Bayesian Analysis (2022) **TBA**, Number TBA, pp. 1

Product-form exchangeable feature probability functions and Gibbs partitions

Lorenzo Ghilotti *

Abstract. Feature allocation models may be seen as generalizations of partitions. In fact, while in a partition model each individual is assigned to one and only one class, in a feature allocation model each individual can belong to multiple groups. In general, the groups can be naturally interpreted as traits or features: an individual belonging to multiple groups corresponds to the individual exhibiting multiple traits or features. Feature allocations with product-form exchangeable partition probability functions may be considered as the counterpart of Gibbs-type partitions. Gnedin and Pitman (2004) propose a characterization of the class of Gibbs-type partitions and Battiston et al. (2018) characterize the class of feature allocations with product-form. In this review, the two characterizations are discussed. In particular, I focus on highlighting a clear common structure between the two.

Keywords: product-form feature allocations, Gibbs partitions, Indian buffet Process, Beta-Bernoulli Process.

1 Introduction to feature allocation models

Consider a set of n individuals (or points). In the setting of feature allocation models, each individual is assumed to possess a random set of traits (or features) among the collection of all possible traits. Formally, consider a set with n points and let the points be indexed by the integers $[n] := \{1, \dots, n\}$. When $n = +\infty$, I am referring to the index set $\mathbb{N} = \{1, 2, \dots\}$.

Definition 1.1. *A feature allocation f_n of $[n]$ is a multiset of non-empty subsets of $[n]$ called features, such that no index i belongs to infinitely many features. I write $f_n = \{A_1, \dots, A_k\}$, where k is the number of features.*

For defining exchangeable feature allocations, let $\mathcal{F}_n$ be the space of all feature allocations of $[n]$. A random feature allocation F_n of $[n]$ is a random element of $\mathcal{F}_n$. Let $\sigma : [n] \longrightarrow [n]$ be a finite permutation. Moreover, for any feature $A \subset [n]$, denote the permutation applied to the feature as follows: $\sigma(A) := \{\sigma(n) : n \in A\}$. For any feature allocation f_n, denote the permutation applied to the feature allocation as follows: $\sigma(f_n) := \{\sigma(A) : A \in f_n\}$. Finally, let F_n be a random feature allocation of $[n]$.

Definition 1.2. *A random feature allocation F_n is exchangeable if $F_n \stackrel{d}{=} \sigma(F_n)$ for every permutation σ of $[n]$.*

Let $(F_n)_n$ be a sequence of exchangeable random feature allocations. In addition to exchangeability (from now on exchangeability will be always assumed but no more

*University of Milano-Bicocca l.ghilotti@campus.unimib.it

specified), a natural property to require is the coherence of the distributions across different values of n. Intuitively, the distribution of F_n should coincide with the one of F_{n+1}, once the last individual is integrated out. In this case, the sequence $(F_n)_n$ is said to be *consistent*. Refer to Broderick et al. (2013) for the formal definition.

When dealing with exchangeable random feature allocations, the so-called *random ordered feature allocation* $\tilde{F}_n$ is useful. Refer to Broderick et al. (2013), Section 3. Note that $\tilde{F}_n$ is a random sequence - rather than a collection - of subsets of $[n]$. This alternative tool allows to avoid some combinatorial factors needed to work with the distribution of F_n (see Broderick et al. (2013), Example 1), while retaining exchangeability and possibly consistency. Moreover, the probability of a random feature allocation, $\mathbb{P}(F_n = f_n)$ is related to the probability of a random ordered feature allocation, $\mathbb{P}(\tilde{F}_n = \tilde{f}_n)$, via Equation (5) in Broderick et al. (2013).

Assume that the feature allocation probability (related to $[n]$) admits the representation

$$\mathbb{P}(\tilde{F}_n = \tilde{f}_n) = \pi_n(|A_1|, \dots, |A_k|) \tag{1.1}$$

for every ordered feature allocation $\tilde{f}_n = (A_1, \dots, A_k)$ and some symmetric function $\pi_n : \cup_{k=0}^{\infty} [n]^k \longrightarrow [0, 1]$, then π_n is called the *exchangeable feature probability function* (EFPF). It is worth noting that random feature allocations may not have an EFPF (Proposition 7, Broderick et al. (2013)). However, in the following I focus on random feature allocations with EFPF. It may be observed that, when dealing with random partitions, an exchangeable random partition distribution always admits an exchangeable partition probability function (EPPF) representation.

