Elect. Comm. in Probab. **13** (2008), 461–474

ELECTRONIC COMMUNICATIONS in PROBABILITY

FRAGMENTING RANDOM PERMUTATIONS

CHRISTINA GOLDSCHMIDT[1]
Department of Statistics, 1 South Parks Road, Oxford OX1 3TG, UK
email: goldschm@stats.ox.ac.uk

JAMES B. MARTIN
Department of Statistics, 1 South Parks Road, Oxford OX1 3TG, UK
email: martin@stats.ox.ac.uk

DARIO SPANÒ[2]
Department of Statistics, University of Warwick, Coventry CV4 7AL, UK
email: D.Spano@warwick.ac.uk

Submitted December 14, 2007, *accepted in final form* July 28, 2008

AMS 2000 Subject classification: 60C05, 05A18
Keywords: Fragmentation process, random permutation, Gibbs partition, Chinese restaurant process

Abstract
Problem 1.5.7 from Pitman's Saint-Flour lecture notes [11]: Does there exist for each n a fragmentation process $(\Pi_{n,k}, 1 \leq k \leq n)$ such that $\Pi_{n,k}$ is distributed like the partition generated by cycles of a uniform random permutation of $\{1, 2, \ldots, n\}$ conditioned to have k cycles? We show that the answer is yes. We also give a partial extension to general exchangeable Gibbs partitions.

1 Introduction

Let $[n] = \{1, 2, \ldots, n\}$, let $\mathcal{P}^n$ be the set of partitions of $[n]$, and let $\mathcal{P}^n_k$ be the set of partitions of $[n]$ into precisely k blocks.
The main result of this paper concerns the partition of $[n]$ induced by the cycles of a uniform random permutation of $[n]$. We begin by putting this in the context of more general *Gibbs partitions*. Suppose that $v_{n,k}$, $n \geq 1$, $1 \leq k \leq n$ is a triangular array of non-negative reals and w_j, $j \geq 1$ is a sequence of non-negative reals. For given n, a random partition Π is said to have the Gibbs$(\mathbf{v}, \mathbf{w})$ distribution on $\mathcal{P}^n$ if, for any $1 \leq k \leq n$ and any partition $\{A_1, A_2, \ldots, A_k\} \in \mathcal{P}^n_k$, we have

$$\mathbb{P}(\Pi = \{A_1, A_2, \ldots, A_k\}) = v_{n,k} \prod_{j=1}^{k} w_{|A_j|}.$$

[1]RESEARCH FUNDED BY EPSRC POSTDOCTORAL FELLOWSHIP EP/D065755/1.
[2]RESEARCH FUNDED IN PART BY EPSRC GRANT GR/T21783/01.

In order for this to be a well-defined distribution, the weights should satisfy the normalisation condition

$$\sum_{k=1}^{n} v_{n,k} Z_{n,k}(\mathbf{w}) = 1,$$

where

$$Z_{n,k}(\mathbf{w}) := \sum_{\{A_1,A_2,\ldots,A_k\}\in\mathcal{P}_k^n} \prod_{j=1}^{k} w_{|A_j|}.$$

$Z_{n,k}(\mathbf{w})$ is a partial Bell polynomial in the variables $w_1, w_2, \ldots$. Let K_n be the number of blocks of Π. Then it is straightforward to see that

$$\mathbb{P}\left(K_n = k\right) = v_{n,k} Z_{n,k}(\mathbf{w})$$

and so the distribution of Π conditioned on the event $\{K_n = k\}$ does not depend on the weights $v_{n,k}$, $n \geq 1$, $1 \leq k \leq n$:

$$\mathbb{P}\left(\Pi = \{A_1, A_2, \ldots, A_k\} | K_n = k\right) = \frac{\prod_{j=1}^{k} w_{|A_j|}}{Z_{n,k}(\mathbf{w})}.$$

By a Gibbs($\mathbf{w}$) partition on $\mathcal{P}_k^n$ we will mean a Gibbs($\mathbf{v}, \mathbf{w}$) partition of $[n]$ conditioned to have k blocks.

For given n, a Gibbs($\mathbf{w}$) fragmentation process is then a process $(\Pi_{n,k}, 1 \leq k \leq n)$ such that $\Pi_{n,k} \in \mathcal{P}_k^n$ for all k, which satisfies the following properties:

(i) for $k = 1, 2, \ldots, n$, we have that $\Pi_{n,k}$ is a Gibbs($\mathbf{w}$) partition;

(ii) for $k = 1, 2, \ldots, n-1$, the partition $\Pi_{n,k+1}$ is obtained from the partition $\Pi_{n,k}$ by splitting one of the blocks into two parts.

If $w_j = (j-1)!$, the Gibbs($\mathbf{w}$) distribution on $\mathcal{P}_k^n$ is the distribution of a partition into blocks given by the cycles of a uniform random permutation of $[n]$, conditioned to have k cycles. Problem 1.5.7 from Pitman's Saint-Flour lecture notes [11] asks whether a Gibbs fragmentation exists for these weights; see also [1, 6] for further discussion and for results concerning several closely related questions.

In Section 2 we will show that such a process does indeed exist. Using the Chinese restaurant process construction of a random permutation, we reduce the problem to one concerning sequences of independent Bernoulli random variables, conditioned on their sum. In Section 3 we describe a more explicit recursive construction of such a fragmentation process. Finally in Section 4 we consider the properties of Gibbs($\mathbf{w}$) partitions with more general weight sequences, corresponding to a class of exchangeable partitions of $\mathbb{N}$ (and including the two-parameter family of (α, θ)-partitions). For these weight sequences we prove that, for fixed n, one can couple partitions of $[n]$ conditioned to have k blocks, for $1 \leq k \leq n$, in such a way that the set of elements which are the smallest in their block is increasing in k. This extends the result from Section 2 on Bernoulli random variables conditioned on their sum; it is a necessary but not sufficient condition for the existence of a fragmentation process.

