



Ranks, copulas, and permutons

R. Grübel¹ 

Received: 26 May 2022 / Accepted: 27 April 2023 / Published online: 12 May 2023
© The Author(s) 2023

Abstract

We review a recent development at the interface between discrete mathematics on one hand and probability theory and statistics on the other, specifically the use of Markov chains and their boundary theory in connection with the asymptotics of randomly growing permutations. Permutations connect total orders on a finite set, which leads to the use of a pattern frequencies. This view is closely related to classical concepts of nonparametric statistics. We give several applications and discuss related topics and research areas, in particular the treatment of other combinatorial families, the cycle view of permutations, and an approach via exchangeability.

Keywords Asymptotics · Boundary theory · Copulas · Exchangeability · Markov chains · Permutations · Pattern frequencies · Ranks

Mathematics Subject Classification 60C05 · 05A05 · 60J50 · 62G20

1 Introduction

Ranks are at the basis of nonparametric statistics, copulas are a standard tool for the description of dependencies, and permutons have recently been introduced as limit objects for large permutations. We survey some of the connections between these concepts, mainly from a probabilistic and statistical point of view. In particular, we explain the use of Markov chains and their boundary theory in the context of sequences of permutations. This approach has received some attention in connection with the asymptotics of other randomly growing discrete structures, such as graphs and trees.

The paper addresses a statistically educated audience. Consequently, we only mention Chapter 13 of van der Vaart (1998) in connection with ranks in statistics, and for copulas we refer to Nelson (2006). In Sect. 2 we recall some basic aspects of permu-

✉ R. Grübel
rgrubel@stochastik.uni-hannover.de

¹ Institute of Actuarial and Financial Mathematics, Leibniz Universität Hannover, Postfach 6009, 30060 Hannover, Germany

tations; see Bona (2004) for a textbook reference. In Sect. 3 we introduce permutons, which are two-dimensional copulas. The basic reference here is the seminal paper Hoppen et al. (2013). The construction of limits is discussed in some generality, again from a probabilistic point of view.

Permutons arise as limits if pattern frequencies are considered. In Sect. 4 we explain an approach to the augmentation of discrete spaces that is based on the construction of suitable Markov chains and the associated boundary theory. The latter can be seen as a discrete variant of classical potential theory, going back to Doob (1959); see Woess (2009) for a general textbook reference and Grübel (2013) for an elementary introduction closer to the present context. We show that the pattern frequency topology arises naturally if a specific Markov chain $(\Pi_n)_{n \in \mathbb{N}}$ is chosen. This approach also provides an almost sure construction, meaning that $\Pi_n \rightarrow \Pi_\infty$ ‘pathwise’ with probability one.

Section 5 collects some applications. Finally, in the last section, we sketch some related problems and research areas, where we use the reviewish character of the present paper as a license for explaining connections that seem to be known to researchers in the field, but not necessarily to an interested newcomer (which the author was, some years ago). It should be clear that we will have to (and do) resist the temptation of always giving the most general results.

Throughout, a specific data set will be used to illustrate the various concepts, and we occasionally point out analogies to statistical concepts.

2 Permutations

For a finite set A we write $\mathbb{S}(A)$ for the set of bijections $\pi : A \rightarrow A$ and abbreviate this to \mathbb{S}_n if $A = [n] := \{1, \dots, n\}$. An element π of \mathbb{S}_n can be represented in one-line notation or in standard cycle notation. For the first $\pi = (\pi(1), \pi(2), \dots, \pi(n))$ (or $\pi(1)\pi(2) \cdots \pi(n)$ if no ambiguities can arise) is simply the tuple of its values, and we note at this early stage that this notation presupposes a total (or linear) order on the base set A . The second one is somewhat closer to the notion of a bijection: Each element a of A will run through a finite cycle if $\pi \in \mathbb{S}(A)$ is applied repeatedly. In the special case $A = [n]$ this leads to a tuple of cycles, and standard notation means that these are individually ordered by putting the respective largest element first and then listing the cycles increasingly according to their first elements. The *transition lemma*, also known as *Foata’s correspondence*, provides a bijection between the two representations by erasing the brackets in the step from cycle to one-line notation, and, roughly, by inserting brackets in front of every increasing record in the one-line notation, such as in

$$\underline{5}87324\underline{9}16 \quad \longrightarrow \quad (5)(87324)(916). \quad (1)$$

Much of the classical material connecting permutations and probabilities refers to the cycle structure, typically under the assumption that the random permutation Π_n is chosen uniformly from \mathbb{S}_n , which we abbreviate to $\Pi_n \sim \text{unif}(\mathbb{S}_n)$. A standard problem, often discussed in elementary probability courses, concerns the probability p_n that Π_n

has at least one fixed point, i.e. a cycle of length 1, which can be evaluated with the inclusion–exclusion formula to $p_n = \sum_{k=1}^n (-1)^{k+1}/k!$. Further, the correspondence in (1) relates the number of cycles on the right to the number of increasing records on the left hand side. As a final example we mention that the distribution of the cycle type $C_n = (C_n(1), C_n(2), \dots)$ of Π_n , where $C_n(i)$ denotes the number of cycles of length i in Π_n , is given by

$$P(C_n(1) = c_1, C_n(2) = c_2, \dots) = \prod_{i=1}^{\infty} \frac{1}{i^{c_i} c_i!}, \tag{2}$$

where the values $c_i, i \in \mathbb{N}$, have to satisfy the obvious constraint that $\sum_{i \in \mathbb{N}} i c_i = n$. In particular, $C_n(i) = 0$ for $i > n$. Multiplication by $n!$ leads to a formula for the number of permutations of a given cycle type. From a probabilistic point of view, rewriting (2) as

$$\mathcal{L}((C_n(1), \dots, C_n(n))) = \mathcal{L}[(Z_1, \dots, Z_n) | T_n = n], \tag{3}$$

with Z_1, \dots, Z_n independent, Z_i Poisson distributed with parameter $1/i$, and $T_n := \sum_{i=1}^n i Z_i$, seems more instructive; for example, it provides access to the distributional asymptotics. We refer to Arratia et al. (2003) for this and related material. Our later aim, however, is to deal with limits for the permutations themselves, but we will return to the cycle view in Sect. 6.2.

We may also describe a permutation $\pi \in \mathbb{S}_n$ by its $n \times n$ permutation matrix $M(\pi)$ with entries $m_{ij} = 1$ if $\pi(j) = i$ and $m_{ij} = 0$ if $\pi(j) \neq i$. For the permutation on the left in (1) we get

$$M(\pi) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}. \tag{4}$$

Below we will return to the fact that the composition of permutations corresponds to matrix multiplication in this description.

Finally, permutations appear as order isomorphisms: If $<_1$ and $<_2$ are total orders on some finite set A , then there is a unique $\pi \in \mathbb{S}(A)$ such that

$$a <_1 b \iff \pi(a) <_2 \pi(b) \quad \text{for all } a, b \in A.$$

A one-line notation such as 587324916 in (1) defines a total order on the set of numerals, with 5 the smallest and 6 the largest element. With ‘ $<$ ’ the canonical order on $[n]$ (or \mathbb{R}) the permutation $\pi \in \mathbb{S}_n$ then provides an order isomorphism between $([n], <)$ and $([n], <_\pi)$, with ‘ $<_\pi$ ’ given by the one-line notation for π .

Suppose now that we have data $x_1, \dots, x_n \in \mathbb{R}$, with $x_i \neq x_j$ if $i \neq j$ (no ties). The rank of x_i in this set is given by $\text{rk}(x_i) = \text{rk}(x_i|x_1, \dots, x_n) = \#\{j \in [n] : x_j \leq x_i\}$, and these can be combined into a permutation $\pi_x := (\text{rk}(1), \dots, \text{rk}(n)) \in \mathbb{S}_n$. For two-dimensional data $(x_1, y_1), \dots, (x_n, y_n)$ with no ties in the two sets of component values we then obtain two permutations π_x, π_y that define two total orders on $[n]$. These are, in the above sense, connected by the permutation

$$\pi = \pi((x_1, y_1), \dots, (x_n, y_n)) = \pi_y \circ \pi_x^{-1}. \tag{5}$$

In this situation, the permutation plot for π is known as the *rank plot* for the data. The order-relating permutation does not depend on the order in which the observation vectors are numbered. Formally, it is invariant under permutations σ of $[n]$ in the sense that the permuted pairs $(x_{\sigma(1)}, y_{\sigma(1)}), \dots, (x_{\sigma(n)}, y_{\sigma(n)})$ lead to the same π . Hence, in statistical terms, π depends on the data only through their empirical distribution.

We close this section with a data example. The 16 largest cities in Germany may be ordered alphabetically, or by decreasing population size, or by location in the west-east or south-north direction. Starting with an ordering by size the corresponding geographic orderings are

$$\begin{aligned} \pi_{\text{long}} &= (15, 11, 13, 3, 7, 9, 2, 6, 4, 8, 16, 14, 10, 12, 1, 5), \\ \pi_{\text{lat}} &= (14, 16, 1, 5, 4, 2, 7, 12, 10, 15, 6, 8, 13, 3, 9, 11). \end{aligned}$$

For example, Berlin is fairly far in the east and fairly far up north within this group. The permutation relating these as total orders is given by

$$\pi_{\text{lat}} \circ \pi_{\text{long}}^{-1} = (9, 7, 5, 10, 11, 12, 4, 15, 2, 13, 16, 3, 1, 8, 14, 6). \tag{6}$$

Fig. 1 shows the true locations (in longitude and latitude degrees) on the left and the rank plot on the right hand side, with the four largest cities in red. The observant reader will have noticed that the rank plot is, essentially, the associated permutation matrix. In fact, the plot of a permutation also leads to an interpretation (or coding) of permutations that will be central for the permuton aspect: Writing δ_x for the one-point measure at x we can associate a distribution $\mu(\pi)$ on the unit square with $\pi \in \mathbb{S}_n$ via

$$\mu(\pi) := \frac{1}{n} \sum_{i=1}^n \delta_{(i/n, \pi(i)/n)}. \tag{7}$$

Clearly, π can be recovered from $\mu(\pi)$. Also, for all $\pi \in \mathbb{S}_n$, the marginals of $\mu(\pi)$ are the discrete uniform distributions on the set $\{1/n, 2/n, \dots, 1\}$.

In view of later developments we note that in the passage from rank plot to permutation plot the relative positions of the points in x - and y -direction do not change. In our data example, a city that is north (up) or east (right) of some other city will stay so if we move from the left to the right part of Fig. 1.




Fig. 1 Scatter plot and rank plot for the city data (red: the four largest cities)

3 Patterns, subsampling and convergence

Our aim in this section is to obtain a formal framework for the informal (and ubiquitous) question, ‘what happens as $n \rightarrow \infty$?’ We may, for example, start with a sequence $(\pi_n)_{n \in \mathbb{N}}$ of permutations with $\pi_n \in \mathbb{S}_n$ for all $n \in \mathbb{N}$ and regard these as elements of the set $\mathbb{S} := \bigsqcup_{n=1}^{\infty} \mathbb{S}_n$, which we then need to augment in order to obtain a limit. A standard procedure begins with a set \mathcal{F} of functions $f : \mathbb{S} \rightarrow [0, 1]$ that separates the points in the sense that, for each pair $\pi, \sigma \in \mathbb{S}$ with $\pi \neq \sigma$, there is a function $f \in \mathcal{F}$ such that $f(\pi) \neq f(\sigma)$. We then identify $\pi \in \mathbb{S}$ with the function $f \mapsto f(\pi)$, which is an element of the space $[0, 1]^{\mathcal{F}}$ which, endowed with the product topology, is compact by Tychonov’s theorem. The closure of the embedded \mathbb{S} provides a compactification $\tilde{\mathbb{S}}$ of \mathbb{S} , where convergence of a sequence $(\pi_n)_{n \in \mathbb{N}}$ is equivalent to the convergence of all real sequences $(f(\pi_n))_{n \in \mathbb{N}}$, $f \in \mathcal{F}$. Moreover, each function $f \in \mathcal{F}$ has a continuous extension to $\tilde{\mathbb{S}}$. In the present situation \mathcal{F} will always be countable, and we start with the discrete topology on \mathbb{S} . Then $\tilde{\mathbb{S}}$ is a separable compact topological space, and the boundary $\partial\mathbb{S}$ is a compact subset.