When $(F_n)_n$ is a consistent sequence of random feature allocations (in the following simply referred to as consistent feature allocation) with EFPFs, then the following consistency condition for the EFPFs holds:

$$\pi_n(m_1, \dots, m_k) = \sum_{j=0}^{\infty} \binom{k+j}{j} \sum_{z \in \{0,1\}^k} \pi_{n+1}(m_1 + z_1, \dots, m_k + z_k, 1, \dots, 1)$$

for all $n \geq 1$, where the sequence of 1s in the second term is of length j.

The most common example of consistent feature allocation with EFPF is the Indian buffet Process (IBP). Specifically, I refer to the 3-parameter (γ, α, θ) IBP, with $\gamma \geq 0, \alpha \in [0, 1], \theta \in [-\alpha, \infty)$, if the EFPF has the form

$$\frac{1}{k!} \left(\frac{\gamma}{(\theta + 1)_{n-1\uparrow}} \right)^k \exp\left(- \sum_{i=1}^{n} \gamma \frac{(\alpha + \theta)_{i-1\uparrow}}{(1 + \theta)_{i-1\uparrow}} \right) \prod_{l=1}^{k} (1 - \alpha)_{m_l - 1\uparrow} (\theta + \alpha)_{n - m_l\uparrow}$$

where $(x)_{m\uparrow} = \prod_{i=0}^{m-1} (x + i)$ and $(x)_{0\uparrow} = 1$. The 2-parameter (γ, θ) IBP is obtained when $\alpha = 0$ and the 1-parameter γ IBP is obtained when $\alpha = 0, \theta = 1$. For the purpose of the paper, I further need to introduce the Beta-Bernoulli model with parameters (N, α, θ), with $N \in \mathbb{N}, \alpha \in (-\infty, 0), \theta \in [-\alpha, \infty)$, whose EFPF writes

$$\binom{N}{k} \left(\frac{-\alpha}{(\theta + \alpha)_{n\uparrow}} \right)^k \left(\frac{(\theta + \alpha)_{n\uparrow}}{(\theta)_{n\uparrow}} \right)^N \prod_{l=1}^{k} (1 - \alpha)_{m_l - 1\uparrow} (\theta + \alpha)_{n - m_l\uparrow}$$

For $\alpha \to -\infty, -\frac{\alpha}{\theta} \to q \in (0,1)$, the homogeneous Bernoulli model of parameters (N, q) is obtained (see Appendix A.1, Battiston et al. (2018)). In the following refer to the homogeneous Bernoulli model of parameters (N, q) as the special Beta-Bernoulli model with $\alpha = -\infty$ and characterize it with $q \in (0,1)$.

The next sections are organized as follows. In Section 2, Gibbs-type random partitions are briefly recalled and the well-known characterizations of Gnedin and Pitman (2004) are reviewed. In the setting of random partitions, Gibbs-type models can be seen as the counterpart of the class of models studied by Battiston et al. (2018) within the setting of random feature allocations, the so-called feature allocations with product-form EFPFs (discussed in Section 3). In Section 3, the class of feature allocations with product-form EFPFs is introduced and the main theorem of Battiston et al. (2018) is reported (Theorem 3.1). It is a characterization for such a class of models which represents the counterpart of results in Section 2 within the feature allocations setting. As an example, the feature allocation distributions induced by the stable beta scaled processes SB-SP (Camerlenghi et al. (2021)) are characterized via Theorem 3.1. Finally, in Section 4, the analogies between the characterization for Gibbs-type partitions of Gnedin and Pitman (2004) (Theorems 2.1 and 2.2) and the characterization for feature allocations with product-form EFPFs of Battiston et al. (2018) (Theorem 3.1) are pointed out.