We note that Granovsky and Erlihson [7] have recently proved a result in the opposite direction, namely that among a more restricted class of fragmentation processes (those with the so-called "mean-field" property) no Gibbs fragmentation process exists for the weights $w_j = (j-1)!$.

2 Existence of a fragmentation process

Consider the *Chinese restaurant process* construction of a uniform random permutation of $[n]$, due to Dubins and Pitman (see, for example, Pitman [11]). Customers arrive in the restaurant one by one. The first customer sits at the first table. Customer i chooses one of the $i-1$ places to the left of an existing customer, or chooses to sit at a new table, each with probability $1/i$. So we can represent our uniform random permutation as follows. Let $C_2, \dots, C_n$ be a sequence of independent random variables such that C_i is uniform on the set $1, 2, \dots, i-1$. Let $B_1, B_2, \dots, B_n$ be independent Bernoulli random variables, independent of the sequence $C_2, C_3, \dots, C_n$ and with B_i having mean $1/i$. If $B_i = 1$ then customer i starts a new table. Otherwise, customer i places himself on the immediate left of customer C_i. Note that it is possible that a later customer will place himself on the immediate left of C_i, in which case customer i will not end up on the immediate left of C_i. The state of the system after n customers have arrived describes a uniform random permutation of $[n]$; each table in the restaurant corresponds to a cycle of the permutation, and the order of customers around the table gives the order in the cycle. Write $\Pi(B_1, B_2, \dots, B_n, C_2, C_3, \dots, C_n)$ for the random partition generated in this way (the blocks of the partition correspond to cycles in the permutation, i.e. to tables in the restaurant).

This construction has two particular features that will be important. Firstly, the number of blocks in the partition is simply $\sum_{i=1}^n B_i$. So if we condition on $\sum_{i=1}^n B_i = k$, we obtain precisely the desired distribution on $\mathcal{P}_k^n$, of the partition obtained from a uniform random permutation of $[n]$ conditioned to have k cycles. Secondly, if one changes one of the B_i from 0 to 1 (hence increasing the sum by 1), this results in one of the blocks of the partition splitting into two parts.

Hence, we can use this representation to construct our sequence of partitions $(\Pi_{n,k}, 1 \le k \le n)$. We will need the following result.

Proposition 2.1. *Let n be fixed. Then there exists a coupling of random variables B_i^k, $1 \le i \le n$, $1 \le k \le n$ with the following properties:*

(i) for each k,

$$(B_1^k, B_2^k, \dots, B_n^k) \stackrel{d}{=} \left(B_1, B_2, \dots, B_n \,\middle|\, \sum_{i=1}^n B_i = k \right);$$

(ii) for all k and i, if $B_i^k = 1$ then $B_i^{k+1} = 1$, with probability 1.

So we fix $C_2, C_3, \dots, C_n$, and define

$$\Pi_{n,k} = \Pi(B_1^k, B_2^k, \dots, B_n^k, C_2, C_3, \dots, C_n).$$

Then $(\Pi_{n,k}, 1 \le k \le n)$ is the desired Gibbs fragmentation process.

It remains to prove Proposition 2.1. Write $B^k = (B_1^k, B_2^k, \dots, B_n^k)$ for the random variable in part (i), let $b = (b_1, b_2, \dots, b_n)$ and

$$p^k(b) := \mathbb{P}\left(B^k = b\right),$$

where $b_i \in \{0, 1\}$ for $1 \le i \le n$. Clearly, $p^k(b)$ is only non-zero for sequences b having exactly k 1's. Write $\mathcal{S}_k$ for the subset of $\{0,1\}^n$ consisting of these sequences and write $b \prec b'$ whenever b' can be obtained from b by replacing one of the co-ordinates of b which is 0 by a 1. We need

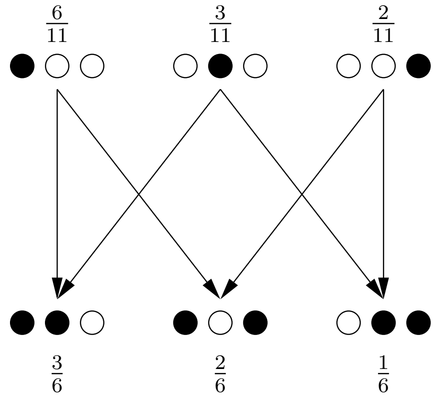

Figure 1: The case $n = 4$ and $k = 2, 3$. Since B_1 is always 1, we omit it from the picture. For $2 \le i \le 4$, a filled circle in position i indicates that $B_i^k = 1$; an empty circle indicates that it is 0. The arrows indicate which states with $k = 3$ can be obtained from those with $k = 2$; the numbers are the probabilities $p^k(b)$ of the states.