We discuss three function classes and the resulting topologies.

First, let \mathcal{F} be the set of indicator functions $f_{\sigma} = 1_{\{\sigma\}}$, $\sigma \in \mathbb{S}$, i.e. $f_{\sigma}(\sigma) = 1$ and $f_{\sigma}(\pi) = 0$ if $\pi \neq \sigma$. Then, if $(\pi_n)_{n \in \mathbb{N}}$ is a sequence that leaves \mathbb{S} in the sense that for each $k \in \mathbb{N}$ there is an $n_0 = n_0(k) \in \mathbb{N}$ such that $\pi_n \notin \mathbb{S}_k$ for all $n \geq n_0$, we always get the function which is equal to 0 as the pointwise limit, hence the boundary consists of a single point. This topology is known as the *one-point compactification*.

For the second example we note first that order relations transfer to subsets by restriction. For $A = [n]$, $\pi \in \mathbb{S}_n$, and $B = [k]$ with $k \leq n$, restricting $<_{\pi}$ from A to B amounts to deleting all entries greater than k in the one-line notation of π , which gives the one-line notation for a permutation $\sigma \in \mathbb{S}_k$. We generally write $\sigma = \pi[B]$ if $(B, <_{\pi})$ is order isomorphic to $([k], <_{\sigma})$. If $B = \{i_1, \dots, i_k\} \subset [n]$ with $i_1 < \dots < i_k$ this is equivalent to

Table 1 Patterns in the city data permutation

Permutation	123	132	213	231	312	321
Pattern frequency	69	57	88	61	130	50
Relative frequency	.152	.125	.193	.134	.286	.110

$$\sigma(j) < \sigma(m) \iff \pi(i_j) < \pi(i_m) \text{ for } j, m \in [k],$$

a condition that will appear repeatedly below. We now define $f_\sigma : \mathbb{S} \rightarrow [0, 1], \sigma \in \mathbb{S}_k$, by $f_\sigma(\pi) = 1$ if $\sigma = \pi[A]$ for $A = [k]$, and $f_\sigma(\pi) = 0$ otherwise. As in the previous case the functions take only the values 0 and 1, so that convergence means that $f_\sigma(\pi_n)$ remains constant from some $n_0 = n_0(\sigma)$ onwards. In view of $\sum_{\sigma \in \mathbb{S}_k} f_\sigma(\pi) = 1$ for all $\pi \in \mathbb{S}_n, n \geq k$, any such sequence leads to a sequence $(\sigma_k)_{k \in \mathbb{N}}, \sigma_k \in \mathbb{S}_k$, that has the property that $f_{\sigma_k}(\pi_n) = 1$ for all $n \geq n_0(k)$. Clearly, this sequence must be projective in the sense that the restriction of σ_k to $[l]$ with $l \leq k$ is equal to σ_l . This implies that there is a total order $<_{\pi_\infty}$ on \mathbb{N} which reduces to $<_{\sigma_k}$ on $[k]$ for all $k \in \mathbb{N}$. The result is known as the *projective topology*, and the set of limits is the set of total orders on \mathbb{N} .

The third topology will lead to permutons. It arises if we use *pattern counting*, another aspect that is closely related to the interpretation of permutations as connecting total orders. Let $M(k, n)$ be the set of strictly increasing functions $f : [k] \rightarrow [n], n \geq k$. We define the *relative frequency of the pattern σ in the permutation π* as

$$t(\sigma, \pi) = \frac{1}{\#M(k, n)} \#\{f \in M(k, n) : \sigma(i) < \sigma(j) \Leftrightarrow \pi(f(i)) < \pi(f(j))\}.$$

Of course, $\#M(k, n) = \binom{n}{k}$. We augment this by putting $t(\sigma, \pi) = 0$ if $k > n$. Using the notation from the previous paragraph and identifying $f \in M(n, k)$ with its range A we see that the numerator is the number of size k subsets A of $[n]$ with $\pi[A] = \sigma$. Table 1 contains the absolute and relative frequencies for the six patterns of length three in the city data introduced at the end of Sect. 1. The four largest cities generate the pattern $\sigma = 2413 \in \mathbb{S}_4$.

From a probabilistic or, in fact, statistical point of view, we may regard the relative frequencies as probabilities in a sampling experiment. With $\xi = \xi_{n,k}$ uniformly distributed on the set of all subsets A of $[n]$ with size k , the function $\sigma \mapsto t(\sigma, \pi)$ is the probability mass function of the random element $\pi[\xi]$ of \mathbb{S}_k . In particular, whenever $k \leq n$,

$$\sum_{\sigma \in \mathbb{S}_k} t(\sigma, \pi) = 1 \text{ for all } \pi \in \mathbb{S}_n. \tag{8}$$

Building on an earlier approach in graph theory, Hoppen et al. (2013) introduced a notion of convergence based on such pattern frequencies. We identify a permutation π with its function $\sigma \rightarrow t(\sigma, \pi)$ on the set \mathbb{S} and with values in the unit interval; it is easy to see that different permutations lead to different functions. The general approach outlined above then leads us to define convergence of the sequence $(\pi_n)_{n \in \mathbb{N}}$ as convergence of $t(\sigma, \pi_n)$ for all $\sigma \in \mathbb{S}$. Below we will refer to this as the *pattern frequency topology*,

For the construction to be of use we need a description of the boundary $\partial\mathbb{S} = \bar{\mathbb{S}} \setminus \mathbb{S}$. In the general setup the limit of a convergent sequence $(\pi_n)_{n \in \mathbb{N}}$ with $|\pi_n| \rightarrow \infty$ is simply the function

$$\pi_\infty : \mathbb{S} \rightarrow [0, 1], \quad \sigma \mapsto \lim_{n \rightarrow \infty} t(\sigma, \pi_n),$$

but obviously, not all functions from \mathbb{S} to $[0, 1]$ appear in that manner. For example, (8) implies that $\sum_{\sigma \in \mathbb{S}_k} \pi_\infty(\sigma) = 1$ for any fixed $k \in \mathbb{N}$, and a conditioning argument leads to the inequality $\pi_\infty(\sigma) \geq t(\sigma, \rho) \cdot \pi_\infty(\rho)$ for all $\sigma \in \mathbb{S}_k, \rho \in \mathbb{S}_n$ with $k \leq n$. In fact, we even have, for all $k, l, m \in \mathbb{N}$ with $k < l < m$ and all $\sigma \in \mathbb{S}_k, \rho \in \mathbb{S}_m$,

$$t(\sigma, \rho) = \sum_{\tau \in \mathbb{S}_l} t(\sigma, \tau) t(\tau, \rho). \tag{9}$$

This follows from viewing the sampling procedure as removing single elements repeatedly, and then decomposing with respect to the permutation obtained at the intermediate level. The decomposition (9) will play an important role later in the context of the dynamics of permutation sequences; see Sect. 4.

Now copulas enter the stage: A (two-dimensional) *copula* $C : [0, 1] \times [0, 1] \rightarrow [0, 1]$ is the distribution function of a probability measure on the unit square that has uniform marginals, meaning that $C(u, 1) = u$ for all $u \in [0, 1]$ and $C(1, v) = v$ for all $v \in [0, 1]$. We define the associated *pattern frequency function* $t(\cdot, C)$ by sampling: For $\sigma \in \mathbb{S}_k$ let $t(\sigma, C)$ be the probability that a sample $(X_1, Y_1), \dots, (X_k, Y_k)$ from C leads to the permutation σ in the sense that

$$\Pi_n((X_1, Y_1), \dots, (X_k, Y_k)) = \sigma,$$

see also (5). As the joint distribution of the sample is invariant under permutations it follows that, for all $k \in \mathbb{N}$ and $\sigma \in \mathbb{S}_k$,

$$t(\sigma, C) = k! P(X_1 < X_2 < \dots < X_k, Y_{\sigma(1)} < Y_{\sigma(2)} < \dots < Y_{\sigma(k)}).$$

We observe (but do not prove) that C is determined by the function $\sigma \mapsto t(\sigma, C)$ on \mathbb{S} , and that for copulas, weak convergence is equivalent to the pointwise convergence on \mathbb{S} of these functions. Note the analogy to the moment problem in the first and to the moment method in the second statement.

We can now collect and rephrase some of the main results in Hoppen et al. (2013). For the definition of the random permutations in part (c) we refer to (5) again, and to the description given there of a suitable algorithm for obtaining the permutation from the data. Ties in the X - resp. Y -values may be ignored as these are, individually, a sample from the uniform distribution on the unit interval.

Theorem 1 (a) *A sequence $(\pi_n)_{n \in \mathbb{N}}$ of permutations in \mathbb{S} with $|\pi_n| \rightarrow \infty$ converges in the pattern frequency topology if and only if there is a copula C such that*

$$\lim_{n \rightarrow \infty} t(\sigma, \pi_n) = t(\sigma, C) \quad \text{for all } \sigma \in \mathbb{S}. \tag{10}$$

Moreover, the copula C is unique.

(b) In the pattern frequency topology, the boundary $\partial\mathbb{S}$ is homeomorphic to the space of copulas, endowed with weak convergence.

(c) Let C be a copula and suppose that (X_i, Y_i) , $i \in \mathbb{N}$, are independent random vectors with distribution function C . For each $n \in \mathbb{N}$ let Π_n be the random permutation associated with the first n pairs $(X_1, Y_1), \dots, (X_n, Y_n)$. Then, in the pattern frequency topology, Π_n converges to C almost surely as $n \rightarrow \infty$.

Below we will refer the stochastic process $(\Pi_n)_{n \in \mathbb{N}}$ in part (c) of Theorem 1 as the permutation sequence generated (or parametrized) by the copula C .