2 Random partitions of Gibbs-type and the main characterization

Random partitions may be regarded as particular cases of random feature allocations where each individual possesses one and only one feature. The notions of consistency and exchangeability (see Pitman (2006)) were first introduced for random partitions and then extended to feature allocations. Any consistent and exchangeable random partition distribution can be characterized by its exchangeable partition probability function (EPPF) p_n. A distinguished class of exchangeable partitions is the two-parameter (α, θ) family (also known as Ewens-Pitman family), with EPPF

$$p_n(\lambda_1, \dots, \lambda_k) = \frac{(\theta+\alpha)_{k-1\uparrow\alpha}}{(\theta+1)_{n-1\uparrow}} \prod_{l=1}^{k} (1-\alpha)_{\lambda_j - 1\uparrow}$$

where $\lambda_1, \dots, \lambda_k$ is a composition of n and $(x)_{m\uparrow\beta} = \prod_{i=0}^{m-1} (x + i\beta)$, with $(x)_{0\uparrow\beta} = 1$. Parameters (α, θ) are such that either (i) $\alpha \in [-\infty, 0)$ and $\theta = m|\alpha|, m = 1, \dots, \infty$ or (ii) $\alpha \in [0,1]$ and $\theta \geq -\alpha$, with proper definition in the limiting cases. Gnedin and Pitman (2004) introduce the class of random partition distributions called *Gibbs-type partitions*, which generalizes the class of Ewens-Pitman models.

Definition 2.1. *An exchangeable random partition is said to be of Gibbs form if, for some nonnegative weights $W = (W_j)$ and $V = (V_{n,k})$, its EPPF satisfies*

$$p_n(\lambda_1, \dots, \lambda_k) = V_{n,k} \prod_{l=1}^{k} W_{\lambda_l} \tag{2.1}$$

for all $1 \le k \le n$ *and all compositions* $(\lambda_1, \ldots, \lambda_k)$ *of* n.

Lemma 2 in Gnedin and Pitman (2004) provides a characterization of the Gibbs-type partitions in terms of the form of W and V, equivalently restated as follows.

Theorem 2.1. *The distribution of a consistent exchangeable random partition can be represented by an EPPF of form (2.1), i.e. it is of Gibbs form, with* $W_1 = V_{1,1}$ *if and only if one of the following two cases holds:*

(i) W *can be uniquely written as* $W_j = (1-\alpha)_{j-1\uparrow}$ *for fixed* $\alpha \in (-\infty, 1)$*, and the elements of* V *satisfy the recursion (with* $V_{1,1} = 1$*)*

$$V_{n,k} = (n - \alpha k)V_{n+1,k} + V_{n+1,k+1}, \qquad 1 \le k \le n$$

(ii) W *can be uniquely written as* $W_j = 1$ *and* V *satisfies the recursion (with* $V_{1,1} = 1$*)*

$$V_{n,k} = kV_{n+1,k} + V_{n+1,k+1}, \qquad 1 \le k \le n$$

corresponding to the limiting case $\alpha \to -\infty$ *(indicate this case with* $\alpha = -\infty$*).*

The theorem could also deal with the extension for $\alpha = 1$, but it corresponds to the trivial singleton partition, so it is excluded here. Similarly, the random partition having a unique cluster almost surely is not taken into consideration. Moreover, the main result of Gnedin and Pitman (2004) is the characterization of each Gibbs partition of type $\alpha \in [-\infty, 1)$ as a mixture of special partitions (depending on α). The theorem is presented as follows.

Theorem 2.2 (Theorem 12 of Gnedin and Pitman (2004))**.** *Each Gibbs partition of fixed type* $\alpha \in [-\infty, 1)$ *is a unique probability mixture of the extreme partitions of type*

(i) $(\alpha, |\alpha|m)$*-partitions with* $m = 1, \ldots, \infty$*, for* $\alpha \in [-\infty, 0)$*;*

(ii) the Ewens $(0, \theta)$*-partitions with* $\theta \in [0, \infty]$*, for* $\alpha = 0$ *;*

(iii) the Poisson-Kingman $(\alpha|s)$*-partitions with* $s \in [0, \infty]$*, for* $\alpha \in (0, 1)$*, where Poisson-Kingman* (α, s)*-partition refers to the partition derived from a Poisson-Kingman discrete distribution denoted in Pitman (2003) by* $PK(\rho_\alpha|t)$ *for* $t = s^{-\alpha}$*.*

3 Feature allocations with product-form EFPFs and the main characterization

Battiston et al. (2018) consider the class of distributions for feature allocations with EFPFs of the form

$$\pi_n(m_1, \ldots, m_k) = V_{n,k} \prod_{l=1}^{k} W_{m_l} U_{n-m_l} \tag{3.1}$$

for an infinite array $V = (V_{n,k} : (n,k) \in \mathbb{N} \times \mathbb{N}_0)$ and two sequences $W = (W_j : j \in \mathbb{N})$ and $U = (U_j : j \in \mathbb{N}_0)$ of nonnegative weights. They refer to them as feature allocations with product-form EFPFs. They provide the following characterization of such a class of distributions in terms of V, W and U.

