

Beta-Binomial stick-breaking non-parametric prior

María F. Gil-Leyva

*IIMAS, Universidad Nacional Autónoma de México
CDMX, México
e-mail: marifer@sigma.iimas.unam.mx*

Ramsés H. Mena

*IIMAS, Universidad Nacional Autónoma de México
CDMX, México
e-mail: ramses@sigma.iimas.unam.mx*

Theodoros Nicolieris

*Department of Economics
National and Kapodistrian University of Athens
Athens, Greece
e-mail: tnicolieris@econ.uoa.gr*

Abstract: A new class of nonparametric prior distributions, termed Beta-Binomial stick-breaking process, is proposed. By allowing the underlying length random variables to be dependent through a Beta marginals Markov chain, an appealing discrete random probability measure arises. The chain's dependence parameter controls the ordering of the stick-breaking weights, and thus tunes the model's label-switching ability. Also, by tuning this parameter, the resulting class contains the Dirichlet process and the Geometric process priors as particular cases, which is of interest for MCMC implementations.

Some properties of the model are discussed and a density estimation algorithm is proposed and tested with simulated datasets.

Keywords and phrases: Beta-Binomial Markov chain, density estimation, Dirichlet process prior, geometric process prior, stick-breaking prior.

Received November 2019.

1. Introduction

Discrete random probability measures and their distributions play a key role in Bayesian nonparametric statistics. The availability of general classes of priors and their different representations are crucial for the study of theoretical properties, as well as for the proposal of simulation and estimation algorithms. This continuously encourages the search of competitive alternatives to the canonical model, Ferguson (1973) Dirichlet process. At the outset, one could consider a (proper) species sampling process (Pitman; 2006) over a measurable Polish

space $(S, \mathcal{B}(S))$,

$$\boldsymbol{\mu} = \sum_{j \geq 1} \mathbf{w}_j \delta_{\boldsymbol{\xi}_j}, \quad (1)$$

where the atoms, $\boldsymbol{\Xi} = (\boldsymbol{\xi}_j)_{j \geq 1}$, and the weights, $\mathbf{W} = (\mathbf{w}_j)_{j \geq 1}$, are independent collections of random variables (r.v.'s), with $\boldsymbol{\xi}_j \stackrel{\text{iid}}{\sim} P_0$, a diffuse measure on $(S, \mathcal{B}(S))$, and $\sum_{j \geq 1} \mathbf{w}_j = 1$, almost surely (a.s.). To fully specify the law of $\boldsymbol{\mu}$, one could assume a form for P_0 and place a distribution over the infinite dimensional simplex $\Delta_\infty = \{(w_1, w_2, \dots) : w_i \geq 0, \sum_{i \geq 1} w_i = 1\}$. An important aspect to note is that

$$\sum_{j \geq 1} \mathbf{w}_j \delta_{\boldsymbol{\xi}_j} \stackrel{d}{=} \sum_{j \geq 1} \mathbf{w}_{\rho(j)} \delta_{\boldsymbol{\xi}_j} \quad (2)$$

for every random permutation of \mathbb{N} , ρ , independent of $\boldsymbol{\Xi}$. This means that once the atom's distribution, P_0 , is fixed, there are infinitely many distributions over Δ_∞ that lead to the exact same prior, hence the need to study orderings for the weights. In particular, one can consider the decreasing ordering of its elements, here denoted by $\mathbf{W}^\downarrow = (\mathbf{w}_j^\downarrow)_{j \geq 1}$, with $\mathbf{w}_1^\downarrow > \mathbf{w}_2^\downarrow > \dots$ a.s., or the size-biased permutation, denoted by $\tilde{\mathbf{W}} = (\tilde{\mathbf{w}}_j)_{j \geq 1}$, which satisfies $\mathbb{P}[\tilde{\mathbf{w}}_1 = \mathbf{w}_j | \mathbf{W}] = \mathbf{w}_j$, and for $n \geq 2$

$$\mathbb{P}[\tilde{\mathbf{w}}_n = \mathbf{w}_j | \mathbf{W}, \tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_{n-1}] = \frac{\mathbf{w}_j}{1 - \sum_{i=1}^{n-1} \tilde{\mathbf{w}}_i} \mathbf{1}_{\{\mathbf{w}_j \notin \{\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_{n-1}\}\}}.$$

Working with decreasing representations of the weights reduces the identifiability problem that arises from (2) in the sense that if $\gamma_1, \gamma_2, \dots$ is sampled i.i.d. from $\boldsymbol{\mu}$, conditionally given $\boldsymbol{\mu}$, then \mathbf{w}_1^\downarrow corresponds to the atom that appears more frequently in the sequence, \mathbf{w}_2^\downarrow corresponds to the second most frequent value, and so on (e.g., Mena and Walker; 2015). On the other hand, the size-biased permutation of the weights is of interest when the focus is in the clusters featured in the sample, i.e. if γ_j^* is the j th distinct value to appear in the sample, then the long-run proportion of elements in $\{n : \gamma_n = \gamma_j^*\}$ coincides precisely with $\tilde{\mathbf{w}}_j$ (Pitman; 1996a).

Different techniques to place distributions on Δ_∞ are available (e.g. Ferguson; 1973; Blackwell and MacQueen; 1973; James et al.; 2009) and connections among such techniques are well known (e.g. Ishwaran and James; 2001; Ishwaran and Zarepour; 2002; Hjort et al.; 2010). Perhaps one of the most practical constructions is enjoyed by the so-called stick-breaking process (McCloskey; 1965; Sethuraman; 1994; Ishwaran and James; 2001) where the weights are decomposed as

$$\mathbf{w}_1 = \mathbf{v}_1, \quad \mathbf{w}_j = \mathbf{v}_j \prod_{i=1}^{j-1} (1 - \mathbf{v}_i), \quad j \geq 2, \quad (3)$$

for some sequence taking values in $[0, 1]$, $\mathbf{V} = (\mathbf{v}_i)_{i \geq 1}$, hereinafter referred to as length variables (l.v.'s). The practical compromise inherent to (3) is relatively little, as most practical classes of priors have a stick-breaking representation, e.g. the Dirichlet process (Ferguson; 1973; Sethuraman; 1994), its two-parameter

generalization (Perman et al.; 1992; Pitman and Yor; 1992), the normalized inverse-Gaussian process (Favaro et al.; 2012) and the more general class of homogeneous normalized random measures with independent increments (Favaro et al.; 2016). In particular, the Dirichlet process is recovered when $\mathbf{v}_i \stackrel{\text{iid}}{\sim} \text{Be}(1, \theta)$, for some $\theta > 0$, and, as shown by Pitman (1996b), the resulting weights coincide with the corresponding size-biased permutation of them, an ideal feature for clustering (Pitman; 1996a). A different stick-breaking prior is the Geometric process, introduced by Fuentes-García et al. (2010). For this case, the decreasing ordering of the weights takes the form

$$\mathbf{w}_j = \lambda(1 - \lambda)^{j-1}, \quad j \geq 1,$$

for some $\lambda \sim \text{Be}(\alpha, \theta)$, with $\alpha, \theta > 0$. Here the random variables $(\mathbf{v}_i)_{i \geq 1}$ are completely dependent, indeed identical, unlike for the Dirichlet process. As mentioned above, the ordering of the weights, or lack of it, is of high relevance when using Bayesian nonparametric priors for density estimation and/or clustering. The dependence on only one random variable makes the Geometric process an attractive choice from a numerical point of view, and also makes it quite simple to generalize to non-exchangeable settings (Fuentes-García et al.; 2009; Mena et al.; 2011; Hatjispyros et al.; 2018). Furthermore, as shown by Bissiri and Ongaro (2014), both the Dirichlet and the Geometric processes have full support.