a process $(B^1, B^2, \dots, B^n)$ whose kth marginal B^k has distribution p^k and such that, with probability 1, $B^k \prec B^{k+1}$ for each k.
It is enough to show that for each k, we can couple B^k with distribution p^k and B^{k+1} with distribution p^{k+1} in such a way that $B^k \prec B^{k+1}$. (The couplings can then be combined, for example in a Markovian way, to give the desired law on the whole sequence). As noted in Section 4 of Berestycki and Pitman [1], a necessary and sufficient condition for such couplings to exist with specified marginals and order properties was given by Strassen; see for example Theorem 1 and Proposition 4 of [10]. (Strassen's theorem may be seen as a version of Hall's marriage theorem [8] and is closely related to the max-flow/min-cut theorem [3, 4]). The required condition may be stated as follows. For $C \subseteq \mathcal{S}_k$, write $N(C) = \{b' \in \mathcal{S}_{k+1} : b \prec b' \text{ some } b \in C\}$. We then need that for all $C \subseteq \mathcal{S}_k$,

$$\sum_{b \in C} p^k(b) \le \sum_{b' \in N(C)} p^{k+1}(b') \tag{1}$$

It is convenient to phrase things a little differently. Suppose $C \subseteq \mathcal{S}_k$ has m elements, say $b^{(1)}, b^{(2)}, \dots, b^{(m)}$. Let $A_j = \{i \in [n] : b_i^{(j)} = 1\}$, $1 \le j \le m$. Then, by construction, each A_j has k elements. Now set

$$E_j = \{B_i = 1 \ \forall \ i \in A_j\}, \ 1 \le j \le m.$$

Then

$$\sum_{b \in C} p^k(b) = \mathbb{P}\left(E_1 \cup E_2 \cup \dots \cup E_m \,\middle|\, \sum_{i=1}^n B_i = k \right)$$

and

$$\sum_{b' \in N(C)} p^{k+1}(b') = \mathbb{P}\left(E_1 \cup E_2 \cup \dots \cup E_m \,\middle|\, \sum_{i=1}^n B_i = k+1 \right).$$

So (1) is equivalent to the following proposition.

Proposition 2.2. *Let $A_1, A_2, \ldots, A_m$ be any collection of distinct k-subsets of $[n]$ and $E_1, E_2 \ldots, E_m$ the corresponding events, as defined above. Then*

$$\mathbb{P}\left(E_1 \cup E_2 \cup \cdots \cup E_m \middle| \sum_{i=1}^n B_i = k\right) \leq \mathbb{P}\left(E_1 \cup E_2 \cup \cdots \cup E_m \middle| \sum_{i=1}^n B_i = k+1\right),$$

$1 \leq k \leq n-1.$

This is a corollary of the following result from Efron [2].

Proposition 2.3. *Let $\phi_n : \{0,1\}^n \to \mathbb{R}_+$ be a function which is increasing in all of its arguments. Let I_i, $1 \leq i \leq n$, be independent Bernoulli random variables (not necessarily with the same parameter). Then*

$$\mathbb{E}\left[\phi_n(I_1, I_2, \ldots, I_n) \middle| \sum_{i=1}^n I_i = k\right] \leq \mathbb{E}\left[\phi_n(I_1, I_2, \ldots, I_n) \middle| \sum_{i=1}^n I_i = k+1\right],$$

for all $0 \leq k \leq n-1$.

Since Efron's proof does not apply directly to the case of discrete random variables, we give a proof here.

Proof. First let $\psi : \mathbb{Z}_+^2 \to \mathbb{R}_+$ be a function which is increasing in both arguments. Fix $n \geq 1$ and write $X_n = \sum_{i=1}^n I_i$. We will first prove that for $0 \leq k \leq n$,

$$\mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = k\right] \leq \mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = k+1\right]. \tag{2}$$

Let $u(i) = \mathbb{P}(X_n = i)$ for $0 \leq i \leq n$. We observe that, as a sum of independent Bernoulli random variables, X_n is *log-concave*, that is

$$u(i)^2 \geq u(i-1)u(i+1), \quad i \geq 1. \tag{3}$$

(This follows because any Bernoulli random variable is log-concave, and the sum of independent log-concave random variables is itself log-concave, as proved by Hoggar [9]).
Since $\psi(0,0) \leq \psi(1,0)$ and $\psi(0,0) \leq \psi(0,1)$, we have

$$\begin{aligned}
\mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = 0\right] &= \psi(0,0) \\
&\leq \frac{p_{n+1}u(0)\psi(1,0) + (1-p_{n+1})u(1)\psi(0,1)}{p_{n+1}u(0) + (1-p_{n+1})u(1)} \\
&= \mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = 1\right].
\end{aligned}$$

Similarly, since $\psi(n-1,1) \leq \psi(n,1)$ and $\psi(n,0) \leq \psi(n,1)$, we have

$$\begin{aligned}
\mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = n\right] &= \frac{p_{n+1}u(n-1)\psi(n-1,1) + (1-p_{n+1})u(n)\psi(n,0)}{p_{n+1}u(n-1) + (1-p_{n+1})u(n)} \\
&\leq \psi(n,1) \\
&= \mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = n+1\right].
\end{aligned}$$

Suppose now that $1 \le k \le n-1$. Then

$$\begin{aligned}
&\mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = k+1\right] - \mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = k\right] \\
&= \frac{p_{n+1}u(k)\psi(k,1) + (1-p_{n+1})u(k+1)\psi(k+1,0)}{p_{n+1}u(k) + (1-p_{n+1})u(k+1)} \\
&\quad - \frac{p_{n+1}u(k-1)\psi(k-1,1) + (1-p_{n+1})u(k)\psi(k,0)}{p_{n+1}u(k-1) + (1-p_{n+1})u(k)} \\
&= \frac{p_{n+1}^2 u(k)u(k-1)[\psi(k,1) - \psi(k-1,1)]}{d} \\
&\quad + \frac{(1-p_{n+1})^2 u(k)u(k+1)[\psi(k+1,0) - \psi(k,0)]}{d} \\
&\quad + \frac{p_{n+1}(1-p_{n+1})[u(k-1)u(k+1)\psi(k+1,0) - u(k)^2\psi(k,0)]}{d} \\
&\quad + \frac{p_{n+1}(1-p_{n+1})[u(k)^2\psi(k,1) - u(k-1)u(k+1)\psi(k-1,1)]}{d},
\end{aligned} \tag{4}$$