Theorem 3.1. *The distribution of a consistent exchangeable feature allocation can be represented by an EFPF of the form (3.1) if and only if one of the following three cases holds:*

(i) *W and U can be uniquely written as $W_m = (1-\alpha)_{m-1\uparrow}$ and $U_m = (\theta+\alpha)_{m\uparrow}$, for fixed α, θ satisfying $\alpha \in (-\infty, 1)$ and $\theta \in (-\alpha, \infty)$, and the elements of V satisfy the recursion*

$$V_{n,k} = \sum_{j=0}^{\infty} \binom{k+j}{j} \left((\theta+\alpha)_{n\uparrow}\right)^j (\theta+n)^k V_{n+1,k+j} \tag{3.2}$$

(ii) *W and U can be uniquely written as $W_m = q^{m-1}$ and $U_m = (1-q)^m$, for some $q \in (0,1)$, and V satisfies the recursion*

$$V_{n,k} = \sum_{j=0}^{\infty} \binom{k+j}{j} (1-q)^{nj} V_{n+1,k+j}$$

corresponding to the limiting case $\alpha \to -\infty$ and $-\frac{\alpha}{\theta} \to q$ (indicate this case with $\alpha = -\infty$ and characterize it with $q \in (0,1)$).

(iii) *One of the following two degenerate cases holds.*

(a) *There is no feature sharing. In this case, $W_m = (1-\alpha)_{m-1\uparrow}$ for $\alpha = 1$, and $\tilde{V}_{n,k} := V_{n,k} U_{n-1}^k$ satisfies $\tilde{V}_{n,k} = \sum_{j=0}^{\infty} \binom{k+j}{j} \tilde{V}_{n+1,k+j}$*

(b) *There is complete feature sharing. In this case, $U_m = (\theta+\alpha)_{m\uparrow}$ for $\theta = -\alpha$, and $\tilde{V}_{n,k} := V_{n,k} W_n^k$ satisfies $\tilde{V}_{n,k} = \tilde{V}_{n+1,k}$*

Moreover, for fixed (α, θ) (or q if $\alpha = -\infty$), the set of solutions of these recursions is:

1. *for $\alpha \in (0,1]$, mixtures over γ of the V of a 3 -parameter IBP;*

2. *for $\alpha = 0$, mixtures over γ of the V of a 2 -parameter IBP;*

3. *for $\alpha \in [-\infty, 0)$, mixtures over N of the V of a Beta-Bernoulli model with N features.*

Note that this theorem may be seen as the counterpart, in the context of feature allocations with product-form EFPFs, of the characterization for Gibbs-type partitions described in Theorem 2.1 and Theorem 2.2.

General idea for the proof of points 1, 2 and 3 (second part) For $\alpha \in (-\infty, 1), \theta \in (-\alpha, \infty)$ (the case $\alpha = -\infty$ and $q \in (0,1)$ is similar), let $\mathcal{V}_{\alpha,\theta}$ be the set of the elements $V \in \mathbb{R}_+^{\mathbb{N}\times\mathbb{N}_0}$ satisfying (3.2). Let $\mathcal{B}_\mathcal{V}$ be the smallest σ-algebra that makes the maps $V \mapsto V_{n,k}$ measurable, and equip $\mathcal{V}_{\alpha,\theta}$ with $\mathcal{B}_\mathcal{V}$. For each measure μ on $\mathcal{B}_\mathcal{V}$, define the barycenter V^μ as