We propose a new class of stick-breaking distributions over Δ_∞ , featured by dependent l.v.'s driven by a strictly stationary Beta Markov chain, thus leading to a novel family of random probability measures, the Beta-Binomial stick-breaking (BBSB) priors. The Beta Markov chain in question has a dependence parameter which modulates the ordering of the corresponding weights, allowing BBSB priors to enjoy a good trade-off between weights identifiability and mixing. For extreme values of the dependence parameter, we find that the Dirichlet process and the Geometric process priors are particular cases of our model. Furthermore, using an extension of the aforementioned result by Bissiri and Ongaro (2014), we will see that BBSB priors also have full support.

The remaining part of the article is organized as follows: In Section 2 we present the construction of the Markov chain with $\text{Be}(\alpha, \theta)$ marginals. In here, we also analyse some special and limiting cases that will subsequently allow to recover the Dirichlet and Geometric processes. This Markov chain then assembles in Section 3 a sequence of l.v.'s, thus leading to Beta-Binomial stick-breaking priors. In Section 4 we derive a sampling scheme for density estimation and, in Section 5 we test it in simulated data. The proofs of the main results are deferred to the Appendix.

2. Beta-Binomial Markov chain

Following Pitt et al. (2002), given a density function $\pi_{\mathbf{v}, \mathbf{x}}(v, x)$ with marginals $\pi_{\mathbf{v}}(v)$ and $\pi_{\mathbf{x}}(x)$, and whose conditional distributions are $\pi_{\mathbf{v}|\mathbf{x}}(v|x)$ and $\pi_{\mathbf{x}|\mathbf{v}}(x|v)$, it is possible to construct two of reversible Markov chains $(\mathbf{v}_i)_{i \geq 1}$ and $(\mathbf{x}_i)_{i \geq 1}$ with stationary distributions $\pi_{\mathbf{v}}$ and $\pi_{\mathbf{x}}$ respectively. The construction considers

the law induced by $\mathbf{v}_1 \sim \pi_{\mathbf{v}}$, and $\{\mathbf{x}_i \mid \mathbf{v}_i\} \sim \pi_{\mathbf{x}|\mathbf{v}}(\cdot|\mathbf{v}_i)$, $\{\mathbf{v}_{i+1} \mid \mathbf{x}_i\} \sim \pi_{\mathbf{v}|\mathbf{x}}(\cdot|\mathbf{x}_i)$, for $i \geq 1$; where \mathbf{v}_{i+1} is conditionally independent of $(\mathbf{v}_1, \mathbf{x}_1, \dots, \mathbf{v}_{i-1}, \mathbf{x}_{i-1}, \mathbf{v}_i)$ given \mathbf{x}_i , and analogously \mathbf{x}_{i+1} is conditionally independent of $(\mathbf{v}_1, \mathbf{x}_1, \dots, \mathbf{v}_i, \mathbf{x}_i)$ given \mathbf{v}_{i+1} . Arising from the Beta-Binomial conjugate model, we take

$$\pi_{\mathbf{v},\mathbf{x}}(v, x) = \text{Bin}(x|\kappa, v)\text{Be}(v|\alpha, \theta),$$

for some $\alpha, \theta > 0$, $\kappa \in \{0, 1, \dots\}$, and where $\text{Bin}(0, p) = \delta_0$. Thus, the dependence induced by $\mathbf{v}_1 \sim \text{Be}(\alpha, \theta)$, and $\{\mathbf{x}_i \mid \mathbf{v}_i\} \sim \text{Bin}(\kappa, \mathbf{v}_i)$, $\{\mathbf{v}_{i+1} \mid \mathbf{x}_i\} \sim \text{Be}(\alpha + \mathbf{x}_i, \theta + \kappa - \mathbf{x}_i)$, for $i \geq 1$ generates Markov chains, $\mathbf{V} = (\mathbf{v}_i)_{i \geq 1}$ and $\mathbf{X} = (\mathbf{x}_i)_{i \geq 1}$, where the former has transition probabilities given by

$$\mathbb{P}[\mathbf{v}_i \in A | \mathbf{v}_{i-1}] = \int_A \sum_{x=0}^{\kappa} \text{Be}(s|\alpha + x, \theta + \kappa - x)\text{Bin}(x|\kappa, \mathbf{v}_{i-1})ds, \tag{4}$$

and stationary distribution $\text{Be}(\alpha, \theta)$, and the latter

$$\begin{aligned} \mathbb{P}[\mathbf{x}_i = x | \mathbf{x}_{i-1}] &= \int_0^1 \text{Bin}(x|\kappa, p)\text{Be}(p|\alpha + \mathbf{x}_{i-1}, \theta + \kappa - \mathbf{x}_{i-1})dp \\ &= \binom{\kappa}{x} \frac{(\alpha + \mathbf{x}_{i-1})_{x\uparrow}(\theta + \kappa - \mathbf{x}_{i-1})_{\kappa-x\uparrow}}{(\alpha + \theta + \kappa)_{\kappa\uparrow}}, \end{aligned} \tag{5}$$

where $(y)_{m\uparrow} = \prod_{j=0}^{m-1} (y + j)$, and its stationary distribution is

$$\mathbb{P}[\mathbf{x}_i = x] = \binom{\kappa}{x} \frac{(\alpha)_{x\uparrow}(\theta)_{\kappa-x\uparrow}}{(\alpha + \theta)_{\kappa\uparrow}}. \tag{6}$$

To any Markov chains, \mathbf{V} , \mathbf{X} and $(\mathbf{V}, \mathbf{X}) = (\mathbf{v}_i, \mathbf{x}_i)_{i \geq 1}$, we refer to them as *Beta*, *Binomial* and *Beta-Binomial* chains. See Nieto-Barajas and Walker (2002) and Mena and Walker (2009) for more on this kind of Markov chains. In what follows, we focus on the the Beta chain and some of its properties, specifically in how the parameter κ affects the dependence of the chain. This will be relevant for our construction of the nonparametric prior in the following section.

Proposition 2.1. *Let (\mathbf{V}, \mathbf{X}) be a Beta-Binomial chain with parameters (κ, α, θ) , then for the Beta chain, \mathbf{V} , and for every $i \geq 1$, we have the following conditional moments*

- a) $\mathbb{E}[\mathbf{v}_{i+1} | \mathbf{v}_i] = \frac{\alpha + \kappa \mathbf{v}_i}{\alpha + \theta + \kappa}$.
- b) $\text{Var}(\mathbf{v}_{i+1} | \mathbf{v}_i) = \frac{(\alpha + \kappa \mathbf{v}_i)(\theta + \kappa(1 - \mathbf{v}_i)) + \kappa \mathbf{v}_i(1 - \mathbf{v}_i)(\alpha + \theta + \kappa)}{(\alpha + \theta + \kappa)^2(\alpha + \theta + \kappa + 1)}$.
- c) $\text{Cov}(\mathbf{v}_i, \mathbf{v}_{i+1}) = \frac{\kappa \alpha \theta}{(\alpha + \theta)^2(\alpha + \theta + 1)(\alpha + \theta + \kappa)}$.
- d) $\rho_{\mathbf{v}_i, \mathbf{v}_{i+1}} = \frac{\text{Cov}(\mathbf{v}_i, \mathbf{v}_{i+1})}{\sqrt{\text{Var}(\mathbf{v}_i)}\sqrt{\text{Var}(\mathbf{v}_{i+1})}} = \frac{\kappa}{\alpha + \theta + \kappa}$.