where the denominator d is given by

$$d = \left[p_{n+1}u(k) + (1-p_{n+1})u(k+1)\right]\left[p_{n+1}u(k-1) + (1-p_{n+1})u(k)\right]$$

and is clearly non-negative. The first two terms in (4) are non-negative because ψ is increasing. The sum of the third and fourth terms is bounded below by

$$\frac{p_{n+1}(1-p_{n+1})}{d}[u(k)^2 - u(k-1)u(k+1)][\psi(k,1) - \psi(k,0)].$$

This is non-negative by the log-concavity property (3). Putting all of this together, we see that

$$\mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = k\right] \le \mathbb{E}\left[\psi(X_n, I_{n+1}) | X_n + I_{n+1} = k+1\right],$$

as required.

We now proceed by induction on n. Note that (2) with $n = 1$ gives the base case. Now define

$$\psi_n\left(\sum_{i=1}^{n-1} I_i, I_n\right) = \mathbb{E}\left[\phi_n(I_1, I_2, \dots, I_n) \middle| \sum_{i=1}^{n-1} I_i, I_n\right].$$

Assume that we have

$$\mathbb{E}\left[\phi_{n-1}(I_1, I_2, \dots, I_{n-1}) \middle| \sum_{i=1}^{n-1} I_i = k\right] \le \mathbb{E}\left[\phi_{n-1}(I_1, I_2, \dots, I_{n-1}) \middle| \sum_{i=1}^{n-1} I_i = k+1\right]$$

for $0 \le k \le n-2$. By this induction hypothesis, $\psi_n\left(\sum_{i=1}^{n-1} I_i, I_n\right)$ is increasing in its first argument. By the assumption that ϕ_n is increasing, $\psi_n\left(\sum_{i=1}^{n-1} I_i, I_n\right)$ is increasing in its second argument. So by (2),

$$\mathbb{E}\left[\psi_n\left(\sum_{i=1}^{n-1} I_i, I_n\right) \middle| \sum_{i=1}^{n-1} I_i + I_n = k\right] \le \mathbb{E}\left[\psi_n\left(\sum_{i=1}^{n-1} I_i, I_n\right) \middle| \sum_{i=1}^{n-1} I_i + I_n = k+1\right]$$

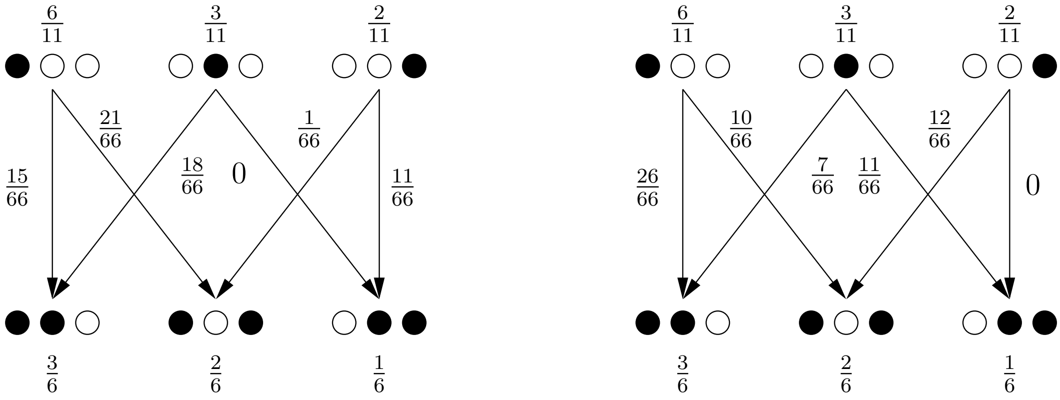

Figure 2: The two extreme solutions to the coupling problem for $n = 4$ and $k = 2, 3$ (as drawn in Figure 1). The arrows are labelled with the joint probabilities $\mathbb{P}\left(B^k = b, B^{k+1} = b'\right)$.

for $0 \le k \le n-1$. But by the tower law, this says exactly that

$$\mathbb{E}\left[\phi_n(I_1, I_2, \ldots, I_n)\middle|\sum_{i=1}^n I_i = k\right] \le \mathbb{E}\left[\phi_n(I_1, I_2, \ldots, I_n)\middle|\sum_{i=1}^n I_i = k+1\right].$$

The result follows by induction. □

Proof of Proposition 2.2. Let $\phi_n(B_1, B_2, \ldots, B_n) = \mathbb{1}_{\{E_1 \cup E_2 \cup \cdots \cup E_m\}}$. Then ϕ_n is increasing in all of its arguments and so by the previous proposition we have

$$\begin{aligned}\mathbb{P}\left(E_1 \cup E_2 \cup \cdots \cup E_m \middle| \sum_{i=1}^n B_i = k\right) &= \mathbb{E}\left[\phi_n(X_1, X_2, \ldots, X_n)\middle|\sum_{i=1}^n B_i = k\right] \\ &\le \mathbb{E}\left[\phi_n(X_1, X_2, \ldots, X_n)\middle|\sum_{i=1}^n B_i = k+1\right] \\ &= \mathbb{P}\left(E_1 \cup E_2 \cup \cdots \cup E_m \middle| \sum_{i=1}^n B_i = k+1\right). \qquad \square\end{aligned}$$

Note that we have only proved the existence of a coupling. There is no reason why it should be unique. Indeed, in general, there is a simplex of solutions. For the example given in Figure 1, the extremes of the one-parameter family of solutions are shown in Figure 2.