$$V^\mu_{n,k} = \int_{\mathcal{V}_{\alpha,\theta}} V_{n,k} d\mu$$

It is shown that $\mathcal{V}_{\alpha,\theta}$ is a convex set, i.e. for all probability measures μ on $\mathcal{B}_\mathcal{V}$, $V^\mu \in \mathcal{V}_{\alpha,\theta}$. Let $\mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$ be the set of all distributions of feature allocations with product-form EFPF (3.1) with W and U set as in point (i) of Theorem 3.1. Consequently, each $\mathbb{P} \in \mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$ induces a distribution for $(K_n)_n$, $n \in \mathbb{N}$, where K_n is the number of features in the corresponding random feature allocation of n individuals. Such a distribution is defined on $(\mathbb{N}_0^\infty, \mathcal{C}(\mathbb{N}_0^\infty))$, where $\mathcal{C}(\mathbb{N}_0^\infty)$ is the cylinder σ-algebra. Consider the map $T : \mathcal{V}_{\alpha,\theta} \longrightarrow \mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$, such that $T(V) = P_V \in \mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$, where P_V is the feature allocation distribution with product-form EFPF (3.1) with weights W and U set as in point (i) of Theorem 3.1, and V. Battiston et al. (2018) show that the map T is a convex isomorphism, that is an isomorphism between convex sets (note that $\mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$ is a convex set), meaning that T is invertible and T and T^{-1} are measurable and preserve the convex structure (see Proposition A.4 in Battiston et al. (2018)). This establishes a bijection between the extreme points of $\mathcal{V}_{\alpha,\theta}$ and $\mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$ (recall that an extreme point of a convex set is a point which cannot be represented as a convex combination of two other points of the set). The importance of extreme points here stems from the following remark: by general theory of Dynkin (1978) and Diaconis and Freedman (1984), each element of $\mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$ can be uniquely represented as a convex mixture of the extreme elements of $\mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$, and the same can be said for $\mathcal{V}_{\alpha,\theta}$. Therefore, the proof proceeds by determining the extreme points of $\mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$ (refer to Battiston et al. (2018)). It turns out that the set of extreme points of $\mathcal{P}_{\mathcal{V}_{\alpha,\theta}}$, and consequently of $\mathcal{V}_{\alpha,\theta}$, remarkably depends on the value of α, as described in points 1, 2 and 3 of Theorem 3.1.

In the next example the novel class of stable beta scaled processes SB-SP introduced by Camerlenghi et al. (2021) is shown to induce a family of feature allocations with product-form EFPFs of type $\alpha \in (0,1)$ (case 1, Theorem 3.1).

Example 3.1. *Consider the SB-SP of parameters (σ, c, β), with $\sigma \in (0,1), c, \beta > 0$ (refer to Camerlenghi et al. (2021)). From Proposition 4 of Camerlenghi et al. (2021), the EFPF of the induced feature allocation is*

$$\pi_n(m_1, \ldots, m_k) = \frac{1}{k!} \frac{\sigma^k \beta^{c+1} \Gamma(k+c+1)}{(\beta + \gamma_0^{(n)})^{k+c+1} \Gamma(c+1)} \prod_{l=1}^{k} \frac{\Gamma(m_l - \sigma)\Gamma(n - m_l + 1)}{\Gamma(n - \sigma + 1)}$$

It can be equivalently expresses as

$$\begin{aligned}\pi_n(m_1, \ldots, m_k) &= \frac{1}{k!} \frac{\sigma^k \beta^{c+1} \Gamma(k+c+1)\Gamma(1-\sigma)}{(\beta + \gamma_0^{(n)})^{k+c+1} \Gamma(c+1)} \prod_{l=1}^{k} \frac{(1-\sigma)_{m_l-1\uparrow}(1)_{n-m_l\uparrow}}{\Gamma(n-\sigma+1)} \\ &= \frac{1}{k!} \frac{\sigma^k \beta^{c+1} (c+1)_{k\uparrow}}{(\beta + \gamma_0^{(n)})^{k+c+1} (1-\sigma)_{n\uparrow}} \prod_{l=1}^{k} (1-\sigma)_{m_l-1\uparrow}(1)_{n-m_l\uparrow}\end{aligned}$$

It is clear that this EFPF is in product-form (3.1). Consider the two sequences of weights W and U as $W_m = (1-\sigma)_{m-1\uparrow}$ and $U_m = (1)_{m\uparrow}$. Identifying $\alpha := \sigma \in (0,1)$, this corresponds to case (i) of Theorem 3.1 with $\alpha \in (0,1)$ and $\theta = 1-\alpha \in (-\alpha,\infty)$. Moreover, since $\alpha \in (0,1)$ and point 1 of Theorem 3.1, each feature allocation distribution induced by a SB-SP (σ, c, β) is a mixture over γ of 3-parameter IBP $(\gamma, \sigma, 1-\sigma)$ distributions.

4 Similarities between the two characterizations

In this section, analogies and differences between the characterizations for Gibbs-type random partitions and feature allocations with product-form EFPFs, without focusing on degenerate cases like (iii) of Theorem 3.1, are discussed. Firstly, I focus on the comparison between the sequence of weights W for the Gibbs-type partitions and the sequences W and U for the feature allocations with product-form EFPFs. Secondly, for fixed values of the parameters defining the already discussed weights, i.e. α for the Gibbs-type partitions and (α,θ) or q for the feature allocations with product-form EFPFs, considerations about the sequence V are presented.