From a probabilistic point of view it is not too difficult to understand this result: Let $\mathbb{M}_1 = \mathbb{M}_1([0, 1] \times [0, 1], \mathcal{B}([0, 1] \times [0, 1]))$ be the space of probability measures on the Borel subsets of the unit square, endowed with the topology of weak convergence, and consider the embedding $\pi \mapsto \mu(\pi)$ of \mathbb{S} into \mathbb{M}_1 given by (7). Suppose now that $(\pi_n)_{n \in \mathbb{N}}$ is a sequence of permutations where we assume for simplicity that $\pi_n \in \mathbb{S}_n$ for all $n \in \mathbb{N}$. As \mathbb{M}_1 is compact with respect to weak convergence this sequence has a limit point μ_∞ , and convergence would follow from the uniqueness of μ_∞ . For this we first note that the marginals of any such limit point are the uniform distributions on the unit interval, as weak convergence implies the convergence of the marginal distributions. Hence any limit point μ_∞ has as its distribution function a copula C . Next we argue that (10) holds: Let $\sigma \in \mathbb{S}_k$ be given and let the function h_σ on $([0, 1] \times [0, 1])^k$ indicate whether or not the order-relating permutation for the marginal rank vectors of $((x_1, y_1), \dots, (x_k, y_k))$ is equal to σ . Then

$$t(\sigma, \pi_n) = \frac{1}{\binom{n}{k}} \sum_{1 \leq i_1 < \dots < i_k \leq n} h_\sigma((x_{i_1}, y_{i_1}), \dots, (x_{i_k}, y_{i_k})). \tag{11}$$

Uniform marginals implies that the set of discontinuities of h_σ is a μ_∞ -null set, and the weak convergence of $\mu(\pi_n)$ to μ_∞ implies the weak convergence of the respective k th measure-theoretic powers. Taken together this shows (10). The desired uniqueness now follows from the fact that a copula C is determined by its function $\sigma \rightarrow t(\sigma, C)$. For (b) we use that weak convergence of copulas is equivalent to pointwise convergence of the associated functions on \mathbb{S} . Finally, part (c) is immediate from the representation (11) of pattern frequencies as U -statistics, and the classical consistency result for the latter.

4 Dynamics

Generating random permutations is an important applied task. In line with the history of probability we recall a standard gaming example: Card mixing requires $\Pi_n \sim \text{unif}(\mathbb{S}_n)$ for some fixed n . Many shuffling algorithms may be modeled as random walks on \mathbb{S}_n and can then be analyzed using the representation theory for this group; see also Sect. 6.2 below.

Here we are interested in sequences $(\Pi_n)_{n \in \mathbb{N}}$ with $\Pi_n \in \mathbb{S}_n$ for all $n \in \mathbb{N}$, and a Markovian dynamic. A classical example is the Markov chain $\Pi^F = (\Pi_n^F)_{n \in \mathbb{N}}$ which starts at $\Pi_1^F \equiv (1)$ and then uses independent random variables $I_n \sim \text{unif}([n + 1])$,

independent of Π_1^F, \dots, Π_n^F , to construct Π_{n+1}^F from $\Pi_n^F = (\pi(1), \dots, \pi(n))$ as follows,

$$\Pi_{n+1}^F = (\pi(1), \dots, \pi(I_n - 1), n + 1, \pi(I_n), \dots, \pi(n)). \tag{12}$$

In words: We insert the new value $n + 1$ at a position in a gap chosen uniformly at random. This appears, for example, on p.132 in Feller (1968). A second possibility, which we will relate to Hoppen et al. (2013), and which we denote by $\Pi^H = (\Pi_n^H)_{n \in \mathbb{N}}$, uses ‘double insertion’ instead. This means that the step from n to $n + 1$ is based on two independent random variables $I = I_n$ and $J = J_n$, both uniformly distributed on $[n + 1]$, and

$$\Pi_{n+1}^H(i) = \begin{cases} \Pi_n^H(i), & \text{if } i < I \text{ and } \Pi_n^H(i) < J, \\ \Pi_n^H(i) + 1, & \text{if } i < I \text{ and } \Pi_n^H(i) \geq J, \\ J, & \text{if } i = I, \\ \Pi_n^H(i - 1), & \text{if } i > I \text{ and } \Pi_n^H(i) < J, \\ \Pi_n^H(i - 1) + 1, & \text{if } i > I \text{ and } \Pi_n^H(i) \geq J. \end{cases} \tag{13}$$

In words: We insert J at position I and increase the appropriate values by 1. Obviously, this process is a Markov chain, again with start at the single element $\sigma = (1)$ of \mathbb{S}_1 . We collect some properties of this process. The notion of a permutation sequence generated by a copula is explained in Theorem 1 (c).

Lemma 1 *Let Π^H be as defined above.*

(a) Π^H is equal in distribution to the permutation sequence generated by the independence copula $C(x, y) = x \cdot y$.

(b) The transition probabilities of Π^H are given by

$$p^H(\sigma, \tau) := P(\Pi_n^H = \tau | \Pi_k^H = \sigma) = \frac{k!}{n!} t(\sigma, \tau) \tag{14}$$

for all $k, n \in \mathbb{N}$ with $k \leq n$, and all $\sigma \in \mathbb{S}_k$ and $\tau \in \mathbb{S}_n$.

(c) $\Pi_n^H \sim \text{unif}(\mathbb{S}_n)$ for all $n \in \mathbb{N}$.

Proof (a) Let $\Pi = (\Pi_n)_{n \in \mathbb{N}}$ be generated by the independence copula. Then the step from Π_n to Π_{n+1} is based on the component ranks I and J of the next pair (X_{n+1}, Y_{n+1}) in the sequence $(X_1, Y_1), \dots, (X_{n+1}, Y_{n+1})$. It is well known that I and J are both uniformly distributed on $[n + 1]$. Here they are also independent, hence Π_{n+1} arises from Π_n as in (13).

(b) The transition from $\sigma \in \mathbb{S}_n$ to $\tau \in \mathbb{S}_{n+1}$ is based on independent random variables I and J , both uniformly distributed on $[n + 1]$. If $I = i$ then, with $\phi_i : [n] \rightarrow [n + 1]$, $\phi_i(k) = k$ if $k \leq i$ and $\phi_i(k) = k + 1$ if $k > i$, we must have $\sigma(l) < \sigma(m)$ if and only if $\tau(\phi_i(l)) < \tau(\phi_i(m))$ for all $l, m \in [n]$. The number of such i ’s is $(n + 1) t(\sigma, \tau)$, and in each case we additionally need $J = \tau(I)$. This proves the one-step version of (14).

To obtain the general case we now use induction, the Markov property, and (9).

(c) This is immediate from (14) with $k = 1$ and $\sigma = 1$. □




Fig. 2 The transition graphs for the Markov chains Π^F (top) and Π^H (bottom)

Hence both Π^F and Π^H consist of uniformly distributed variables. A trivial third possibility to achieve this is to use independent Π_n 's with $\Pi_n \sim \text{unif}(\mathbb{S}_n)$. The *transition graph* for these chains has \mathbb{S} as its set of nodes and an edge $\{\sigma, \tau\}$ for each pair $\sigma \in \mathbb{S}_n, \tau \in \mathbb{S}_{n+1}, n \in \mathbb{N}$, with the property that $P(\Pi_{n+1} = \tau | \Pi_n = \sigma) > 0$. Figure 2 shows the first four levels of the transition graphs for Π^F and Π^H . These graphs are the *Hasse diagrams* associated with two specific partial orders on \mathbb{S} . In the first, $\sigma <_F \rho$ for $\sigma \in \mathbb{S}_k$ and $\rho \in \mathbb{S}_n$ with $k < n$ means that the one-line notation for σ arises from the one-line notation for ρ by simply deleting all values greater than k in the latter, whereas in connection with Π^H we have $\sigma <_H \rho$ if and only if σ appears as a pattern in ρ . Clearly, the graph for Π^F is a tree: All previous states are deterministic functions of the current state. Informally speaking, all three chains eventually leave the state space \mathbb{S} , and this is what provides the connection to the previous section.

We need an excursion. The Markov property is often stated informally as ‘the future development depends on the history of the process only via its present state’. It may be instructive to carry out a simple calculation: If X_1, X_2, \dots, X_n are random variables on some probability space with values in some countable set, then their joint distribution can always be evaluated step by step via

$$\begin{aligned}
 P(X_1 = x_1, \dots, X_n = x_n) &= P(X_1 = x_1) \cdot \prod_{i=2}^n P(X_i = x_i | X_{i-1} \\
 &= x_{i-1}, \dots, X_1 = x_1),
 \end{aligned}$$

whenever the probability on the left hand side is positive. If the tuple has the Markov property then elementary manipulations of conditional probabilities lead to

$$P(X_1 = x_1, \dots, X_n = x_n) = P(X_1 = x_1) \cdot \prod_{i=2}^n P(X_i = x_i | X_{i-1} = x_{i-1})$$

$$\begin{aligned}
 &= P(X_1 = x_1) \cdot \prod_{i=2}^n P(X_{i-1} = x_{i-1} | X_i = x_i) \frac{P(X_i = x_i)}{P(X_{i-1} = x_{i-1})} \\
 &= P(X_n = x_n) \prod_{i=n-1}^1 P(X_i = x_i | X_{i+1} = x_{i+1}).
 \end{aligned}$$

Hence the Markov property then also holds ‘in the other direction’. Moreover, the converse statement follows on reading the above from bottom to top. For a sequence $(X_n)_{n \in \mathbb{N}}$, and with $\mathcal{G}_n, \mathcal{H}_n$ the σ -fields generated by the variables X_1, \dots, X_n and X_n, X_{n+1}, \dots respectively, the Markov property can formally be written as $\mathcal{L}(X_{n+1} | \mathcal{G}_n) = \mathcal{L}(X_{n+1} | X_n), n \in \mathbb{N}$, and the above argument shows that this implies $\mathcal{L}(X_{n-1} | \mathcal{H}_n) = \mathcal{L}(X_{n-1} | X_n), n \in \mathbb{N}$, a procedure that may be interpreted as *time reversal*. In fact, a somewhat neutral statement would be that, for a Markov process, past and future are conditionally independent, given the present.

The distribution of a Markov chain is specified by the distribution of the first variable and the (forward) transition probabilities; for the reverse chain we speak of the *cotransition probabilities*. The tree structure of its transition graph implies that for Π^F the cotransitions are degenerate in the sense only the values 0 and 1 occur. For Π^H we use Lemma 1 to obtain

$$P(\Pi_n^H = \sigma | \Pi_{n+1}^H = \tau) = \frac{P(\Pi_n^H = \sigma)}{P(\Pi_{n+1}^H = \tau)} p^H(\sigma, \tau) = t(\sigma, \tau) \tag{15}$$

for all $n \in \mathbb{N}, \sigma \in \mathbb{S}_n$ and $\tau \in \mathbb{S}_{n+1}$. The decomposition (9) can now be interpreted as describing the backwards movement of Π^H : *Pattern frequencies arise as cotransitions for this chain*.

Next we sketch Markov chain boundary theory, which grew out of the historically very important probabilistic approach to classical potential theory. For the highly transient chains considered here, where no state can be visited twice and where the time parameter is a function of the state and thus time homogeneity is implicit, this takes on a particularly simple form. Suppose that $X = (X_n)_{n \in \mathbb{N}}$ is a Markov chain with graded state space $S = \bigsqcup S_n$, that X is adapted to the grading in the sense that $X_n \in S_n$ for all $n \in \mathbb{N}$, and that S is minimal in the sense that $P(X_n = x) > 0$ for all $x \in S_n, n \in \mathbb{N}$. The *Martin kernel* $K : S \times S \rightarrow \mathbb{R}_+$ for X is then given by

$$K(x, y) := \frac{P(X_n = y | X_m = x)}{P(X_n = y)} \text{ for all } x \in S_m, y \in S_n \tag{16}$$

if $n \geq m$, and $K(x, y) = 0$ otherwise. For each $x \in S_m$ let $f_x : S \rightarrow [0, 1]$ be defined by

$$f_x(y) = P(X_m = x) K(x, y) = P(X_m = x | X_n = y) \tag{17}$$

and let $\mathcal{F} := \{f_x : x \in S\}$. The *Doob-Martin compactification* \bar{S} and the *Martin boundary* ∂S of the state space are the results of the procedure outlined at the beginning of Sect. 3, with this choice of \mathcal{F} as space of separating functions.

We give two central results.