Fixing the value of κ and increasing either α or θ , the correlation coefficient, $\rho_{\mathbf{v}_i, \mathbf{v}_{i+1}}$ goes to 0. Conversely, if we fix α and θ , for large values of κ , $\rho_{\mathbf{v}_i, \mathbf{v}_{i+1}} \approx 1$.

Also, if α and θ are very small with respect to κ

$$\mathbb{E}[\mathbf{v}_{i+1}|\mathbf{v}_i] \approx \mathbf{v}_i \quad \text{and} \quad \text{Var}(\mathbf{v}_{i+1}|\mathbf{v}_i) \approx \frac{2\mathbf{v}_i(1-\mathbf{v}_i)}{\kappa+1}.$$

Hence, intuition tells us that the conditional distribution of \mathbf{v}_{i+1} given \mathbf{v}_i , tends to $\delta_{\mathbf{v}_i}$, as κ grows, see Figure 1. The following result generalizes this intuition.

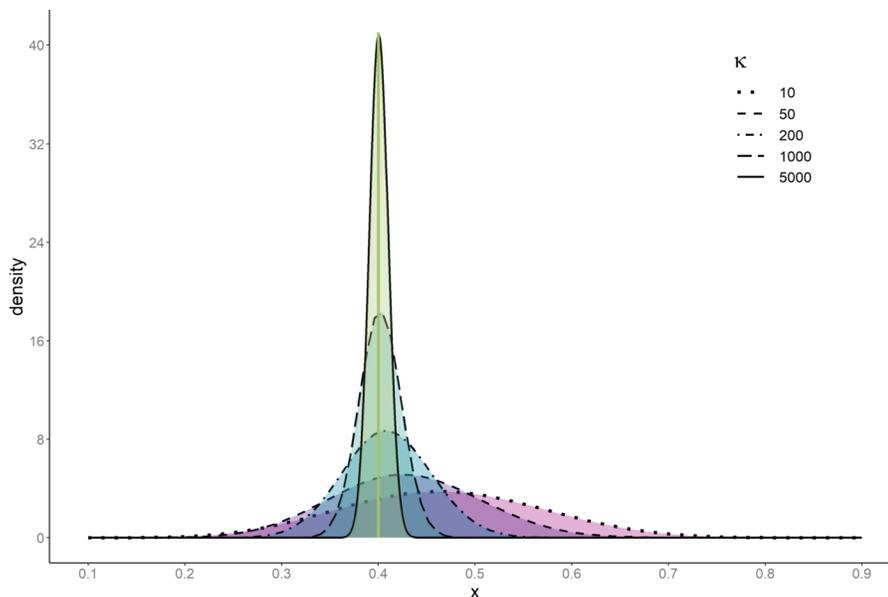


FIG 1. Conditional densities of \mathbf{v}_{i+1} given $\mathbf{v}_i = 0.4$, for distinct values of κ . We vary κ in the set $\{10, 50, 200, 1000, 5000\}$, the values of α and θ were fixed to 10.

Proposition 2.2. Let $\mathbf{V}^{(\kappa)} = (\mathbf{v}_i^{(\kappa)})_{i \geq 1}$ be a Beta-chain with parameters (κ, α, θ) .

- (i) For $\kappa = 0$, $\mathbf{V}^{(0)}$ is a sequence of i.i.d. random variables with distribution $\text{Be}(\alpha, \theta)$.
- (ii) As $\kappa \rightarrow \infty$, $\mathbf{V}^{(\kappa)}$ converges in distribution to $(\boldsymbol{\lambda}, \boldsymbol{\lambda}, \dots)$, where $\boldsymbol{\lambda} \sim \text{Be}(\alpha, \theta)$.

3. Beta-binomial stick-breaking prior

We call Beta-Binomial stick-breaking prior to any species sampling process, $\boldsymbol{\mu}$, with weights sequence as in (3) for some l.v.'s, \mathbf{V} , driven by a Beta chain with transition density (4). As usual, the parameters of the l.v.'s are inherited to the prior, adding to the latter, the diffuse probability measure, P_0 , as an additional

parameter. The first property to check is that the corresponding weights add up to one.

Proposition 3.1. *Let \mathbf{W} be as in equation (3), for some Beta chain, \mathbf{V} . Then*

$$\sum_{j \geq 1} \mathbf{w}_j \stackrel{a.s.}{=} 1.$$

Moreover, notice that for every $0 < \delta < \varepsilon < 1$ and $n \geq 1$, any Beta-Binomial chain, (\mathbf{V}, \mathbf{X}) , with parameters (κ, α, θ) , satisfies

$$\begin{aligned} \mathbb{P} \left[\bigcap_{i=1}^n (\delta < \mathbf{v}_i < \varepsilon) \right] &= \mathbb{E} \left[\prod_{i=1}^n \mathbb{P}[\delta < \mathbf{v}_i < \varepsilon | \mathbf{X}] \right] \\ &= \mathbb{E} \left[\mathbb{P}[\delta < \mathbf{v}_1 < \varepsilon | \mathbf{x}_1] \prod_{i=2}^n \mathbb{P}[\delta < \mathbf{v}_i < \varepsilon | \mathbf{x}_{i-1}, \mathbf{x}_i] \right] > 0, \end{aligned}$$

as conditionally given \mathbf{X} , the elements of \mathbf{V} are independent and Beta distributed. As shown by Bissiri and Ongaro (2014), the above observation shows that any Beta-Binomial prior has full support, and thus feasible for nonparametric inference. The following results, which follow from Proposition 2.2, motivate their study.

Theorem 3.2. *Let $\boldsymbol{\mu}^{(\kappa)}$ be a BBSB prior with parameters $(\kappa, \alpha, \theta, P_0)$ then*

- (i) *For $\kappa = 0$ and $\alpha = 1$, $\boldsymbol{\mu}^{(0)}$ is a Dirichlet process with parameters (θ, P_0) .*
- (ii) *For any α and θ fixed, as $\kappa \rightarrow \infty$, $\boldsymbol{\mu}^{(\kappa)}$ converges in distribution to the Geometric process, $\boldsymbol{\mu}$, with parameters (α, θ, P_0) .*

In terms of the ordering of the corresponding weights, we have the following corollary.

