Moreira, B. Ráth, and R. M. Sampaio, Limits of permutation sequences. *J. Combin. Theory Ser. B*, **103** (2013), 93–113.
- [11] R. Kenyon, D. Král’, C. Radin, P. Winkler, Permutations with fixed pattern densities. *Random Structures Algorithms*, **56** (2020), 220–250.
- [12] T. Mansour and M. Shattuck, Pattern avoidance in flattened permutations. *Pure Math. Appl. (P.U.M.A.)*, **22** (2011), 75–86.
- [13] T. Mansour, M. Shattuck, and S. Wagner, Counting subwords in flattened partitions of sets. *Discrete Math.*, **338** (2015), 1989–2005.

- [14] B. Maurey, Construction de suites symétriques. *C. R. Math. Acad. Sci Paris, Sér. A*, **288** (1979), 679–681.
- [15] V. Milman and G. Schechtman, *Asymptotic Theory of Infinite Dimensional Normed Spaces*. Springer (2016).
- [16] S. Miner and I. Pak, The shape of random pattern-avoiding permutations. *Adv. Appl. Math.*, **55** (2014), 86–130.
- [17] O. Nabawanda, F. Rakotondrajao, and A. S. Bamunoba, Run distribution over flattened partitions. *J. Integer Seq.*, **23** (2020).

DEPARTMENT OF MATHEMATICS, PRINCETON UNIVERSITY, PRINCETON, NJ 08544, USA AND SCHOOLS OF MATHEMATICS AND COMPUTER SCIENCE, TEL AVIV UNIVERSITY, TEL AVIV 69978, ISRAEL
Email address: `nalon@math.princeton.edu`

DEPARTMENT OF MATHEMATICS, PRINCETON UNIVERSITY, PRINCETON, NJ 08540, USA
Email address: `cdefant@princeton.edu`

DEPARTMENT OF MATHEMATICS, PRINCETON UNIVERSITY, PRINCETON, NJ 08540, USA
Email address: `nkravitz@princeton.edu`