Theorem 2 *In the Doob-Martin compactification, X_n converges almost surely to a random variable X_∞ with values in the Martin boundary.*

Proof Let $f = f_x \in \mathcal{F}$ with some $x \in S_k$. In view of the above remarks on the Markov property and time reversal,

$$f(X_n) = P(X_k = x|X_n) = E[1_{\{x\}}(X_k)|\mathcal{H}_n] \quad \text{for all } n \geq k,$$

so that, as $\mathcal{H}_{n+1} \subset \mathcal{H}_n$,

$$E[f(X_n)|\mathcal{H}_{n+1}] = E[E[1_{\{x\}}(X_k)|\mathcal{H}_n]|\mathcal{H}_{n+1}] = E[1_{\{x\}}(X_k)|\mathcal{H}_{n+1}] = f(X_{n+1}).$$

This shows that $(f(X_n), \mathcal{H}_n)_{n \geq k}$ is a reverse martingale, hence $f(X_n)$ converges almost surely. As $x \in S$ was arbitrary, the first statement follows by construction of the compactification.

As $P(X_\infty \in \bigcup_{k=1}^n S_k) = 0$ for all $n \in \mathbb{N}$, we also get $P(X_\infty \in \partial S) = 1$. □

The general compactification procedure implies that the Martin kernel can be extended continuously from $S \times S$ to $S \times \bar{S}$ via $K(x, \alpha) = \lim_{n \in \mathbb{N}} K(x, y_n)$ where $(y_n)_{n \in \mathbb{N}}$ is such that $y_n \rightarrow \alpha \in \partial S$ in the Doob-Martin topology. The following is a special case of *Doob's h-transform*.

Theorem 3 *Let $X, S, \bar{S}, \partial S$ be as above and let $p(x, y) = P(X_{n+1} = y|X_n = x)$, $x \in S_n, y \in S_{n+1}$, be the transition function of the Markov chain X . For $\alpha \in \partial S$ let p^α be defined by*

$$p^\alpha(x, y) = \frac{K(y, \alpha)}{K(x, \alpha)} p(x, y) \quad \text{for all } n \in \mathbb{N}, x \in S_n, y \in S_{n+1}.$$

Then p_α is the transition function for a Markov chain $X^\alpha = (X_n^\alpha)_{n \in \mathbb{N}}$, and X_n^α converges almost surely to α in the Doob-Martin topology.

The picture that emerges from these results is that of an *exit boundary* and, with $K(\cdot, \alpha)$ as basis for an *h-transform*, of the resulting transformed chain as the process conditioned on the value $\alpha \in \partial S$ for the limit X_∞ . Further, the relation to time reversal and the two-sidedness of the Markov property outlined above shows that convergence $y_n \rightarrow \alpha \in \partial S$ in the Doob-Martin topology is equivalent to the convergence in distribution of the initial segments of the conditioned chain,

$$\mathcal{L}[X_1, \dots, X_k|X_n = y_n] \rightarrow_{\text{distr}} \mathcal{L}[X_1, \dots, X_k|X_\infty = \alpha] \quad \text{for all } k \in \mathbb{N}.$$

Finally, we note that the topological construction depends on the chain via the associated cotransitions only. In the familiar forward case, the distribution of X is specified by the distribution of the first variable X_1 and the transition function p . Here we may consider the distribution of X_∞ as given, and then use the cotransitions to obtain the distribution of X as a mixture of the distributions $\mathcal{L}[X|X_\infty = \alpha], \alpha \in \partial S$.

We now return to processes $(\Pi_n)_{n \in \mathbb{N}}$ of permutations. This is a special case of the above framework, with $S = \mathbb{S}$ and $S_n = \mathbb{S}_n$. First, for Π^F it is clear from the transition

mechanism that, for $\sigma \in \mathbb{S}_k$ and $\pi \in \mathbb{S}_n$, the conditional probability $P(\Pi_k^F = \sigma | \Pi_n^F = \pi)$ has value 1 if the (order) restriction of π to $[k]$ is equal to σ , and that it is 0 otherwise. As functions on \mathbb{S} these are the same as the functions that generate the projective topology, see Sect. 3. Of course, convergence in the projective topology also follows directly from the fact that Π^F is consistent in the sense that, for all $n \in \mathbb{N}$, the restriction of Π_{n+1}^F to $[n]$ is equal to Π_n^F . At the other end of the spectrum, with the chain consisting of independent random variables, we obtain the one-point compactification, where each sequence of permutations that does not have a limit point in \mathbb{S} itself converges to the single point ∞ . An intermediate case is Π^H or, more generally, the copula-based process introduced in part (c) of Theorem 1. Recall that Π^H appears with the independence copula $C(x, y) = x \cdot y$.

Theorem 4 *The process $\Pi = (\Pi_n)_{n \in \mathbb{N}}$ generated by a copula C is a Markov chain with transition probabilities*

$$P(\Pi_{n+1} = \tau | \Pi_n = \sigma) = \frac{(n + 1) t(\tau, C)}{t(\sigma, C)} p^H(\sigma, \tau), \quad n \in \mathbb{N}, \sigma \in \mathbb{S}_n, \tau \in \mathbb{S}_{n+1}, \tag{18}$$

and the associated Doob-Martin topology is the same as the pattern frequency topology. Further, Π is an h -transform of Π^H .

Proof We consider Π^H first. By (15) the functions in (17) evaluate to $f_\sigma(\pi) = t(\sigma, \pi)$, which implies that the Doob-Martin topology coincides with the pattern frequency topology.

For an arbitrary copula C the permutations are generated by a sequence of independent random vectors (X_i, Y_i) , $i \in \mathbb{N}$, with distribution function C , and Π_n is a deterministic function of the first n pairs. The conditional distribution of Π_n given all Π_k , $k > n$, thus depends only on Π_{n+1} and the pair (i, j) in the permutation plot of Π_{n+1} that has been generated by (X_{n+1}, Y_{n+1}) . As the initial segments are invariant in distribution under permutations, each $i \in [n + 1]$ is equally likely. Taken together this shows that $(\Pi_n)_{n \in \mathbb{N}}$ is a Markov chain with cotransitions that do not depend on C . This completes the proof of the topological assertion.

In order to obtain the expression for the one-step transition probabilities we note that, by definition of $t(\cdot, C)$,

$$P(\Pi_k = \sigma) = t(\sigma, C) \quad \text{for all } k \in \mathbb{N}, \sigma \in \mathbb{S}_k.$$

Using this together with (14), (15) and some elementary manipulations leads to (18).

For the proof of the last assertion we need the extended Martin kernel $K(\cdot, \cdot)$ associated with Π^H . First, by the familiar arguments,

$$K(\sigma, \pi) = \frac{P(\Pi_k^H = \sigma | \Pi_n^H = \pi)}{P(\Pi_k^H = \sigma)} = k! t(\sigma, \pi)$$

for all $\sigma \in \mathbb{S}_k, \pi \in \mathbb{S}_n$ with $k \geq n$. Suppose now that $(\pi_n)_{n \in \mathbb{S}}$ is such that $\pi_n \rightarrow C \in \partial \mathbb{S}$. As the extension of $K(\sigma, \cdot)$ to $\tilde{\mathbb{S}}$ is continuous, it follows that

$$K(\sigma, C) = \lim_{n \rightarrow \infty} K(\sigma, \pi_n) = k! \lim_{n \rightarrow \infty} t(\sigma, \pi_n) = k! t(\sigma, C),$$

where we have used Theorem 1 (a) in the last step. Using this (18) can be written as

$$P(\Pi_{n+1} = \tau | \Pi_n = \sigma) = \frac{K(\tau, C)}{K(\sigma, C)} P(\Pi_{n+1}^H = \tau | \Pi_n^H = \sigma)$$

for all $n \in \mathbb{N}, \sigma \in \mathbb{S}_n$ and $\tau \in \mathbb{S}_{n+1}$. □

We presented a purely topological approach to state space augmentation that is based on an embedding of \mathbb{S} into $[0, 1]^{\mathcal{F}}$ with a suitable family \mathcal{F} of functions $f : \mathbb{S} \rightarrow [0, 1]$ and the subsequent use of Tychonov’s theorem. This is a variant of the famous Stone-Čech compactification; see e.g. (Kelley 1955, p.152). Alternatively, we could endow \mathbb{S} with a metric d and then pass to the completion of (\mathbb{S}, d) . A suitable choice of d , based on the Martin kernel, leads to a totally bounded metric space that is topologically homeomorphic to the compactification discussed above.

5 Applications

We give a range of examples where permutations and their limits are important, from classical statistics to stochastic modeling to theoretical computer science.

5.1 Tests for independence

If the X - and Y -variables are independent then we obtain the copula $C(x, y) = x \cdot y, 0 \leq x, y \leq 1$. For the asymptotic frequencies this means that, for all $k \in \mathbb{N}$ and $\sigma \in \mathbb{S}_k$,

$$t(\sigma, C) = k! P(X_1 < X_2 < \dots < X_k, Y_1 < Y_2 < \dots < Y_k),$$

as a permutation of the Y -variables alone does not change the joint distribution of the pairs $(X_1, Y_1), \dots, (X_n, Y_n)$. In particular, all patterns of the same length have the same asymptotic frequencies, so that $t(\sigma, C) = 1/k!$ for all $\sigma \in \mathbb{S}_k$.

This may be used for testing independence. A famous classical procedure is based on Kendall’s τ -statistic which, in the notation introduced in Sect. 3, may be written as $\tau = (t(12, \Pi_n) - t(21, \Pi_n)) / \binom{n}{2}$ and thus makes use of the two complementary patterns of length two. A central limit theorem for U -statistics leads to a test for independence with asymptotically correct level, see e.g. Example 12.5 in van der Vaart (1998). In our running data example a total of 61 of the 120 pairs are concurrent: Clearly, the test based on Kendall’s τ would not reject the hypothesis of independence for these data.

An analogous procedure based on longer patterns requires a multivariate central limit theorem for vectors of U -statistics. We consider patterns of length three, with a view towards the values given in Table 1. For the asymptotic covariances we use Theorem 12.3 in van der Vaart (1998) and thus have to compute $\zeta(\sigma, \tau) := 3^2 \text{cov}(h_\sigma(Z_1, Z_2, Z_3), h_\tau(Z_1, Z'_2, Z'_3))$ for $\sigma, \tau \in \mathbb{S}_3$, where $Z_1, Z_2, Z_3, Z'_2, Z'_3$ are independent and uniformly distributed on the unit square. For this we first determine the conditional distribution $\mathcal{L}[\Pi_3 | X_1 = x, Y_1 = y]$ under the assumption of independence of horizontal and vertical positions. If $\pi = 123$, for example, we either have $X_2 < x, Y_2 < y, X_3 > x, Y_3 > y$ or $X_3 < x, Y_3 < y, X_2 > x, Y_2 > y$ if (x, y) is ‘in the middle’ of the rank plot, and similar conditions apply in the other two cases. With $\phi_\sigma(x, y) := P(\Pi_3 = \sigma | X_1 = x, Y_1 = y)$ elementary calculations along these lines lead to

$$\begin{aligned} \phi_{123}(x, y) &= 2x(1-x)y(1-y) + x^2y^2/2 + (1-x)^2(1-y)^2/2, \\ \phi_{132}(x, y) &= x^2y(1-y) + x(1-x)y^2 + (1-x)^2(1-y)^2/2, \\ \phi_{213}(x, y) &= x(1-x)(1-y)^2 + (1-x)^2y(1-y) + x^2y^2/2, \\ \phi_{231}(x, y) &= x(1-x)y^2 + (1-x)^2y(1-y) + x^2(1-y)^2/2, \\ \phi_{312}(x, y) &= x(1-x)(1-y)^2 + x^2y(1-y) + (1-x)^2y^2/2, \\ \phi_{321}(x, y) &= 2x(1-x)y(1-y) + x^2(1-y)^2/2 + (1-x)^2y^2/2. \end{aligned}$$