Corollary 3.3. *Let $(\mathbf{w}_j^{(\kappa)})_{j \geq 1}$ be as in equation (3), for some Beta chain, $(\mathbf{v}_i^{(\kappa)})_{i \geq 1}$, with parameters (κ, α, θ) . Then*

- (i) *For $\alpha = 1$, $\kappa = 0$, and any choice of θ , $(\mathbf{w}_j^{(\kappa)})_{j \geq 1}$ is size-biased ordered.*
- (ii) *For any choices of α and θ , and for every $j \geq 1$*

$$\lim_{\kappa \rightarrow \infty} \mathbb{P} \left[\mathbf{w}_{j+1}^{(\kappa)} < \mathbf{w}_j^{(\kappa)} \right] = 1.$$

If we fix $\alpha = 1$, the choice $\kappa = 0$ implies that $\mathbf{W} = \tilde{\mathbf{W}}$ is size-biased ordered. In general for such sequences $\mathbb{E}[\tilde{\mathbf{w}}_j] \geq \mathbb{E}[\tilde{\mathbf{w}}_{j+1}]$, even though $\tilde{\mathbf{w}}_j \geq \tilde{\mathbf{w}}_{j+1}$ does not occur with probability 1. On the other extreme, as $\kappa \rightarrow \infty$ we have the decreasing ordering of the Geometric weights $\mathbf{W} = \mathbf{W}^\downarrow$, which satisfy $\mathbb{P} \left[\mathbf{w}_j^\downarrow \geq \mathbf{w}_{j+1}^\downarrow \right] = 1$. Roughly speaking, by increasing the parameter κ , we make the weights sequence more likely to be decreasingly ordered. Figure 2 shows some simulations of $(\mathbf{w}_j)_{j=1}^{25}$ and their corresponding l.v.'s that illustrate the afore-

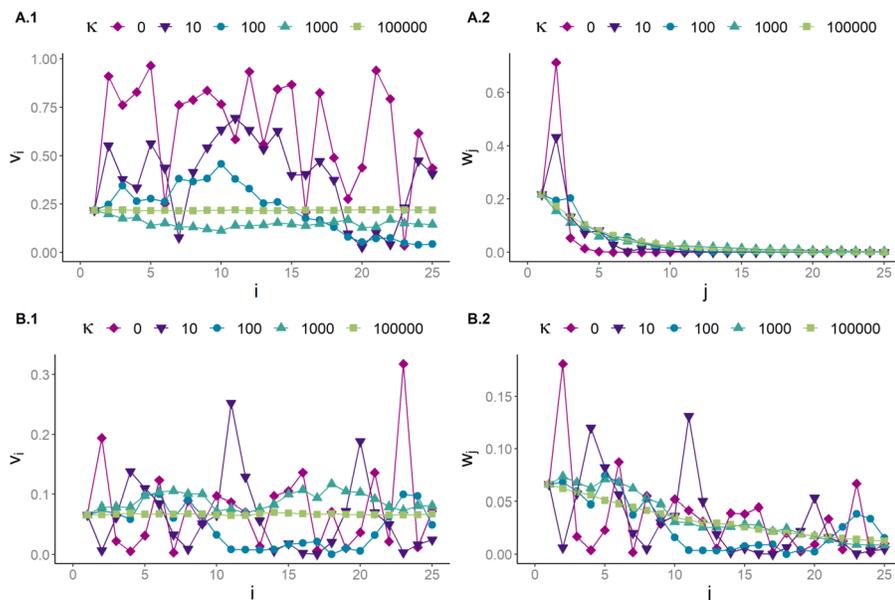


FIG 2. Simulations of $(\mathbf{w}_j)_{j=1}^{25}$ (A.2 and B.2) and their corresponding l.v.'s (A.1 and B.1 respectively) for distinct values of κ . For the Beta chains in A.1, we fixed $\alpha = 1$ and $\theta = 1$, for the ones in B.1 we used the same value of α , whilst $\theta = 10$. The chains in a single graph share the same initial r.v. for the sake of a simpler analysis.

mentioned behaviour. The initial value, \mathbf{v}_1 , of the Beta chain strongly affects the behaviour of the complete sequence of weights, this is particularly evident for large values of κ . Recall that if κ is sufficiently large we have $\mathbf{v}_1 \approx \mathbf{v}_2$, so for instance if \mathbf{v}_1 is close to 0, then $(1 - \mathbf{v}_1) \approx 1$ and $\mathbf{w}_2 = \mathbf{v}_2(1 - \mathbf{v}_1) \approx \mathbf{v}_1 = \mathbf{w}_1$, which means that if $\mathbf{v}_2 > \mathbf{v}_1$ even slightly, we might obtain $\mathbf{w}_2 > \mathbf{w}_1$. Alternatively, a large value of \mathbf{v}_1 , translates to a small value of $(1 - \mathbf{v}_1)$, so in order to obtain $\mathbf{w}_2 > \mathbf{w}_1$, it would require \mathbf{v}_2 to be significantly larger than \mathbf{v}_1 , which under the assumption that κ is large, is not very likely to happen, as $\mathbf{v}_1 \approx \mathbf{v}_2$. The same intuition is inherited to the subsequent indexes since we also have $\mathbf{v}_2 \approx \mathbf{v}_3 \approx \dots$, for large values of κ . Hence, the larger/smaller \mathbf{v}_1 is, the larger/smaller we expect \mathbf{v}_i to be, for $i > 1$. Moreover, for large values of the parameter θ we expect \mathbf{v}_1 to take small values, thus in general, a bigger value of θ requires an even larger value of κ , to induce a stochastically decreasing ordering of the weights.

3.1. Distribution of the number of groups

When working with any species sampling process, $\boldsymbol{\mu}$, such as a Dirichlet, BBSB or Geometric process..., a r.v. of interest is the number of distinct values, \mathbf{K}_n , that a sample $\{\gamma_1, \dots, \gamma_n\}$ driven by $\boldsymbol{\mu}$ exhibits. Although for some priors it

is possible to compute or characterize the probabilistic behaviour of \mathbf{K}_n (see for instance Pitman; 2006), in general this is not an easy task to do. Despite this, whenever it is feasible to obtain samples from the weights sequence, \mathbf{W} , as is the case of any BBSB prior, obtaining samples from \mathbf{K}_n can be easily achieved as follows: Sample n independent $U(0, 1)$ r.v.'s, $(\mathbf{u}_k)_{k=1}^n$, and $(\mathbf{w}_j)_{j=1}^\varphi$ where φ is some constant satisfying $\sum_{j=1}^\varphi \mathbf{w}_j > \max_k \mathbf{u}_k$. For $k \in \{1, \dots, n\}$ and $i \in \{1, \dots, \varphi\}$, let $\mathbf{d}_k = i$ if and only if $\sum_{j=1}^{i-1} \mathbf{w}_j < \mathbf{u}_k < \sum_{j=1}^i \mathbf{w}_j$ (with the convention that the empty sum equals 0) then the number of distinct values $(\mathbf{d}_1, \dots, \mathbf{d}_n)$ exhibits is precisely a sample from \mathbf{K}_n .

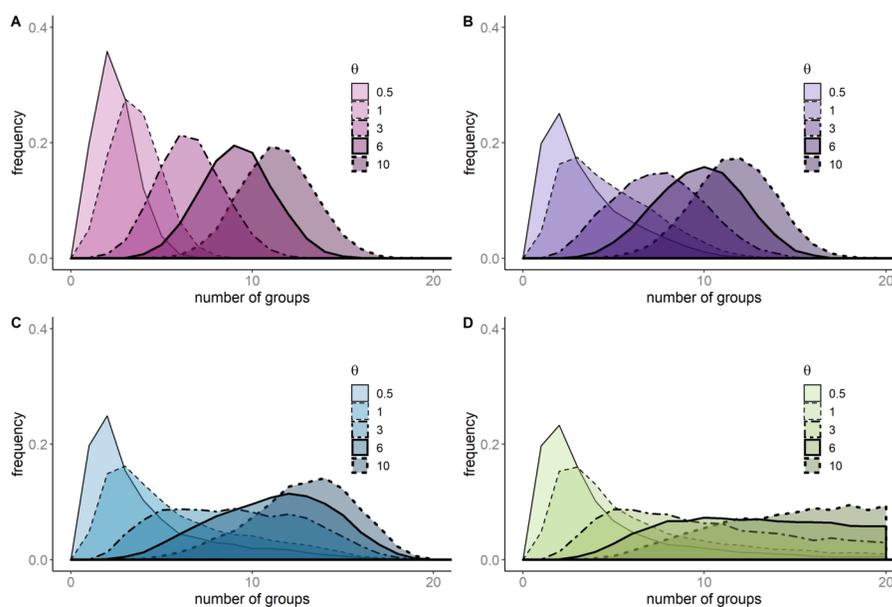


FIG 3. Frequency polygons of samples of size 10000 from \mathbf{K}_{20} for distinct values of κ and θ and fixing $\alpha = 1$. For the frequency polygons in A, B and C we fixed κ to 0, 10 and 100 respectively, whilst the frequency polygons in D correspond to the Geometric prior. For each fixed value of κ , we vary θ in the set $\{0.5, 1, 3, 6, 10\}$.