3 A recursive construction

In this section, we describe a more explicit construction of the Gibbs fragmentation processes, which is recursive in n and possesses a certain consistency property as n varies.

The basic principle is the following simple observation. Suppose we want to create a uniform random permutation of $[n]$ conditioned to have k cycles. Then n either forms a singleton, or is contained in some cycle with other individuals. If it forms a singleton, then the rest of the

permutation is a uniform random permutation of $[n-1]$, conditioned to have $k-1$ cycles. If, on the other hand, n is not a singleton, then we take a uniform random permutation of $[n-1]$ with k cycles and insert n into a uniformly chosen position.
So we proceed as follows. A Gibbs fragmentation process on $[n]$ for $n=1$ or $n=2$ is trivial. Suppose we have constructed a process $(\Pi_1^{n-1}, \Pi_2^{n-1}, \ldots, \Pi_{n-1}^{n-1})$ on $[n-1]$, with the required marginal distributions and splitting properties. We will derive from it a process $(\Pi_1^n, \Pi_2^n, \ldots, \Pi_n^n)$ on $[n]$.
For each k, the partition Π_k^n of $[n]$ into k parts will come from either

(a) adding a singleton block $\{n\}$ to Π_{k-1}^{n-1}; or

(b) adding the element n to one of the blocks of Π_k^{n-1}, by choosing an element C_n uniformly at random from $[n-1]$ and putting n in the same block as C_n.

If we do (a) with probability $\mathbb{P}(\{n\} \text{ is a singleton in } \Pi_k^n)$ and (b) otherwise then we will obtain Π_k^n having the desired distribution. Now note that

$$\mathbb{P}(\{n\} \text{ is a singleton in } \Pi_k^n) = \mathbb{P}\left(B_n = 1 \,\middle|\, \sum_{i=1}^n B_i = k \right) = \frac{Z_{n-1,k-1}(\mathbf{w})}{Z_{n,k}(\mathbf{w})},$$

where $w_j = (j-1)!$ and so $Z_{n,k}(\mathbf{w})$ is a Stirling number of the first kind. Note that this probability is increasing in k, by the monotonicity results proved in the previous section. So first let R_n be a random variable which is independent of $(\Pi_1^{n-1}, \Pi_2^{n-1}, \ldots, \Pi_{n-1}^{n-1})$ and whose distribution is given by

$$\mathbb{P}(R_n \le k) = \mathbb{P}(\{n\} \text{ is a singleton in } \Pi_k^n).$$

Then for $k < R_n$, add n to one of the blocks of Π_k^{n-1} by choosing C_n uniformly at random from $[n-1]$ and putting n in the same block as C_n. (Note that we use the same value of C_n for each such k). For $k \ge R_n$, create Π_k^n by taking Π_{k-1}^{n-1} and adding $\{n\}$ as a singleton block. By construction, Π_k^n has the same distribution as the partition derived from a random permutation of $[n]$ conditioned to have k blocks. Moreover, the partitions $(\Pi_k^n, 1 \le k \le n)$ are nested because, firstly, the partitions $(\Pi_k^{n-1}, 1 \le k \le n-1)$ were nested (by assumption) and, secondly, once $\{n\}$ has split off (at time R_n) it remains a singleton block.
A possible realisation of this recursive process is illustrated in Figure 3.
To explain exactly what we mean by the *consistency in* n which this construction possesses, let

$$\tilde{\Pi}_k^n = \begin{cases} \Pi_k^n & \text{if } 1 \le k < R_n \\ \Pi_{k+1}^n & \text{if } R_n \le k \le n-1. \end{cases}$$

Then if we restrict the sequence $(\tilde{\Pi}_k^n, 1 \le k \le n-1)$ of partitions of $[n]$ to $[n-1]$, we obtain exactly $(\Pi_k^{n-1}, 1 \le k \le n-1)$.

4 Partial extension to exchangeable Gibbs partitions

At the beginning of this paper, we introduced the notion of a partition of $[n]$ with Gibbs$(\mathbf{v}, \mathbf{w})$ distribution. In general, there is no reason why these partitions should be consistent as n varies. That is, it is not necessarily the case that taking a Gibbs$(\mathbf{v}, \mathbf{w})$ partition of $[n+1]$ and deleting $n+1$ gives rise to a Gibbs$(\mathbf{v}, \mathbf{w})$ partition of $[n]$. If, however, this *is* the case, we can

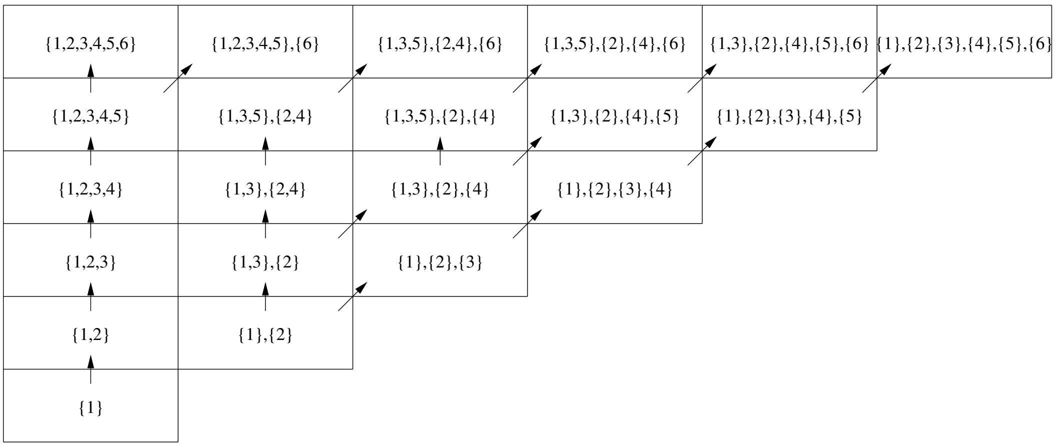

Figure 3: A possible realisation of the coupling.