Note that $\phi_\sigma(x, y) = \phi_{\sigma^{-1}}(y, x)$ for all $\sigma \in \mathbb{S}, 0 < x, y < 1$. With these conditional probabilities we get

$$\begin{aligned} &\text{cov}(h_\sigma(Z_1, Z_2, Z_3), h_\tau(Z_1, Z'_2, Z'_3)) \\ &= E(h_\sigma(Z_1, Z_2, Z_3)h_\tau(Z_1, Z'_2, Z'_3)) - (Eh_\sigma(Z_1, Z_2, Z_3))(Eh_\tau(Z_1, Z_2, Z_3)) \\ &= \int_0^1 \int_0^1 \phi_\sigma(x, y)\phi_\tau(x, y) dx dy \\ &\quad - \left(\int_0^1 \int_0^1 \phi_\sigma(x, y) dx dy \right) \left(\int_0^1 \int_0^1 \phi_\tau(x, y) dx dy \right). \end{aligned}$$

The observed value 0.228 for $\sigma = 312$ deviates notably from the theoretical $1/6$. For this permutation the asymptotic variance is $\zeta = 9(19/600 - 1/36) = 7/200$. The corresponding asymptotic p -value for the observation would be the probability that a standard normal variable exceeds the value

$$\sqrt{16} \sqrt{\frac{200}{7}} \left(0.228 - \frac{1}{6} \right) \approx 1.311,$$

which is about 0.094 and thus significant at the level 0.1. This, of course, is an act of ‘data snooping’, but it should be clear from the above how to obtain the 6×6 asymptotic covariance matrix for the vector $(t(\sigma, \Pi_n))_{\sigma \in \mathbb{S}_3}$ of pattern frequencies

under the hypothesis of independence, and how then to apply an appropriate test with the full data set in Table 1.

The use of ever longer patterns is a straightforward if computationally demanding generalization. In this context another interesting connection between discrete mathematics and statistics appears: A (deterministic) sequence $(\pi_n)_{n \in \mathbb{N}}$ of permutations with $|\pi_n| \rightarrow \infty$ is said to be *quasirandom* if $t(\sigma, \pi_n)$ converges to $1/|\sigma|!$ for all $\sigma \in \mathbb{S}$. The (random) sequence $(\Pi_n)_{n \in \mathbb{N}}$ generated by the independence copula has this property with probability one, so such sequences exist (there is an obvious parallel with normal numbers). Is it enough for quasirandomness if the limit statement holds for all permutations of some fixed length k ? It follows from (9) that this ‘forcing property’ would then automatically hold for all $j \geq k$ and, clearly, $k = 2$ is too small. This question has been answered in Král and Pikhurko (2013), who showed that quasirandomness follows with $k = 4$, but not with $k = 3$. Their proof of the sufficiency part is based on the permuton approach discussed in Sect. 3. In fact, as every limit point C must satisfy $t(\sigma, C) = 1/4!$ for all $\sigma \in \mathbb{S}_4$, the statement boils down to a characterization property of the independence copula. Remarkably, in a statistical context, such a characterization had been obtained much earlier in Yanagimoto (1970), and it leads to an extension of Kendall’s τ that is consistent against all alternatives. In both areas, related questions are active fields of research. In connection with quasirandomness we refer to Crudele et al. (2023), for the statistical side see e.g. Shi et al. (2022).

5.2 Delay models

Objects arrive at a system at times U_1, U_2, \dots and leave at times $U_1 + X_1, U_2 + X_2, \dots$, where we assume that the arrivals are independent and uniformly distributed on the unit interval, that the delay times are independent with distribution function G , and that arrival and delay times are independent. Let Π_n be the random permutation that connects the orders of the first n arrivals and departures. For example, in a queuing context, arriving customers might receive consecutively numbered tickets which they return on departure, and the tickets are put on a stack.

In Baringhaus and Grübel (2022) the resulting permutons, the delay copulas, are discussed. From a probabilistic and statistical point of view the information about the delay distribution contained in the permutations is of interest, together with the question of how to estimate (aspects of) the delay distribution. Another point of interest are second order approximations or, in general terms, the transitions from a strong law of large number to a central limit theorem. For the set of patterns of a fixed size this has already been important in Sect. 5.1. In Baringhaus and Grübel (2022) this is taken further by considering $\sqrt{n}(\Pi_n(\sigma) - t(\sigma, C))_{\sigma \in \mathbb{S}}$ as a stochastic process with $\sigma \in \mathbb{S}$ as ‘time parameter’ and then establishing a functional central limit theorem.

5.3 Sparseness

The delay models considered in Sect. 5.2 can be related to the $M/G/\infty$ queue where customers arrive at the times of a Poisson process with intensity λ . There are infinitely many servers so, strictly speaking, there is no queuing. The limits refer to an increasing

arrival rate and thus generate a ‘dense’ sequence of permutations as the service time distribution is kept fixed. Here we regard the standard $M/G/1$ model with only one server and fixed arrival rate; for convenience we start with an empty queue at time 0. From the first $n \in \mathbb{N}$ arrivals and departures we still obtain a sequence $(\Pi_n)_{n \in \mathbb{N}}$ of random permutations, with $\Pi_n \in \mathbb{S}_n$ for all $n \in \mathbb{N}$. We argue that this sequence is ‘sparse’, and that the pattern frequency topology does not provide much insight.

We assume that the service time distribution has finite first moment μ and that $\rho < 1$, where $\rho := \lambda\mu$ is known as the *traffic intensity*. The queue is then stable and the queue length process $Q = (Q_t)_{t \geq 0}$, with Q_t the number of customers in the system at time t , can be decomposed into alternating busy and idle periods. These are independent, and the individual sequences of busy and idle pieces of Q are identically distributed. Obviously, arrival and departure times of any particular customer belong to the same busy period. Let K_i be the number of customers in the i th period. We recall the definition of the direct sum of finite permutations: For $\sigma \in \mathbb{S}_k, \tau \in \mathbb{S}_l$, we obtain the sum $\sigma \oplus \tau \in \mathbb{S}_{l+k}$ in one-line notation as

$$\sigma \oplus \tau = (\sigma_1, \sigma_2, \dots, \sigma_k, \tau_1 + k, \tau_2 + k, \dots, \tau_l + k).$$

Suppose now that $K_1 = k$. Then, for all $n \geq k, \Pi_n(i) \leq k$ if $i \leq k$, and $\Pi_n(i) > k$ if $k < i \leq n$, hence $\Pi_n = \sigma \oplus \tau$ with some $\sigma \in \mathbb{S}_k$. Iterating this argument we see that, as $n \rightarrow \infty, \Pi_n$ decomposes into subpermutations related to the sequence of busy cycles. Also, for any given $i, \Pi_n(i)$ remains constant as soon as n exceeds the number of the first customer no longer in the same busy period as customer i . Taken together this leads to the following result.

Proposition 1 *Let $(\Pi_n)_{n \in \mathbb{N}}$ be the sequence of permutations generated by an $M/G/1$ queue with traffic intensity $\rho < 1$. Then $\Pi_n \rightarrow \Pi_\infty$ almost surely in the projective topology, with $\Pi_\infty =_{\text{distr}} \bigoplus_{i=1}^\infty \Psi_i$, where $\Psi_i \in \mathbb{S}, i \in \mathbb{N}$, are independent and identically distributed and $|\Psi_i| = K_i$.*

The distribution $\mathcal{L}(\Pi_\infty)$ of the limit Π_∞ is fully characterized by $\mathcal{L}(\Psi_1)$, which may in turn be specified by $\mathcal{L}(K_1)$ and the conditional distribution $\mathcal{L}(\Psi_1|K_1)$. Nothing is known (to me) about the latter, apart from the obvious fact that it is concentrated on the subset of \mathbb{S}_{K_1} of permutations that cannot be written as a direct sum.

The sequence $(\Pi_n)_{n \in \mathbb{N}}$ in Proposition 1 is contained in the compactification of \mathbb{S} that we used in the ‘dense’ situation. Under a mild moment condition we obtain the corresponding set of limit points.

Proposition 2 *Suppose that, in addition to the assumptions in Proposition 1, the service time distribution has finite third moment. Then, in the pattern frequency topology, $\Pi_n \rightarrow C_\infty$ almost surely, where the permuton C_∞ is given by $C_\infty(u, v) = \min\{u, v\}$ for all $u, v \in [0, 1]$.*

Proof The path-wise argument used for Proposition 1 shows that for any pair (i, j) with $i < j$ and $\Pi_n(i) > \Pi_n(j)$ customers i and j must belong to the same busy period. In particular,

$$\#\{1 \leq i < j \leq n : \Pi_n(i) > \Pi_n(j)\} \leq n M_n, \quad \text{with } M_n := \max\{K_1, \dots, K_n\},$$

so that the relative frequency of the inversion $\tau = 21$ satisfies $t(\tau, \Pi_n) \leq 2M_n/(n-1)$. It follows from Section 5.6, p158 in Cox and Smith (1961) that $E K_1^3 < \infty$ under the above moment assumption. By a standard argument this implies $M_n/n \rightarrow 0$ almost surely.

Now let σ be a permutation that contains an inversion. The relative frequency $t(\sigma, \pi)$ of σ in $\pi \in \mathbb{S}_n$, $n \geq k := |\sigma|$, can be interpreted as the probability that $\pi[A] = \sigma$, if $A \subset [n]$ with $\#A = k$ is chosen uniformly at random; see also the remarks at the beginning of Sect. 3. Further, choosing $\{i, j\} \subset [k]$ leads to an inversion with probability $p > 0$, where p does not depend on n . In the two-stage experiment, with independent steps, we would thus find an inversion in Π_n with probability $p \cdot t(\sigma, \Pi_n)$. This is bounded from above by $t(\sigma, \Pi_n)$, and the first step of the proof therefore implies that $t(\sigma, C) = \lim_{n \rightarrow \infty} t(\sigma, \Pi_n) = 0$ for all permutations that contain an inversion. As the subsampling must lead to *some* element of \mathbb{S}_k we see that any limit permutation C must assign the value 1 to each identity permutation. This holds for the copula in the assertion of the theorem, and an appeal to uniqueness of the limit completes the proof. \square

Thus, from the pattern frequency point of view, there is asymptotically no difference between $M/G/1$ and the queue where customers depart in order of their arrival, such as in a system with immediate service and constant service times.

5.4 Pattern avoiding and stack sorting

In Bassino et al. (2018) the authors deal with *separable permutations*, which are those that avoid the two patterns 2413 and 3142. This class is of interest as it appears in connection with stack sorting, for example. (The permutation in our data example is not separable; see the red dots in Fig. 1).

Starting with the sets $AV_n := \{\pi \in \mathbb{S}_n : t(2413, \pi) = 0, t(3142, \pi) = 0\}$ the main question is the limit distribution (in the weak topology related to the pattern frequency topology) of $\Pi_n \sim \text{unif}(AV_n)$ as $n \rightarrow \infty$. Whereas we often obtain a limit that is concentrated at one point, such as the independence copula in connection with the unrestricted case where $\Pi_n = \text{unif}(\mathbb{S}_n)$ in Sect. 5.1, or in the delay context in Sect. 5.2, it turns out that the limit distribution for uniformly distributed separable permutations is ‘truly random’; in fact, this distribution is an interesting object on its own. The analysis in Bassino et al. (2018) is based on the fact that separable permutations may be coded by a specific set of trees, where the operation \oplus mentioned in Sect. 5.3 together with its ‘skew’ variant \ominus plays a key role. I do not know at present if this fits into the Markovian framework. Is there a Markov chain adapted to \mathbb{S} with marginal distributions uniform on AV_n and such that the Doob-Martin approach leads to the pattern frequency topology?