On the sequence of weights W (and U for feature allocations) For Gibbs-type partitions, the sequence of weights W (refer to Definition 2.1) is parametrized by the single parameter $\alpha \in [-\infty, 1)$. The case $\alpha = -\infty$ denotes the limiting case described in point (ii) of Theorem 2.1, where W is a sequence of 1s.

For feature allocations with product-form EFPFs, the two sequences of weights W and U are parametrized by two parameters (α,θ) such that $\alpha \in [-\infty,1), \theta \in (-\alpha,\infty)$. The case $\alpha = -\infty$ denotes the limiting case described in point (ii) of Theorem 3.1, where W and U are still parametrized by the single parameter $q \in (0,1)$.

On the sequence V for fixed values of the parameters determining W (and U for feature allocations) For Gibbs-type partitions, V must satisfy a recursion depending on α (Theorem 2.1) and the set of solutions $\mathcal{V}_\alpha$ forms a convex set (see Gnedin and Pitman (2004)). The extreme points of $\mathcal{V}_\alpha$ distinguishably depend on α. For $\alpha \in [-\infty, 0)$, the set of extreme points of $\mathcal{V}_\alpha$ is countably infinite and coincide with the family of V of the Ewens-Pitman models $(\alpha, m|\alpha|)$, with $m = 1,\dots,\infty$. For $\alpha = 0$, this set coincides with the family of V of the Ewens family θ, $\theta \in (0,\infty)$ (i.e. the Ewens-Pitman family $(0,\theta)$). For $\alpha \in (0,1)$, the set of extremes coincides with the family of V of the Poisson-Kingman $(\alpha|s)$-partitions, with $s \in [0,\infty]$.

For feature allocations with product-form EFPFs, V must satisfy a recursion depending on (α,θ) or q (Theorem 3.1) and the set of solutions $\mathcal{V}_{\alpha,\theta}$ (or $\mathcal{V}_q$) forms a convex set (see Battiston et al. (2018)). Similarly to the Gibbs-type partitions setting, the extreme points of such a convex set strongly depend on α (and a minor role is played by θ or q). For $\alpha \in [-\infty, 0)$, the set of extreme points is countably infinite and coincides with the family of V of the Beta-Bernoulli models (N,α,θ), with $N \in \mathbb{N}$ (when $\alpha = -\infty$, the Beta-Bernoulli is intended as the homogeneous Bernoulli model of parameters (N, q), review Section 1). For $\alpha = 0$, the set of extreme points coincides with the family of V of the 2-parameter IBP family (γ, θ), $\gamma \geq 0$ (i.e. the 3-parameter IBP $(\gamma, 0, \theta)$). For $\alpha \in (0, 1)$, the set of extremes coincides with the family of V of the 3-parameter IBP family (γ, α, θ), $\gamma \geq 0$.

References

Battiston, M., Favaro, S., Roy, D. M., and Teh, Y. W. (2018). "A characterization of product-form exchangeable feature probability functions." *The Annals of Applied Probability*, 28(3): 1423 – 1448. 1, 3, 4, 6, 7

Broderick, T., Pitman, J., and Jordan, M. I. (2013). "Feature Allocations, Probability Functions, and Paintboxes." *Bayesian Analysis*, 8(4): 801 – 836. 2

Camerlenghi, F., Favaro, S., Masoero, L., and Broderick, T. (2021). "Scaled process priors for Bayesian nonparametric estimation of the unseen genetic variation." 3, 6

Diaconis, P. and Freedman, D. (1984). "Partial exchangeability and sufficiency." *Statistics: applications and new directions*, 205–236. 6

Dynkin, E. (1978). "Sufficient statistics and extreme points." *The Annals of Probability*, 6(5): 705–730. 6

Gnedin, A. and Pitman, J. (2004). "Exchangeable Gibbs partitions and Stirling triangles." 1, 3, 4, 7

Pitman, J. (2003). "Poisson-kingman partitions." *Lecture Notes-Monograph Series*, 1–34. 4

— (2006). *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Berlin: Springer-Verlag. Lectures from the 32nd Summer School on Probability Theory held in Saint-Flour, July 7–24, 2002, With a foreword by Jean Picard. 3