define an exchangeable random partition Π of $\mathbb{N}$ as the limit of the sequence of projections onto $[n]$ as $n \to \infty$. In this case, we refer to Π as an *exchangeable Gibbs*$(\mathbf{v}, \mathbf{w})$ *partition*.
An important subfamily of the exchangeable Gibbs$(\mathbf{v}, \mathbf{w})$ partitions are the (α, θ)-partitions, whose asymptotic frequencies have the Poisson-Dirichlet(α, θ) distribution, for $0 \le \alpha < 1$ and $\theta > -\alpha$, or $\alpha < 0$ and $\theta = m|\alpha|$, some $m \in \mathbb{N}$ (see Pitman and Yor [12] or Pitman [11] for a wealth of information about these distributions). Here, the corresponding weight sequences are

$$w_j = (1-\alpha)_{j-1\uparrow 1}, \quad v_{n,k} = \frac{(\theta+\alpha)_{(k-1)\uparrow\alpha}}{(\theta+1)_{(n-1)\uparrow 1}},$$

where $(x)_{m\uparrow\beta} := \prod_{j=1}^{m}(x+(j-1)\beta)$ and $(x)_{0\uparrow\beta} := 1$. In the first part of this paper, we have treated the case $\alpha = 0$: a $(0,1)$-partition of $[n]$ has the same distribution as the partition derived from the cycles of a uniform random permutation. In view of the fact that the array $\mathbf{v}$ does not influence the partition conditioned to have k blocks, we have also proved that a fragmentation process having the distribution at time k of a $(0, \theta)$-partition conditioned to have k blocks exists. The Gibbs$(\mathbf{w})$ distribution corresponding to the case $\alpha = -1$ (i.e. to weight sequence $w_j = j!$) is discussed in Berestycki and Pitman [1]; in particular, it is known that a Gibbs$(\mathbf{w})$ fragmentation exists in this case.
Gnedin and Pitman [5] have proved that the exchangeable Gibbs$(\mathbf{v}, \mathbf{w})$ partitions all have $\mathbf{w}$-sequences of the form

$$w_j = (1-\alpha)_{j-1\uparrow 1}$$

for some $-\infty \le \alpha \le 1$, where for $\alpha = -\infty$ the weight sequence is interpreted as being identically equal to 1 for all j. Furthermore, the array $\mathbf{v}$ must solve the backward recursion

$$v_{n,k} = \gamma_{n,k} v_{n+1,k} + v_{n+1,k+1}, \quad 1 \le k \le n,$$

where $v_{1,1} = 1$ and

$$\gamma_{n,k} = \begin{cases} n - \alpha k & \text{if } -\infty < \alpha < 1 \\ k & \text{if } \alpha = -\infty. \end{cases} \tag{5}$$

The case $\alpha = 1$ corresponds to the trivial partition into singletons and will not be discussed any further.
It seems natural to consider the question of whether Gibbs fragmentations exist for other weight sequences falling into the exchangeable class. We will give here a partial extension of our results to the case of a general $\alpha \in [-\infty, 1)$.
Fix n and consider a Gibbs$(\mathbf{v}, \mathbf{w})$ partition of $[n]$ with $w_j = (1-\alpha)_{j-1\uparrow 1}$ for some $\alpha \in [-\infty, 1)$. Let B_i be the indicator function of the event that i is the smallest element in its block, for $1 \leq i \leq n$, and let $K_n = \sum_{i=1}^n B_i$, the number of blocks. Then we have the following extension of the earlier Proposition 2.1.

Proposition 4.1. *For each n, there exists a random sequence of vectors $\left(B_1^{k,n}, \ldots, B_n^{k,n}\right)$, $k = 1, 2, \ldots, n$ such that for each k,*

$$\left(B_1^{k,n}, \ldots, B_n^{k,n}\right) \stackrel{d}{=} (B_1, \ldots, B_n | K_n = k),$$

and such that $B_i^{n,k+1} \geq B_i^{n,k}$ for all k.

This is precisely what we proved earlier for the case $\alpha = 0$ (and, indeed, for more general sequences of independent Bernoulli random variables). Namely we show that Gibbs$(\mathbf{w})$ partitions of $[n]$, conditioned to have k blocks, can be coupled over k in such a way that the set of elements which are the smallest in their block is increasing in k. In the context of random permutations, this was enough to prove the existence of a coupling of *partitions* with the desired fragmentation property, using the fact that the random variables C_i in the Chinese restaurant process, which govern the table joined by each arriving customer, were independent of each other and from the random variables B_i. For $\alpha \neq 0$, a Chinese restaurant process construction is still possible (see Pitman [11]), although it gives rise only to a partition rather than a permutation. However, it is no longer true that the "increments" or "records" are independent of the rest of the partition and so we cannot deduce the full result for partitions from Proposition 4.1. In fact, the stronger result is not always true, for example when $\alpha = -\infty$ and so $w_j = 1$ for all j. In this case it is known that no Gibbs fragmentation process exists for $n = 20$ and for all large enough n (see for example [1] for a discussion). By a continuity argument, one can show similarly that if $-\alpha$ is sufficiently large, then for certain n no Gibbs fragmentation process exists.
As in Section 3, Proposition 4.1 can be proved by induction over n. Suppose we have carried out the construction for $n-1$, and wish to extend to n. Conditioned on $K_n = k$, we need to consider two cases: either $B_n = 1$ and $K_{n-1} = k-1$, or $B_n = 0$ and $K_n = k$. Depending on which of these cases we choose, we will set either