Again, this is an active area of research. We refer the reader to Janson (2020) for a recent contribution, which is of particular interest in the context of the present paper because of its use of a variety of probabilistic techniques.

6 Complements

We sketch some related developments, restricting references essentially to those that can serve as entry points for a more detailed study.

6.1 Other combinatorial families

As pointed out in Hoppen et al. (2013) the construction of permutons as limits of permutations via pattern frequencies has been influenced by the earlier construction of graphons as limits of graphs via subgraph frequencies; see Lovasz (2012) for a definite account and, in addition, Diaconis and Janson (1988) for a probabilistic view. It was observed in Grübel (2015) that, with a suitable Markov chain, the topology based on subgraph frequencies is the same as the Doob-Martin topology arising from the associated boundary, in analogy to Theorem 4 above. A notable difference between the permutation and the graph situation is that graphs are considered modulo isomorphisms. In Grübel (2015) a randomization step was used to take care of this, Hagemann (2016) showed that one can work directly with the equivalence classes.

Markov chain boundary theory has been used in connection with limits of random discrete structures in a variety of combinatorial families other than \mathbb{S} ; see Grübel (2013) for an elementary introduction and references. We single out the case of binary trees, where \mathbb{B}_n is the set of binary trees with n (internal) nodes and $\mathbb{B} = \bigsqcup_{n=1}^{\infty} \mathbb{B}_n$ is the graded state space. Processes $X = (X_n)_{n \in \mathbb{N}}$ with values in \mathbb{B} arise naturally in two different situations. First, if we apply the BST (binary search tree) algorithm sequentially to a sequence of independent and identically distributed real random variables with continuous distribution function, and secondly, in the context of Rémy's algorithm, where X_n is uniformly distributed on \mathbb{B}_n . Both processes are Markov chains adapted to \mathbb{B} , but the associated transition graphs and, consequently, the cotransitions and the Doob-Martin compactifications of \mathbb{B} , are quite different. For search trees the role of pattern or subgraph frequencies is taken over by relative subtree sizes, and the permuton or graphon analogues are the probability distributions on the set of infinite 0-1 sequences; see Evans al. (2012). In the Rémy case a class of real trees appears, together with a specific sampling mechanism; see Evans et al. (2017).

Another combinatorial family that is related in several ways to the topics discussed in the present paper is the family $\mathbb{Y} = \bigsqcup_{n \in \mathbb{N}} \mathbb{Y}_n$ of (number) partitions. Here $\lambda = (\lambda_1, \dots, \lambda_k) \in \mathbb{Y}_n$ is a *partition* of n if $\lambda_i \in \mathbb{N}$, $\lambda_1 \geq \dots \geq \lambda_k$, and $\sum_{i=1}^k \lambda_i = n$. We abbreviate this to $\lambda \vdash n$. Each $\sigma \in \mathbb{S}_n$ defines a partition $\lambda \vdash n$ via the length of its cycles, in decreasing order. The right hand side of (1), for example, leads to $\lambda = (5, 3, 1)$. On \mathbb{Y} a partial order can be defined as follows: If $\lambda = (\lambda_1, \dots, \lambda_k) \in \mathbb{Y}_m$ and $(\eta_1, \dots, \eta_l) \in \mathbb{Y}_n$ with $m \leq n$ then $\lambda \leq \eta$ means that $k \leq l$ and $\lambda_j \leq \eta_j$ for $1 \leq j \leq k$. Figure 3 shows the first five levels of the associated Hasse diagram. We obtain weights for the edges of the graph by 'atom removal': For example, to go from $(3, 2)$ to the two predecessors $(3, 1)$ and $(2, 2)$ respectively there are two possibilities to decrease 2 to 1 in the first and three to decrease 3 to 2 in the second case. Normalizing these values so that they sum to 1 we obtain the cotransitions for a family of Markov chains adapted to \mathbb{Y} .




Fig. 3 The transition graph for integer partitions

Suppose now that $(\lambda(n))_{n \in \mathbb{N}}$ is a sequence in \mathbb{Y} where we again assume for simplicity that $\lambda(n) \in \mathbb{Y}_n$ for all $n \in \mathbb{N}$. It turns out that the Doob-Martin topology associated with these cotransitions is equivalent to the convergence of the normalized partition parts $\lambda(n)_i/n$ as $n \rightarrow \infty$, for all $i \in \mathbb{N}$, and that $\partial\mathbb{Y} =: \partial\mathbb{Y}_c$ is homeomorphic to the space of weakly decreasing sequences $(\alpha_i)_{i \in \mathbb{N}}$ with $\sum_{i=1}^\infty \alpha_i \leq 1$, endowed with the trace of the product topology on $[0, 1]^\infty$.

6.2 The cycle view

Above we worked with the order isomorphism aspect of permutations and the associated partial order on \mathbb{S} given by pattern containment. Using cycles instead leads to another partial order on \mathbb{S} where, for $\sigma \in \mathbb{S}_m$ and $\tau \in \mathbb{S}_n$ with $m \leq n$, the relation $\sigma \leq \tau$ means that in the cycle notation, σ arises from τ by deleting all numbers greater than m (some rearrangement may be necessary to obtain the standard notation). As an example we consider the two sides of (1) again: The order restriction to the set [4] in the one-line notation on the left leads to $3241 \in \mathbb{S}_4$, which is $(134)(2)$ in standard cycle notation. The cycle restriction to [4] of the permutation on the right hand side is $(1)(432)$ in standard notation, which is 1423 in one-line notation. In particular, Foata’s correspondence is not consistent with the two notions of restriction.

Nevertheless, the correspondence, together with the construction of Π^F on the left hand side, can be used to motivate a Markov chain $\Pi^c = (\Pi_n^c)_{n \in \mathbb{N}}$ that is adapted to \mathbb{S} and compatible with this order, and has $\Pi_n^c \sim \text{unif}(\mathbb{S}_n)$ for all $n \in \mathbb{N}$. In this construction, known as the *Chinese restaurant process*, cycles are interpreted as circular tables and customer $n + 1$ selects one of the n already seated customers as right neighbor (successor in the cycle notation), or starts a new table, where the $n + 1$ possibilities are chosen with the same probability. Note that Π^c is projective in the sense that previous values are functions of the present state.

The cycle view is closely related to number partitions: Mapping a permutation to its ordered cycle lengths leads to a function $\Psi : \mathbb{S} \rightarrow \mathbb{Y}$ that preserves the grading and the respective partial orders. Let $\Lambda^c = (\Lambda_n^c)_{n \in \mathbb{N}}$ be defined by $\Lambda_n^c := \Psi(\Pi_n^c)$ for all $n \in \mathbb{N}$. Then Λ^c is again a Markov chain, and it has the cotransitions introduced in Sect. 6.1 for the family \mathbb{Y} . Hence, almost surely as $n \rightarrow \infty$, Λ_n^c converges to an element $(\alpha_i)_{i \in \mathbb{N}}$ of $\partial\mathbb{Y}_c$ in the sense that $\lim_{n \rightarrow \infty} (\Lambda_n^c)_i/n = \alpha_i$ for all $i \in \mathbb{N}$. The distribution

of the limit, which is not concentrated at one point of the boundary, can be obtained by *stick breaking*: Starting with independent and $\text{unif}(0, 1)$ -distributed random variables $U_i, i \in \mathbb{N}$, let $(V_i)_{i \in \mathbb{N}}$ be defined by $V_1 = U_1$ and $V_{n+1} = (1 - V_n)U_{n+1}$ for all $n \in \mathbb{N}$. Putting these in decreasing order provides a random element of $\partial \mathbb{Y}_c$ that has the same distribution as Λ_∞^c .

Note that this approach refers to the big cycles. Locally, at the small end, the limit can be obtained from (3): For $i \in \mathbb{N}$ fixed, the counts of cycles of length $k, 1 \leq k \leq i$, are asymptotically independent and Poisson distributed with parameters $1/k$.

The cycle notation displays a permutation as a product of cyclic permutations that act on disjoint parts of the base set. In contrast to the general case such products are commutative. The cycle view is closely connected to the group aspects of permutations, in particular to the representation theory of finite groups. Many probabilists and statisticians became aware of this connection and the resulting ‘non-commutative Fourier analysis’ through Diaconis (1988); more recent and extensive presentations are Ceccherini-Silberstein et al. (2008) and Méliot (2017). We briefly summarize some aspects that are relevant to the asymptotics of large permutations. All representations below refer to base field \mathbb{C} , $\text{GL}(V)$ denotes the general linear (or automorphism) group of the complex vector space V , and $U(n)$ is the group of unitary $n \times n$ -matrices.

As with compactifications, we begin with an embedding: We regard the elements g of a finite group G as elements of the vector space \mathbb{C}^G of functions $f : G \rightarrow \mathbb{C}$ via $\sigma \mapsto \delta_\sigma$, where $\delta_\sigma(\sigma) = 1$ and $\delta_\sigma(\tau) = 0$ if $\tau \neq \sigma$. With the composition of $\sigma, \tau \in G$ written as $\sigma\tau$ we then define convolution on \mathbb{C}^G by

$$f \star g(\sigma) = \sum_{\tau \in G} f(\sigma\tau^{-1}) g(\tau), \quad \sigma \in G.$$

Seen as multiplication, convolution makes \mathbb{C}^G an algebra, the *group algebra* $\mathbb{C}[G]$ associated with G . Also, $(\sigma, f) \mapsto \sigma.f := \delta_\sigma \star f$ defines an action of G on V , and the function $\rho : G \rightarrow \text{GL}(V), \rho(\sigma)(f) := \sigma.f$, is a group homomorphism. With respect to the canonical basis $\{\delta_\sigma : \sigma \in G\}$ the values $\rho(\sigma)$ are all represented by unitary matrices. Thus, somewhat reminiscent of the fact that every finite group G with n elements is a subgroup of \mathbb{S}_n , every such G is a subgroup of $U(n)$, both up to group isomorphism. This is known as the (left) *regular representation* of G . More generally, a representation (ρ, W) of G consists of a vector space W and a homomorphism ρ from G to $\text{GL}(W)$. For example, representing permutations by their permutation matrices as in (4), we obtain the *permutation representation* of $G = \mathbb{S}_n$.

In a representation (ρ, W) of G a subspace U of W is invariant if $\rho(\sigma)(U) = U$ for all $\sigma \in G$, and the representation is *irreducible* if only $U = \{0\}$ and $U = W$ are invariant. Necessarily, the range of ρ is then the full group $\text{GL}(U)$. It is a crucial fact that the regular representation $(\rho, \mathbb{C}[G])$ can be decomposed into the direct sum $(\bigoplus_{i \in I} \rho_i, \bigoplus_{i \in I} W_i)$ of irreducible representations (ρ_i, W_i) . Let $d_i := \dim W_i$. For later use we note that, with n the size of G , counting dimensions leads to $\sum_{i \in I} d_i^2 = n$. The *character* $\chi : G \rightarrow \mathbb{C}$ of a representation ρ is given by $\chi(\sigma) := \text{Tr}(\rho(\sigma))$, where Tr denotes the trace (which does not depend on the basis chosen to represent $\rho(\sigma)$ by a matrix). For the permutation representation $\chi(\sigma)$ is easily seen to be the number of fixed points of σ .