To understand how the parameters of a BBSB prior affect the distribution of \mathbf{K}_n , we sampled as aforementioned varying the values of κ , α and θ . Particularly, Figure 3(A) shows the distribution of \mathbf{K}_n corresponding to the Dirichlet process, for which is well known that $\mathbb{E}[\mathbf{K}_n]$ increases when θ grows. This location behavior is also observed for other fixed values of κ (B, C and D). Figures 3 and 4, illustrate how for fixed α and θ , an increment on κ contributes to the distribution of \mathbf{K}_n with a heavier right tail, and thus a larger mean and variance, say the prior on \mathbf{K}_n is less informative. In Figure 3, where we fixed $\alpha = 1$, it can be observed that for bigger values of θ , the distribution of \mathbf{K}_n is more sensitive to an increment of κ . The same can be seen in Figure 4, for fixed $\theta = 1$ and smaller values of α .

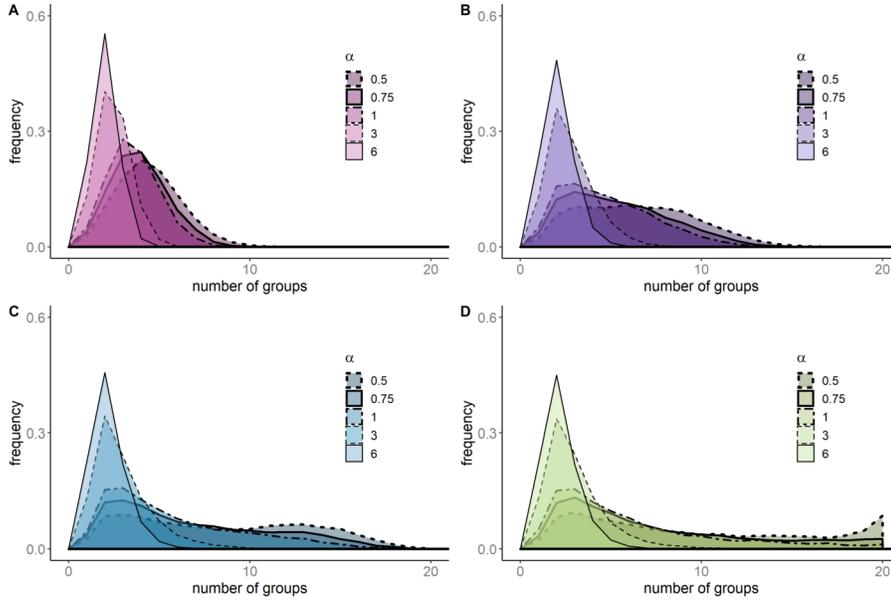


FIG 4. Frequency polygons of samples of size 10000 from \mathbf{K}_{20} for distinct values of κ and α and fixing $\theta = 1$. For the frequency polygons in A, B and C we fixed κ to 0, 10 and 100 respectively, whilst the frequency polygons in D correspond to the Geometric prior. For each fixed value of κ , we vary α in the set $\{0.5, 0.75, 1, 3, 6\}$.

4. Density estimation for Beta-Binomial mixtures

Given a BBSB prior, $\boldsymbol{\mu}$, and a density kernel, $g(\cdot|s)$, with parameter space S , we can consider BBSB mixtures. Namely, we can model elements in $\mathbf{y}^{(n)} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$ as i.i.d. sampled from the random density

$$\phi(\mathbf{y}) := \pi(\mathbf{y}|\mathbf{W}, \Xi) = \int_S g(\mathbf{y}|s) \boldsymbol{\mu}(ds) = \sum_{j \geq 1} \mathbf{w}_j g(\mathbf{y}|\boldsymbol{\xi}_j). \quad (7)$$

For MCMC implementation purposes, and following Walker (2007), this random density can be augmented as

$$\pi(\mathbf{y}, \mathbf{u}|\mathbf{W}, \Xi) = \sum_{j \geq 1} \mathbf{1}_{\{\mathbf{u} < \mathbf{w}_j\}} g(\mathbf{y}|\boldsymbol{\xi}_j), \quad (8)$$

where it can be easily deduced

$$\pi(\mathbf{u}|\mathbf{W}) = \sum_{j \geq 1} \mathbf{1}_{\{\mathbf{u} < \mathbf{w}_j\}}. \quad (9)$$

As in the Dirichlet process case, given \mathbf{u} , the number of components in the mixture is finite, with indexes being the elements of $A_{\mathbf{u}}(\mathbf{W}) = \{j : \mathbf{u} < \mathbf{w}_j\}$, that is

$$\pi(\mathbf{y}|\mathbf{u}, \mathbf{W}, \Xi) = \frac{1}{|A_{\mathbf{u}}(\mathbf{W})|} \sum_{j \in A_{\mathbf{u}}(\mathbf{W})} g(\mathbf{y}|\xi_j). \quad (10)$$

Using the membership variable \mathbf{d} , i.e. $\mathbf{d} = j$ iff \mathbf{y} is sampled from $g(\cdot|\xi_j)$, one can further consider the augmented joint density

$$\pi(\mathbf{y}, \mathbf{u}, \mathbf{d}|\mathbf{W}, \Xi) = \mathbf{1}_{\{\mathbf{u} < \mathbf{w}_{\mathbf{d}}\}} g(\mathbf{y}|\xi_{\mathbf{d}}). \quad (11)$$

The complete data likelihood based on a sample of size n from (11) is easily seen to be

$$\mathcal{L}_{\xi, \mathbf{w}}((\mathbf{y}_k, \mathbf{u}_k, \mathbf{d}_k)_{k=1}^n) = \prod_{k=1}^n \mathbf{1}_{\{\mathbf{u}_k < \mathbf{w}_{\mathbf{d}_k}\}} g(\mathbf{y}_k|\xi_{\mathbf{d}_k}), \quad (12)$$

and under the assumption P_0 has a density, p_0 , with respect to a suitable measure, the full joint density of every variable involved is

$$\begin{aligned} & \pi((\mathbf{y}_k, \mathbf{u}_k, \mathbf{d}_k)_{k=1}^n, (\mathbf{v}_i)_{i \geq 1}, (\xi_j)_{j \geq 1}) \\ &= \left(\prod_{k=1}^n \mathbf{1}_{\{\mathbf{u}_k < \mathbf{w}_{\mathbf{d}_k}\}} g(\mathbf{y}_k|\xi_{\mathbf{d}_k}) \right) \prod_{j \geq 1} p_0(\xi_j) \\ & \quad \times \left(\text{Be}(\mathbf{v}_1|\alpha, \theta) \prod_{i \geq 1} \sum_{x=0}^{\kappa} \text{Be}(\mathbf{v}_{i+1}|\alpha + x, \theta + \kappa - x) \text{Bin}(x|\kappa, \mathbf{v}_i) \right), \end{aligned} \quad (13)$$

recall $\mathbf{w}_{\mathbf{d}_k} = \mathbf{v}_{\mathbf{d}_k} \prod_{i=1}^{\mathbf{d}_k-1} (1 - \mathbf{v}_i)$ with the convention that the empty product equals 1.