$$\left(B_1^{n,k}, \ldots, B_n^{n,k}\right) = \left(B_1^{n-1,k-1}, \ldots, B_{n-1}^{n-1,k-1}, 1\right)$$

or

$$\left(B_1^{n,k}, \ldots, B_n^{n,k}\right) = \left(B_1^{n-1,k}, \ldots, B_{n-1}^{n-1,k}, 0\right).$$

Precisely as in Section 3, this can be made to work successfully provided that the following lemma holds.

Lemma 4.2. *For any n, the probability $\mathbb{P}(B_n = 1 | K_n = k)$ is increasing in k.*

Equivalently, we are showing that the expected number of singletons in Π_n, conditioned on $K_n = k$, is increasing in k (since, by exchangeability, the probability that n is a singleton in Π_n is the same as the probability that i is a singleton, for any $i \in [n]$).
The rest of this section is devoted to the proof of this lemma.
Define $S_\alpha(n,k) = Z_{n,k}(\mathbf{w})$. Then

$$\mathbb{P}(K_n = k) = v_{n,k} S_\alpha(n,k), \tag{6}$$

where the S_α obey the recursion

$$S_\alpha(n+1,k) = \gamma_{n,k} S_\alpha(n,k) + S_\alpha(n,k-1), \tag{7}$$

with boundary conditions $S_\alpha(1,1) = 1$, $S_\alpha(n,0) = 0$ for all n, and $S_\alpha(n,n+1)$=0 for all n. These are *generalized Stirling numbers.* We have

$$\mathbb{P}(B_n = 1 | K_n = k) = \mathbb{P}(K_{n-1} = k-1 | K_n = k) = \frac{S_\alpha(n-1,k-1)}{S_\alpha(n,k)}$$

(see Section 3 of [5] for further details). Using the recursion, this is equal to

$$\left(1 + \frac{\gamma_{n-1,k} S_\alpha(n-1,k)}{S_\alpha(n-1,k-1)}\right)^{-1}.$$

For this to be increasing in k, it is equivalent that

$$\frac{S_\alpha(n-1,k-1)}{\gamma_{n-1,k} S_\alpha(n-1,k)}$$

should be increasing in k. This is implied by the following proposition.

Lemma 4.3. *For all $\alpha < 1$ and all n and k,*

$$\gamma_{n,k} S_\alpha(n,k)^2 \geq \gamma_{n,k+1} S_\alpha(n,k+1) S_\alpha(n,k-1).$$

Proof. If $0 \leq \alpha < 1$ then it is sufficient to prove the statement that $(S_\alpha(n,k), 0 \leq k \leq n)$ is log-concave, that is

$$S_\alpha(n,k)^2 \geq S_\alpha(n,k-1) S_\alpha(n,k+1),$$

since in that case $\gamma_{n,k}$, defined at (5), is decreasing in k. Theorem 1 of Sagan [13] states that whenever $t_{n,k}$ is a triangular array satisfying

$$t_{n,k} = c_{n,k} t_{n-1,k-1} + d_{n,k} t_{n-1,k}$$

for all $n \geq 1$, where $t_{n,k}$, $c_{n,k}$ and $d_{n,k}$ are all integers and such that

- $c_{n,k}$ and $d_{n,k}$ are log-concave in k,
- $c_{n,k-1} d_{n,k+1} + c_{n,k+1} d_{n,k-1} \leq 2 c_{n,k} d_{n,k}$ for all $n \geq 1$,

then $t_{n,k}$ is log-concave in k. These conditions are clearly satisfied for the generalized Stirling numbers $S_\alpha(n,k)$, with the exception that the sequence $d_{n,k}$ is not integer-valued. However, Sagan's argument extends immediately to this case also, and so we will not give a proof here.

We turn now to the case $-\infty < \alpha < 0$. We proceed by induction on n. For $n = 2$ the statement is trivial. Suppose that we have

$$(n - \alpha k)S_\alpha(n,k)^2 \geq (n - \alpha(k+1))S_\alpha(n,k-1)S_\alpha(n,k+1),$$

for $0 \leq k \leq n$. Using the recurrence

$$S_\alpha(n+1,k) = (n - \alpha k)S_\alpha(n,k) + S_\alpha(n,k-1),$$

we obtain that

$$\begin{aligned}
&(n+1-\alpha k)S_\alpha(n+1,k)^2 - (n+1-\alpha(k+1))S_\alpha(n+1,k-1)S_\alpha(n+1,k+1) \\
&= \Big\{(n+1-\alpha k)(n-\alpha k)^2 S_\alpha(n,k)^2 \\
&\qquad - (n+1-\alpha(k+1))(n-\alpha(k-1))(n-\alpha(k+1))S_\alpha(n,k-1)S_\alpha(n,k+1)\Big\} \\
&\quad + \Big\{2(n+1-\alpha k)(n-\alpha k)S_\alpha(n,k)S_\alpha(n,k-1) \\
&\qquad - (n+1-\alpha(k+1))(n-\alpha(k+1))S_\alpha(n,k+1)S_\alpha(n,k-2) \\
&\qquad - (n+1-\alpha(k+1))(n-\alpha(k-1))S_\alpha(n,k)S_\alpha(n,k-1)\Big\} \\
&+ \Big\{(n+1-\alpha k)S_\alpha(n,k-1)^2 - (n+1-\alpha(k+1))S_\alpha(n,k-2)S_\alpha(n,k)\Big\}. \qquad (8)
\end{aligned}$$