We recall that general groups can be decomposed into conjugacy classes, which are the equivalence classes obtained when σ, τ are equivalent if $\sigma\pi = \pi\tau$ for some $\pi \in G$. A function $f \in \mathbb{C}[G]$ is a *class function* or *central* if it is constant on conjugacy classes. The characters of representations are such class functions, and it turns out that the characters of the irreducible representations constitute a basis for the space (in fact, a convolution algebra) of class functions. Moreover, different irreducible representations lead to different characters. For $G = \mathbb{S}_n$ the conjugacy classes correspond to cycle structures and may thus be represented by partitions of n , i.e. elements of \mathbb{Y}_n . Remarkably, the characters of irreducible representations can also be parametrized by such partitions.

This finishes our excursion, and we return to permutation asymptotics. The algebraic point of view mainly refers to characters and thus to partitions, where it leads to a close relative of the Markov chain $A^c = (A_n^c)_{n \in \mathbb{N}}$ discussed above. For this, our main source is the book Kerov (2003); see also the references given there.

It is convenient to augment the permutations of $[n]$ by fixed points so that a bijection of \mathbb{N} is obtained that leaves all $k > n$ invariant. We continue to write \mathbb{S}_n for these ‘padded’ versions, and regard them as an increasing sequence of subgroups of the group \mathbb{S}_∞ that consists of all finite bijections $\sigma : \mathbb{N} \rightarrow \mathbb{N}$, meaning that $\#\{k \in \mathbb{N} : \sigma(k) \neq k\} < \infty$. A representation ρ_{n+1} of \mathbb{S}_{n+1} then leads to a representation ρ_n of \mathbb{S}_n by restriction. This will in general destroy irreducibility, but the character of ρ_n is a class function and may thus be written as a linear combination of the irreducible characters on level n . It turns out that the non-zero coefficients are all equal to one. Consider now the graph on \mathbb{Y} with edges connecting the character of an irreducible representations on level $n + 1$ to the contributing characters of the irreducible representations on level n . This is the same graph as the transition graph obtained above for number partitions, see Fig. 3, but the edge weights are now all equal to 1. The resulting cotransitions then amount to the rule that, from a current state $\eta \in \mathbb{Y}_n$, a state $\lambda \in \mathbb{Y}_{n-1}$ is chosen uniformly at random from its predecessors. The dimension numbers $d_\lambda, \lambda \in \mathbb{Y}$, turn out to be equal to the number of paths in the diagram that lead from its root to λ . As a consequence the cotransition associated with an edge $\{\lambda, \eta\}, \lambda < \eta$, of the transition graph is given by d_λ/d_η .

When does a sequence of characters $(\chi_n)_{n \in \mathbb{N}}$ of irreducible representations or, equivalently, a sequence of partitions $(\lambda(n))_{n \in \mathbb{N}}$, converge in the Doob-Martin topology associated with these cotransitions? Or, in other words, what is the cycle view analog of the pattern frequency topology? The answer requires one more definition: The *conjugate* (or transpose) of a partition $\lambda = (\lambda_1, \dots, \lambda_k)$ is given by $\lambda^* = (\lambda_1^*, \dots, \lambda_j^*)$ with $j = \lambda_1$ and $\lambda_j^* = \#\{m \in [k] : \lambda_m \geq j\}$. Convergence then means that the normalized partition components of $\lambda(n) \in \mathbb{Y}_n$ and $\lambda^*(n)$ converge as elements of the unit interval: $\lambda(n)_i/n \rightarrow \alpha_i, \lambda^*(n)_i/n \rightarrow \beta_i$ as $n \rightarrow \infty$, for all $i \in \mathbb{N}$. The boundary is known as the *Thoma simplex*,

$$\partial\mathbb{Y}_T = \left\{ (\alpha, \beta) \in [0, 1]^{\mathbb{N}} \times [0, 1]^{\mathbb{N}} : \alpha_i \geq \alpha_{i+1}, \beta_i \geq \beta_{i+1} \text{ for all } i \in \mathbb{N}, \sum_{i=1}^{\infty} (\alpha_i + \beta_i) \leq 1 \right\},$$

together with the topology of pointwise convergence. The boundary $\partial\mathbb{Y}_c$ for the partition chain obtained above in connection with the Chinese restaurant process can be

identified with the compact subset of $\partial\mathbb{Y}_T$ that arises if all β -values are equal to zero.

A specific combinatorial Markov chain adapted to \mathbb{Y} is given by the *Plancherel growth process* $\Lambda^P = (\Lambda_n^P)_{n \in \mathbb{N}}$, with transition probabilities

$$P(\Lambda_{n+1}^P = \eta | \Lambda_n^P = \lambda) = \frac{d_\eta}{(n+1)d_\lambda}, \quad \text{for all } \lambda \in \mathbb{Y}_n, \eta \in \mathbb{Y}_{n+1} \text{ with } \lambda < \eta,$$

and marginal distributions $P(\Lambda_n^P = \lambda) = d_\lambda^2/n!$, $n \in \mathbb{N}$, $\lambda \in \mathbb{Y}_n$. It follows from the above dimension counting that the latter are indeed probability mass functions on \mathbb{Y}_n , and the calculation between Lemma 1 and Theorem 2 confirms that Λ^P has the cotransitions associated with character restrictions. In contrast to the partition chain $(\Lambda_n^C)_{n \in \mathbb{N}}$ arising in the cycle length context, the limit distribution for the Plancherel chain $(\Lambda_n^P)_{n \in \mathbb{N}}$ is concentrated at the one point of the boundary $\partial\mathbb{Y}_T$ given by $\alpha_i = \beta_i = 0$ for all $i \in \mathbb{N}$. The analysis of this chain leads to some quite spectacular results, such as the arcsine law for the limit shape of Young diagrams, and the solution of Ulam’s problem on the longest increasing subsequences in $\Pi_n \sim \text{unif}(\mathbb{S}_n)$. We refer to Kerov (2003) again, for Ulam’s problem see also Romik (2015).

In connection with trees we had two different graph structures on \mathbb{B} . Here the graphs on \mathbb{Y} are the same, but the weights are different. A generalization that includes both the cycle model and the Plancherel process is given in Kerov et al. (1998). There, an important aspect is the appearance of classical families of symmetric polynomials in the context of the respective extended Martin kernel.

6.3 Exchangeability

In Sect. 4 we worked with a general Markov chain $X = (X_n)_{n \in \mathbb{N}}$ that is adapted to a graded state space $S = \bigsqcup_{n \in \mathbb{N}} S_n$ and we used the cotransitions of X to obtain a compactification $\bar{S} = S \sqcup \partial S$ of the state space, together with a limit variable X_∞ for the X_n ’s as $n \rightarrow \infty$. Each boundary value $\alpha \in \partial S$ induces a probability measure P^α , which is the model for X conditioned on $X_\infty = \alpha$. Further, an arbitrary distribution μ on the Borel subsets of ∂S may be used to construct a Markov chain that has the same cotransitions as X and limit distribution $\mathcal{L}(X_\infty) = \int P^\alpha \mu(d\alpha)$, the μ -mixture of the P^α ’s. We note that this is the time reversal version of the familiar ‘forward’ situation where the distribution of a Markov chain is specified by the distribution of the first variable and the (forward) transition probabilities. Here, however, the distribution of X_∞ is usually defined on a non-discrete space, and there is of course no step from ∞ to ‘ $\infty - 1$ ’.

The above shows some similarities to the archetypical exchangeability result, de Finetti’s theorem: We start with a sequence $X = (X_n)_{n \in \mathbb{N}}$ of real random variables that is exchangeable in the sense that the distribution $\mathcal{L}(X)$ of X is invariant under all $\sigma \in \mathbb{S}_\infty$, the set of finite permutations of \mathbb{N} introduced in Sect. 6.2. We then obtain a limit M_∞ for the empirical distributions $M_n := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and it holds that $\mathcal{L}(X | M_\infty = \mu) = \mathcal{L}(\tilde{X})$ where $\tilde{X} = (\tilde{X}_i)_{i \in \mathbb{N}}$ with $\tilde{X}_i, i \in \mathbb{N}$, independent random variables with distribution μ . This area has developed into a major branch of modern probability, dealing with structural results for distribution families with

certain invariance properties. Kallenberg (2005) is a standard reference, but see also Aldous (1985) and Austin (2008). An excellent introduction is also given in the first part of the lecture notes Austin (2013).

Exchangeability may provide an approach to the asymptotics of random discrete structures; see e.g. Section 11.3.3 in Lovasz (2012) for graphs and Evans et al. (2017) for Rémy trees. (In the other direction, a proof for the classical de Finetti theorem can be given using Markov chain boundary theory; see Gerstenberg et al. (2016)). In their treatment of random permutations and permutons Hoppen et al. (2013) emphasize the connection to random graphs and graphons. Our aim here is to extend this to the exchangeability aspect.

In the subgraph frequency topology for sequences of (random) graphs the limits are described by a (possibly random) *graphon*, a measurable and symmetric function $W : [0, 1] \times [0, 1] \rightarrow [0, 1]$. Given W , we use a sequence $U = (U_i)_{i \in \mathbb{N}}$ of independent uniforms to construct a random binary $\mathbb{N} \times \mathbb{N}$ -matrix Z by choosing $Z_{ij} = Z_{ji} = 1$ with probability $W(U_i, U_j)$, $1 \leq i < j < \infty$, independently for different pairs, and $Z_{ii} = 0$ for all $i \in \mathbb{N}$. This matrix is jointly exchangeable in the sense that, for all $\sigma \in \mathbb{S}_\infty$, $Z^\sigma := (Z_{\sigma(i), \sigma(j)})_{i, j \in \mathbb{N}}$ has the same distribution as Z . The upper $n \times n$ -corner of Z is the incidence matrix of a random graph G_n . As $n \rightarrow \infty$, these converge in the subgraph frequency topology, and the limit is represented by W . The graphon may be regarded as an analogue of the measure M_∞ in the classical case of exchangeable sequences, and the second step as an analogue of sampling from M_∞ .

Is there a similar representation for random permutations, specifically for the copula-based models in part (c) of Theorem 1? We need an infinite matrix that represents the complete sequence, with the upper left $n \times n$ -corner for the result of the first n steps. Starting with independent random vectors (X_i, Y_i) , $i \in \mathbb{N}$, with distribution function C , we define an infinite random binary array $Z = (Z_{ij})_{i, j \in \mathbb{N}}$ by

$$Z_{ij} = \begin{cases} 1, & \text{if } i < j \text{ and } Y_i < Y_j, \\ 1, & \text{if } i > j \text{ and } X_i < X_j, \\ 0, & \text{else.} \end{cases}$$

For example, if the cities in our data set arrive in order of decreasing population size (Berlin, Hamburg, Munich, Cologne . . .) then the upper left 4×4 -corner of Z is

$$(Z_{ij})_{i, j=1}^4 = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Within this group, there is one city with more citizens than Munich that is further to the east, and none of them is further to the south. The permutation Π_n can be obtained from $(Z_{ij})_{1 \leq i, j \leq n}$ as follows: With

$$\#\{1 \leq i \leq n : X_j \leq X_i\} = \sum_{i=1}^{j-1} Z_{ij} + (n - j - 1) - \sum_{i=j+1}^n Z_{ij}, \quad 1 \leq j \leq n,$$

we get the absolute ranks of the variables $X_i, 1 \leq i \leq n$, within $\{X_1, \dots, X_n\}$, and a similar formula holds for the Y -components. Then Π_n can be determined from the two rank vectors as in (6). Equivalently, we may work with an array \tilde{Z} indexed by subsets $\{i, j\}$ of \mathbb{N} of size two, $i < j$, by using the values 0, 1, 2, 3 for indicating the relative position of (X_j, Y_j) with respect to (X_i, Y_i) through the four quadrants North-East (NE), SE, SW, and NW.