4.1. Full conditionals

The full conditional distributions, required for posterior inference via a Gibbs sampler implementation, are proportional to (13), and given as follows.

1. Updating Ξ :

$$\pi(\xi_j | \dots) \propto p_0(\xi_j) \prod_{k \in D_j} g(\mathbf{y}_k|\xi_j), \quad j \geq 1,$$

where $D_j = \{k \geq 1 : \mathbf{d}_k = j\}$. If p_0 and g form a conjugate pair, the above is easy to sample from.

2. Updating \mathbf{V} and $\mathbf{U} = (\mathbf{u}_k)_{k=1}^n$ as a block:

$$\pi(\mathbf{U}, \mathbf{V} | \dots) \propto \left(\prod_{k=1}^n \mathbf{w}_{\mathbf{d}_k}^{-1} \mathbf{1}_{\{\mathbf{u}_k < \mathbf{w}_{\mathbf{d}_k}\}} \mathbf{w}_{\mathbf{d}_k} \right) \times \left(\text{Be}(\mathbf{v}_1 | \alpha, \theta) \prod_{i \geq 1} \sum_{x=0}^{\kappa} \text{Be}(\mathbf{v}_{i+1} | \alpha + x, \theta + \kappa - x) \text{Bin}(x | \kappa, \mathbf{v}_i) \right).$$

As $\mathbf{w}_{\mathbf{d}_k} = \mathbf{v}_{\mathbf{d}_k} \prod_{i=1}^{\mathbf{d}_k-1} (1 - \mathbf{v}_i)$, with the convention $\prod_{i=1}^0 (\cdot) = 1$, then

$$\pi(\mathbf{U}, \mathbf{V} | \dots) \propto \left[\prod_{k=1}^n \mathbf{w}_{\mathbf{d}_k}^{-1} \mathbf{1}_{\{\mathbf{u}_k < \mathbf{w}_{\mathbf{d}_k}\}} \right] \left[\mathbf{v}_1^{\alpha_1} (1 - \mathbf{v}_1)^{\theta_1} \text{Be}(\mathbf{v}_1 | \alpha, \theta) \right] \times \left[\prod_{i \geq 1} \sum_{x=0}^{\kappa} (\mathbf{v}_{i+1})^{\alpha_{i+1}} (1 - \mathbf{v}_{i+1})^{\theta_{i+1}} \text{Be}(\mathbf{v}_{i+1} | \alpha + x, \theta + \kappa - x) \text{Bin}(x | \kappa, \mathbf{v}_i) \right]$$

where

$$\alpha_i = \sum_{k=1}^n \mathbf{1}_{\{\mathbf{d}_k = i\}} \quad \text{and} \quad \theta_i = \sum_{k=1}^n \mathbf{1}_{\{\mathbf{d}_k > i\}}.$$

Recalling that for $m \in \mathbb{N}$, and $z > 0$, $\Gamma(m + z) = (z)_{m \uparrow} \Gamma(z)$, we obtain

$$\pi(\mathbf{U}, \mathbf{V} | \dots) \propto \left[\prod_{k=1}^n \text{U}(\mathbf{u}_k | 0, \mathbf{w}_{\mathbf{d}_k}) \right] \left[\text{Be}(\mathbf{v}_1 | \alpha + \alpha_1, \theta + \theta_1) \right] \times \left[\prod_{i \geq 1} \sum_{x=0}^{\kappa} \text{Be}(\mathbf{v}_{i+1} | \alpha_{i+1} + \alpha + x, \theta_{i+1} + \theta + \kappa - x) \right] \times \left[\frac{(\alpha + x)_{\alpha_{i+1} \uparrow} (\theta + \kappa - x)_{\theta_{i+1} \uparrow} \text{Bin}(x | \kappa, \mathbf{v}_i)}{(\alpha + \theta + \kappa)_{(\alpha_{i+1} + \theta_{i+1}) \uparrow}} \right],$$

with the convention $(z)_{0 \uparrow} = 1$. Thus, to update \mathbf{V} and \mathbf{U} , we first sample \mathbf{V} from

$$\pi(\mathbf{V} | \dots (\text{exclude } \mathbf{U}) \dots) \propto \left[\text{Be}(\mathbf{v}_1 | \alpha + \alpha_1, \theta + \theta_1) \right] \times \left[\prod_{i \geq 1} \sum_{x=0}^{\kappa} \text{Be}(\mathbf{v}_{i+1} | \alpha_{i+1} + \alpha + x, \theta_{i+1} + \theta + \kappa - x) \right] \times \left[\frac{(\alpha + x)_{\alpha_{i+1} \uparrow} (\theta + \kappa - x)_{\theta_{i+1} \uparrow} \text{Bin}(x | \kappa, \mathbf{v}_i)}{(\alpha + \theta + \kappa)_{(\alpha_{i+1} + \theta_{i+1}) \uparrow}} \right],$$

which can be normalized to a product of Beta densities mixtures, and latter sample \mathbf{U} from $\pi(\mathbf{U} | \dots) = \prod_{k=1}^n \text{U}(\mathbf{u}_k | 0, \mathbf{w}_{\mathbf{d}_k})$.

3. Updating $\mathbf{D} = (\mathbf{d}_k)_{k=1}^n$:

$$\pi(\mathbf{d}_k = j | \dots) \propto g(\mathbf{y}_k | \boldsymbol{\xi}_j) \mathbf{1}_{\{\mathbf{u}_k < \mathbf{w}_j\}}, \quad k \in \{1, \dots, n\},$$

which is a discrete distribution with finite support, hence easy to sample from.

Remark 4.1 (For the updating of $\boldsymbol{\Xi}$ and \mathbf{V}). *As it is well-known for this algorithm, we do not need to sample \mathbf{v}_j and $\boldsymbol{\xi}_j$ for every $j \geq 1$, it suffices to sample enough of them so that step 3 can take place. Explicitly, it suffices to sample $\boldsymbol{\xi}_j$ and \mathbf{v}_j for $j \leq \varphi$, where φ is a constant such that $\sum_{j=1}^{\varphi} \mathbf{w}_j \geq \max_k(1 - \mathbf{u}_k)$, then it is not possible that $\mathbf{w}_j > \mathbf{u}_k$ for any $k \leq n$ and $j > \varphi$.*