We take each of the three terms in braces separately. For the first term, we note that

$$(n+1-\alpha k)(n-\alpha k) - (n+1-\alpha(k+1))(n-\alpha(k-1)) = \alpha(\alpha-1).$$

Hence, by the induction hypothesis, the first term is greater than or equal to

$$\alpha(\alpha-1)(n-k\alpha)S_\alpha(n,k)^2.$$

Since $\alpha < 0$, this is non-negative. For the second term in (8), we note that applying the induction hypothesis twice entails that

$$(n-\alpha(k-1))S_\alpha(n,k)S_\alpha(n,k-1) \geq (n-\alpha(k+1))S_\alpha(n,k+1)S_\alpha(n,k-2).$$

So

$$\begin{aligned}
&2(n+1-\alpha(k+1))(n-\alpha(k-1))S_\alpha(n,k)S_\alpha(n,k-1) \\
&\geq (n+1-\alpha(k+1))(n-\alpha(k+1))S_\alpha(n,k+1)S_\alpha(n,k-2) \\
&\qquad + (n+1-\alpha(k+1))(n-\alpha(k-1))S_\alpha(n,k)S_\alpha(n,k-1).
\end{aligned}$$

It follows that the second term in braces is bounded below by

$$\begin{aligned}
&2S_\alpha(n,k)S_\alpha(n,k-1)\,[(n+1-\alpha k)(n-\alpha k) - (n+1-\alpha(k+1))(n-\alpha(k-1))] \\
&\quad = 2\alpha(\alpha-1)S_\alpha(n,k)S_\alpha(n,k-1).
\end{aligned}$$

This is non-negative. Finally, we turn to the third term in braces in (8). This is equal to

$$\begin{aligned}
&[(n-\alpha(k-1))S_\alpha(n,k-1)^2 - (n-\alpha k)S_\alpha(n,k-2)S_\alpha(n,k)] \\
&- (\alpha-1)[S_\alpha(n,k-1)^2 - S_\alpha(n,k-2)S_\alpha(n,k)].
\end{aligned}$$

By the induction hypothesis (at n and $k-1$), the first term is non-negative and so this quantity is bounded below by

$$(1-\alpha)[S_\alpha(n,k-1)^2 - S_\alpha(n,k-2)S_\alpha(n,k)].$$

That $S_\alpha(n,k-1)^2 - S_\alpha(n,k-2)S_\alpha(n,k) \geq 0$ is implied by the induction hypothesis (in fact it is a weaker statement) and so, since $\alpha < 0$, the above quantity is non-negative. So (8) is non-negative, as required. The case $\alpha = -\infty$ follows similarly. □

Acknowledgments. We are grateful to Jay Taylor and Amandine Véber for organising the reading group on Pitman's Saint-Flour notes [11] which led to the work in this paper. We would like to thank Sasha Gnedin and Nathanaël Berestycki for valuable discussions and an anonymous referee for helpful comments.

References

[1] N. Berestycki and J. Pitman. Gibbs distributions for random partitions generated by a fragmentation process. *J. Stat. Phys.*, 127(2):381–418, 2007. MR2314353

[2] B. Efron. Increasing properties of Pólya frequency functions. *Ann. Math. Statist.*, 36:272–279, 1965. MR0171335

[3] P. Elias, A. Feinstein, and C. Shannon. A note on the maximum flow through a network. *Institute of Radio Engineers, Transactions on Information Theory*, IT-2:117–119, 1956.

[4] L. R. Ford, Jr. and D. R. Fulkerson. Maximal flow through a network. *Canad. J. Math.*, 8:399–404, 1956. MR0079251

[5] A. Gnedin and J. Pitman. Exchangeable Gibbs partitions and Stirling triangles. *J. Math. Sci. (N. Y.)*, 138(3):5677–5685, 2006. Translated from *Zapiski Nauchnykh Seminarov POMI*, Vol. 325, 2005, pp. 83–102. MR2160320

[6] A. Gnedin and J. Pitman. Poisson representation of a Ewens fragmentation process. *Combin. Probab. Comput.*, 16(6):819–827, 2007. MR2351686

[7] B. Granovsky and M. Erlihson. On time dynamics of coagulation-fragmentation processes. `arXiv:0711.0503v2 [math.PR]`, 2007.

[8] P. Hall. On representatives of subsets. *J. London Math. Soc.*, 10:26–30, 1935.

[9] S. G. Hoggar. Chromatic polynomials and logarithmic concavity. *J. Combinatorial Theory Ser. B*, 16, 1974. MR0342424

[10] T. Kamae, U. Krengel, and G. L. O'Brien. Stochastic inequalities on partially ordered spaces. *Ann. Probab.*, 5(6):899–912, 1977. MR0494447

[11] J. Pitman. *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2006. Lectures from the 32nd Summer School on Probability Theory held in Saint-Flour, July 7–24, 2002. MR2245368

[12] J. Pitman and M. Yor. The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann. Probab.*, 25(2):855–900, 1997. MR1434129

[13] B. E. Sagan. Inductive and injective proofs of log concavity results. *Discrete Math.*, 68(2-3):281–292, 1988. MR0926131