If $\sigma : \mathbb{N} \rightarrow \mathbb{N}$ is strictly increasing, then the matrix Z^σ with entries $Z_{\sigma(i),\sigma(j)}$ depends on $(X_{\sigma(i)}, Y_{\sigma(i)})$, $i \in \mathbb{N}$, in the same (deterministic) way as Z depends on (X_i, Y_i) , $i \in \mathbb{N}$. As the subsequence (X^σ, Y^σ) is equal in distribution to the original sequence (X, Y) , the arrays Z^σ and Z are then equal in distribution, which shows that Z is contractable. An analogous statement holds for the corresponding array \tilde{Z} indexed by size two subsets of \mathbb{N} . It follows from the general Aldous-Hoover-Kallenberg representation theorem that a contractable array can be extended to an exchangeable array; see Corollary 7.16 in Kallenberg (2005).

In order to construct a suitable graphon analogue W we note that the copula C represents a distribution on the unit square and thus may be written as a composition of the distribution function of the first variable and the conditional distribution function $G(x, \cdot)$ of the second variable, given the value x for the first. Let

$$W(x, y) := \inf\{z \in [0, 1] : G(x, z) \geq y\}, \quad 0 \leq y \leq 1,$$

be the corresponding quantile function. Further, let $(U_i)_{i \in \mathbb{N}}$ and $(V_i)_{i \in \mathbb{N}}$ be two independent sequences of independent $\text{unif}(0, 1)$ -distributed random variables. By construction, $(X_i, Y_i) =_{\text{distr}} (U_i, W(U_i, V_i))$ for all $i \in \mathbb{N}$, and by independence, this even holds for the two sequences. Thus the random binary matrix $\tilde{Z} = (\tilde{Z}_{ij})_{i,j \in \mathbb{N}}$ given by

$$\tilde{Z}_{ij} = \begin{cases} 1, & \text{if } i < j \text{ and } W(U_i, V_i) < W(U_j, V_j), \\ 1, & \text{if } i > j \text{ and } U_i < U_j, \\ 0, & \text{else,} \end{cases}$$

has the same distribution as Z . Of course, the conditional construction could have been done in the other direction. It is indeed quite common in these constructions that uniqueness of the representation only holds up to some equivalence, which may be difficult to describe.

It is of interest to know whether an exchangeable array is ergodic. For exchangeable sequences $(X_i)_{i \in \mathbb{N}}$ this means that $M_\infty \equiv \mu$ for some fixed distribution μ , so that the X_i 's are independent. In the permuton situation such extreme points correspond to a fixed copula, and this is always the case for the models in Theorem 1 (c). One of the interesting aspects of the results in Bassino et al. (2018) is the fact that, for separable permutations, the limit object is not degenerate.

Above we have used random matrices in order to point out the similarity to random graphs and graphons. A more abstract and more general approach has been developed in Gerstenberg (2018), where cartesian products of total orders on \mathbb{N} are considered and ergodicity is discussed in some detail.

6.4 Some connections to statistical concepts

Markov chain boundaries and exchangeable sequences both lead to parametrized families of probability distributions, where the parameter is the value of X_∞ in the first case and the value of M_∞ in the second. A relationship between the Martin boundary of Markov chains and parametric families has been pointed out early by Abrahamse (1970), together with a connection to exponential families which, roughly, appear as h -transforms in certain situations. Section 18 in Aldous (1985) deals with sufficiency and mixtures in the context of exchangeable random structures. In the framework of combinatorial Markov chains $(X_n)_{n \in \mathbb{N}}$ discussed above, where we regard the value of X_∞ in the Doob-Martin compactification as a parameter θ and the values of the first n variables X_1, \dots, X_n as data, it is clear that X_n is a sufficient statistic for θ as the conditional law of the data given X_n can be reconstructed from the cotransitions, and these do not depend on θ . Lauritzen (1974) used this as a basis for his concept of *total sufficiency*; the paper also discusses a variety of related aspects. Lauritzen (1988) emphasizes the role of the Martin boundary as a basic concept.

Further, both the set of distributions on the boundary and the set of directing measures are convex, and ‘pure parameters’ or ‘extreme models’ correspond to the ergodic case, which are extremal elements of the respective convex set. This aspect has appeared repeatedly in the above examples. Dynkin (1978) introduced *H-sufficiency*, dealing with the question of whether the convex set is a simplex (in a barycentric sense). The latter points to a connection with geometry; see Baringhaus and Grübel (2021) for a recent contribution.

Two final comments may be in order, both are somewhat loose. First, as seen above, large random discrete structures can be investigated through the Doob-Martin boundary of Markov chains, or through the representation of exchangeable distributions. The first is based on the topological approach of constructing a limit for a sequence of growing objects, the second is based on the construction of an ‘asymptotic template’ from which the sequence can be obtained by sampling, and is of a more measure-theoretical nature. (Poisson boundaries, which we did not discuss, may be seen as a measure-theoretical variant of the former.) A similar distinction appears in connection with convex sets, with Choquet’s theorem for topological vector spaces, and the barycenter approach motivated by potential theory. As a second comment, we note that permutations appear naturally in classical nonparametric statistics, where ranks are a basic tool, but that other combinatorial families may similarly be analyzed with a view towards statistical applications. Often the results on the discrete structures themselves can serve as a basis for obtaining asymptotics for a variety of specific aspects, in some resemblance to using functional limit theorems to obtain distributional limit theorems for a variety of specific statistics.

Acknowledgements Over the years, discussions with Ludwig Baringhaus, Steve Evans, Julian Gerstenberg, Klaas Hagemann, Anton Wakolbinger and Wolfgang Woess were instrumental for my understanding of many of the above concepts. I am also grateful to the referees for their helpful and encouraging comments.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The author declares that he has no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abrahamse AF (1970) A comparison between the Martin boundary theory and the theory of likelihood ratios. *Ann Math Stat* 41:1064–1067
- Aldous DJ (1985) Exchangeability and related topics, vol 1117. Springer, Berlin, pp 1–198
- Arratia R, Barbour AD, Tavaré S (2003) Logarithmic combinatorial structures: a probabilistic approach. European Mathematical Society, Zürich
- Austin T (2008) On exchangeable random variables and the statistics of large graphs and hypergraphs. *Probab Surv* 5:80–145
- Austin T (2013) Exchangeable random arrays. Lecture Notes, www.math.ucla.edu/~tim/notes.html
- Baringhaus L, Grübel R (2021) Discrete mixture representations of parametric distribution families: geometry and statistics. *Electron J Stat* 15:37–70
- Baringhaus L, Grübel R (2023) Random permutations generated by delay models and estimation of delay distributions. Submitted for publication
- Bassino F, Bouvel M, Féray V, Gerin L, Pierrot A (2018) The Brownian limit of separable permutations. *Ann Probab* 46:2134–2189
- Bona M (2004) Combinatorics of permutations. Chapman and Hall, Boca Raton
- Ceccherini-Silberstein T, Scarabotti F, Tolli F (2008) Harmonic analysis on finite groups. Cambridge University Press, Cambridge
- Cox DR, Smith WL (1961) Queues. Chapman and Hall, London
- Crudele G, Dukes P, Noel JA (2023) Six permutation patterns force quasirandomness. Available at [arXiv:2303.04776](https://arxiv.org/abs/2303.04776)
- Diaconis P (1988) Group representations in probability and statistics. Institute of Mathematical Statistics Lecture Notes—Monograph Series 11. Hayward, CA
- Diaconis P, Janson S (2008) Graph limits and exchangeable random graphs. *Rend Mat Appl* 28:33–61
- Doob JL (1959) Discrete potential theory and boundaries. *J Math Mech* 8:433–458
- Dynkin EB (1978) Sufficient statistics and extreme points. *Ann Probab* 6:705–730
- Evans SN, Grübel R, Wakolbinger A (2012) Trickle-down processes and their boundaries. *Electron J Probab* 17:58
- Evans SN, Grübel R, Wakolbinger A (2017) Doob-Martin boundary of Rémy's tree growth chain. *Ann Probab* 45:225–277
- Feller W (1968) An introduction to probability theory and its applications, vol I, 3rd edn. Wiley, New York
- Gerstenberg J (2018) Austauschbarkeit in Diskreten Strukturen: Simplexes und Filtrationen. Dissertation. Leibniz Universität Hannover
- Gerstenberg J, Grübel R, Hagemann K (2016) A boundary theory approach to de Finetti's theorem. Available at [arXiv:1610.02561](https://arxiv.org/abs/1610.02561)
- Grübel R (2013) Kombinatorische Markov-Ketten. *Math Semesterber* 60:185–215
- Grübel R (2015) Persisting randomness in randomly growing discrete structures: graphs and search trees. *Discret Math Theor Comput Sci* 18:23
- Hagemann, K. (2016) Doob-Martin-Theorie diskreter Markov-Ketten: Struktur und Anwendungen. PhD thesis, Leibniz Universität Hannover

- Hoppen C, Kohayakawa Y, Moreira CG, Ráth B, Sampaio RM (2013) Limits of permutation sequences. *J Comb Theory Ser B* 103:93–113
- Janson S (2020) Patterns in random permutations avoiding some sets of multiple patterns. *Algorithmica* 82:616–641
- Kallenberg O (2005) Probabilistic symmetries and invariance principles. Springer, New York
- Kelley JL (1955) General topology. Van Nostrand Company, Toronto-New York-London
- Kerov SV, Okounkov A, Olshanski G (1998) The boundary of the Young graph with Jack edge multiplicities. *Int Math Res Not* (4):173–199
- Kerov SV (2003) Asymptotic representation theory of the symmetric group and its applications in analysis. *Translations of Mathematical Monographs* 219. American Mathematical Society, Providence, RI
- Král D, Pikhurko O (2013) Quasirandom permutations are characterized by 4-point densities. *Geom Funct Anal* 23:570–579
- Lauritzen SL (1974) Sufficiency, prediction and extreme models. *Scand J Stat* 1:128–134
- Lauritzen SL (1988) Extremal families and systems of sufficient statistics. *Lecture Notes in Statistics* 49. Springer, Berlin
- Lovász L (2012) Large networks and graph limits. *American Mathematical Society Colloquium Publications* 60. American Mathematical Society, Providence, RI
- Méliot P-L (2017) Representation theory of symmetric groups. CRC Press, Boca Raton
- Nelson RB (2006) An introduction to copulas, 2nd edn. Springer, New York
- Romik D (2015) The surprising mathematics of longest increasing subsequences. *Institute of Mathematical Statistics Textbooks* 4. Cambridge University Press, New York
- Shi H, Drton M, Han F (2022) On the power of Chatterjee's rank correlation. *Biometrika* 109:317–333
- van der Vaart AW (1998) Asymptotic statistics. Cambridge University Press, Cambridge
- Woess W (2009) Denumerable Markov chains. European Mathematical Society, Zürich
- Yanagimoto T (1970) On measures of association and a related problem. *Ann Inst Stat Math* 22:57–63

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