4.2. Posterior distribution analysis

Given $\left\{ (\boldsymbol{\xi}_j^{(t)})_j, (\mathbf{w}_j^{(t)})_j, (\mathbf{u}_k^{(t)})_k, (\mathbf{d}_k^{(t)})_k \right\}_{t=1}^T$, from $\{\boldsymbol{\Xi}, \mathbf{W}, \mathbf{U}, \mathbf{D} | \mathbf{y}^{(n)}\}$ obtained after T iterations of the Gibbs sampler, following (10) we estimate the density of the data by

$$\mathbb{E} \left[\phi | \mathbf{y}^{(n)} \right] \approx \frac{1}{T} \sum_{t=1}^T \frac{1}{n} \sum_{k=1}^n \frac{1}{|A_k^{(t)}|} \sum_{j \in A_k^{(t)}} g(\cdot | \boldsymbol{\xi}_j^{(t)}), \quad (14)$$

where $A_k^{(t)} = \{j : \mathbf{u}_k^{(t)} < \mathbf{w}_j^{(t)}\}$. Furthermore, we can also estimate the posterior distribution of $\{\mathbf{K}_n | \mathbf{y}^{(n)}\}$ through

$$\mathbb{P} \left[\mathbf{K}_n = m | \mathbf{y}^{(n)} \right] \approx \frac{1}{T} \sum_{t=1}^T \mathbf{1}_{\{\mathbf{K}_n^{(t)} = m\}}, \quad (15)$$

where $\mathbf{K}_n^{(t)}$ is the number of distinct values $(\mathbf{d}_k^{(t)})_k$ exhibits. As usual, when working with mixtures of densities, \mathbf{K}_n can be interpreted as the number of components of the mixture featured in the sample $\mathbf{y}^{(n)}$, that is the number of elements in $\{g(\cdot | \boldsymbol{\xi}_j)\}_{j \geq 1}$ such that \mathbf{y}_k is sampled from $g(\cdot | \boldsymbol{\xi}_j)$, for some $\mathbf{y}_k \in \mathbf{y}^{(n)}$. This way, the estimates (14) together with (15), give us information of how well a model performs for the given data set. Among the models for which (14) adjusts well to the data, those for which (15) favours smaller values of m might be preferred, as this means the model is mixing the components, $\{g(\cdot | \boldsymbol{\xi}_j)\}_{j \geq 1}$, more efficiently.

4.3. Posterior inference for the dependence parameter

In order to highlight the role of the dependence parameter, κ , we incorporate its posterior inference. Namely, we consider this parameter random and endow it with a prior distribution, $\boldsymbol{\kappa} \sim \pi_{\kappa}$. For this case, the likelihood (12) remains identical and the joint distribution (13) is multiplied by $\pi_{\kappa}(\boldsymbol{\kappa})$. It can easily be seen that, conditionally given $\boldsymbol{\kappa}$, the full conditionals $\{\pi(\boldsymbol{\xi}_j | \dots)\}_{j \geq 1}$, $\{\pi(\mathbf{d}_k | \dots)\}_{k=1}^n$ and $\pi(\mathbf{V}, \mathbf{U} | \dots)$ also remain the same. As to the full conditional of $\boldsymbol{\kappa}$ given the

rest of the r.v.'s, we have that

$$\pi(\boldsymbol{\kappa} = \kappa | \dots) \propto \pi_\kappa(\kappa) \prod_{i \geq 1} \sum_{x=0}^{\kappa} \text{Be}(\mathbf{v}_{i+1} | \alpha + x, \theta + \kappa - x) \text{Bin}(x | \kappa, \mathbf{v}_i), \quad (16)$$

which is easy to sample from if π_κ has finite support. Summarizing, at each iteration of the Gibbs sampler, we update $\boldsymbol{\Xi}$, \mathbf{V} , \mathbf{U} and \mathbf{D} as above and add a fourth step in which we sample $\boldsymbol{\kappa}$ from (16).

Finally, given the samples $(\boldsymbol{\kappa}^{(t)})_{t=1}^T$ obtained after T iterations of the Gibbs sampler, once the burn-in period has elapsed, we estimate the posterior distribution of $\boldsymbol{\kappa}$ by

$$\mathbb{P}[\boldsymbol{\kappa} = \kappa | \mathbf{y}^{(n)}] \approx \frac{1}{T} \sum_{t=1}^T \mathbf{1}_{\{\boldsymbol{\kappa}^{(t)} = \kappa\}}$$

5. Illustrations

In principle, every choice of κ leads to robust posterior MCMC estimates, after an appropriate burn-in period and enough valid iterations. However, depending on the sample, initial conditions, and current parameter values in the Gibbs sampler, the need to more/less ordered weights, thus different values of κ , might be required. To test the performance of BBSB priors for density estimation, we first conduct a small experiment in which we fix the value of κ to 0, 10, 100 and ∞ and compare the results provided by the 4 distinct models. Secondly, in order to choose the optimal value of κ for a dataset and given that the rest of the parameters are fixed, we place a prior distribution on the dependence parameter and analyse its posterior distribution. Here we also compare our models to another well-known stick-breaking prior, the Pitman-Yor process (Perman et al.; 1992; Pitman and Yor; 1992). In all cases we assume a Gaussian kernel with random location and scale parameters, i.e., for each $j \geq 1$, $\boldsymbol{\xi}_j = (\mathbf{m}_j, \mathbf{p}_j)$, and $g(\mathbf{y} | \boldsymbol{\xi}_j) = \text{N}(\mathbf{y} | \mathbf{m}_j, \mathbf{p}_j^{-1})$. To attain a conjugate pair for p_0 and g , we assume $p_0(\boldsymbol{\xi}_j) = \text{N}(\mathbf{m}_j | \vartheta, \tau \mathbf{p}_j^{-1}) \text{Ga}(\mathbf{p}_j | a, b)$, where $a = b = 0.5$, $\tau = 100$ and $\vartheta = n^{-1} \sum_{k=1}^n \mathbf{y}_k$.

5.1. Analysis for BBSB mixtures with fixed dependence parameter

For this exercise we simulated a data set (database 1) containing 200 observations and featuring 11 modes equally spaced. As it is well known for this type of data, and if the parameter θ is not carefully chosen, the Dirichlet mixture underestimates the number modes featured in the sample. Alternatively, Geometric mixtures do recognize every mode, but they tend to use a large number of mixture components. In order to study how BBSB priors perform in this context, and to compare them with the Dirichlet and Geometric processes, we fixed $\alpha = 1$, $\theta = 1$ and vary κ in the set $\{0, 10, 100, \infty\}$. No burn-in period was

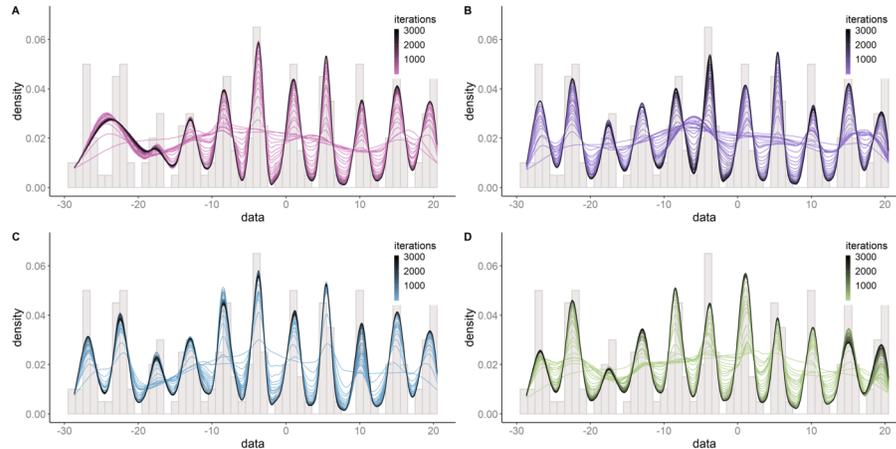


FIG 5. Evolution of the estimated densities for **database 1**, through the first 3000 iterations of the Gibbs sampler, for four distinct BBSB mixtures. The estimated densities in A, B, C and D correspond to BBSB mixtures with κ fixed to 0, 10, 100 and ∞ respectively, in the four cases $\alpha = \theta = 1$.

considered, so that one may analyse the number of iterations required by the model to provide a good estimate.

In Figure 5 we observe that the Dirichlet process (A) fails to recover the eleven modes featured in the dataset, the three remaining models are able to capture the 11 well-separated modes. In terms of the number of iterations required to recognize the modes, we observe that BBSB mixtures with larger values of κ (C and D) perform better. Consistently with the prior analysis of the number of groups, in Figure 6 we observe that the posterior mean and variance increase as κ does. Comparing Figures 5 and 6 we note that the model with $\kappa = 10$ (B) mixes better the components of the mixture than the other ones in the sense that fewer components were needed in order to capture every mode. Overall, the cases $\kappa = 10$ (B) and $\kappa = 100$ (C), seem to inherit desirable properties from the limiting cases, i.e. $\kappa = 0$ (A) and $\kappa = \infty$ (D). From the Dirichlet process they inherit a more efficient component mixing, while from the Geometric process they inherit the flexibility to adapt even if the parameter θ is not carefully chosen.

5.2. Analysis for BBSB mixtures with random dependence parameter

The main objective of this analysis is to determine the optimal value of κ for different datasets. To this aim, we first we consider a very simple data set (**database 2**) consisting of 200 observations that were sampled from a mixture of two Gaussian distributions. And secondly, we examine a more complicated set of data (**database 3**) that contains 200 observations sampled from a mixture of seven Gaussian kernels with distinct means, variances and weights,

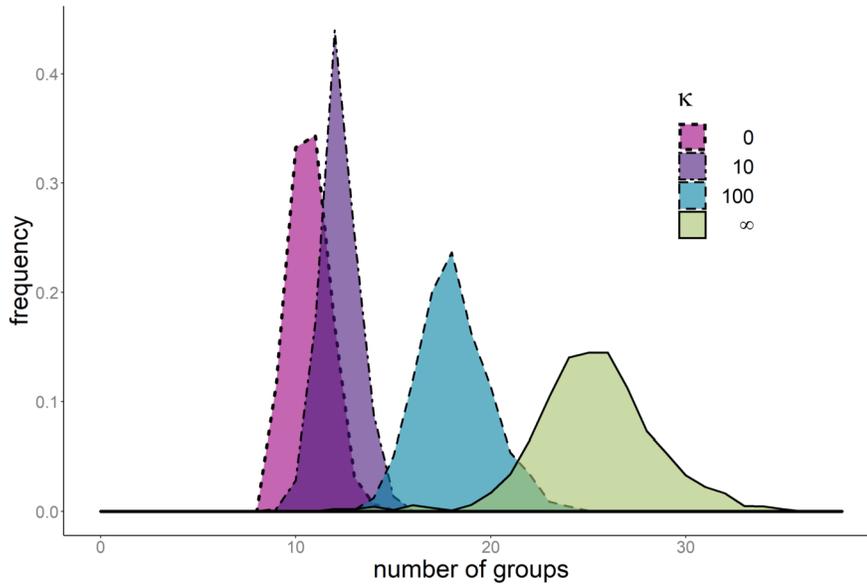


FIG 6. Frequency polygon of the estimated posterior distributions of \mathbf{K}_n given database 1 for the four BBSB mixtures which share the parameters $\alpha = \theta = 1$, and differ on the parameter κ , same one that varies in the set $\{0, 10, 100, \infty\}$.

this database was created and studied before by Lijoi et al. (2007). For each, database 2 and database 3, we study three BBSB mixtures with parameters α and θ fixed to distinct values, and compare the estimations with the ones provided by a Pitman-Yor mixture. Recall that this two-parameter generalization of the Dirichlet process has stick-breaking representation with independent l.v.'s $\mathbf{v}_i \sim \text{Be}(1 - \sigma, \theta + i\sigma)$ where $0 \leq \sigma < 1$ and $\theta > -\sigma$ (see for instance Perman et al.; 1992; Pitman and Yor; 1992; Pitman; 2006, for further details). In particular the Dirichlet process is recovered when $\sigma = 0$. For this mixture we fixed θ and consider the other parameter random with a uniform distribution over $[0, 1]$, this way the model is allowed to choose the best value of σ for the data set. In a similar spirit, for every BBSB mixture considered here, the parameter κ was considered random with a uniform prior distribution over $\{0, 1, \dots, 100\}$.

5.2.1. Results for database 2

In Figure 7 we observe that the estimated densities for the four mixtures adjust well to the data and do not differ significantly. In Figure 8 we see that every posterior distribution is asymmetrical, hence we will estimate the corresponding randomized parameter by the mode rather than the mean. For the BBSB models with parameter $\alpha = 1$ (A and B), the posterior mode of κ equals 0, suggesting that for this simple data set, the Dirichlet process is an excellent choice. In D we see that for the Pitman-Yor mixture the posterior distribution of σ also assigns

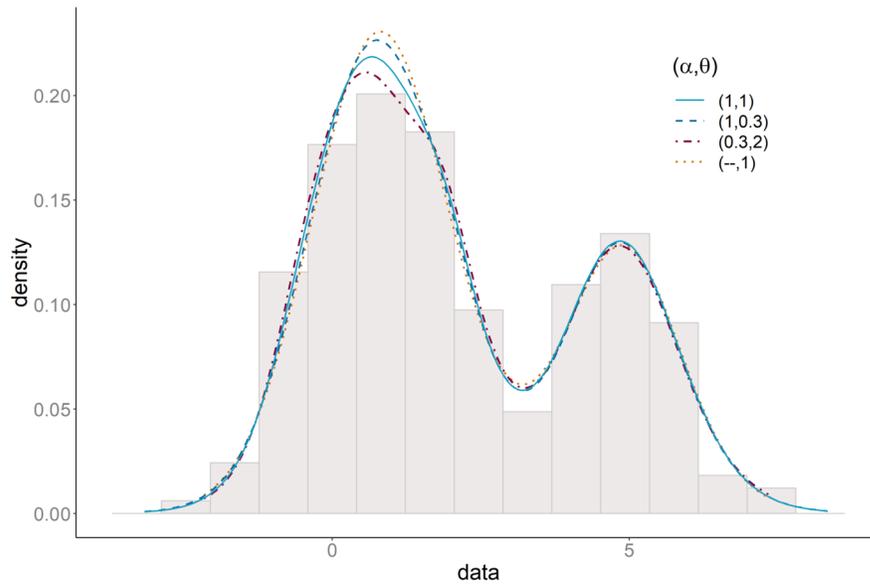


FIG 7. Estimated densities for database 2, taking into account 5000 iterations of the Gibbs sampler after a burn-in period of 3000, for three distinct BBSB mixtures with parameters (α, θ) fixed to $(1, 1)$, $(1, 0.3)$ and $(0.3, 2)$, and a Pitman-Yor mixture with parameter $\theta = 1$.

a bigger probability to values closer to 0, so it agrees with our models that the Dirichlet process adjust well to this data set. As for the BBSB model with $\alpha = 0.3$ and $\theta = 2$, we observe that the posterior distribution of κ (C) prefers a value bigger than 0. Explicitly, the posterior mode of this distribution is $\kappa = 6$. This could be due to the fact that for $\alpha = 0.3$ and $\theta = 2$ the stick-breaking mixture with completely independent l.v.'s is not a good choice for this dataset, so the BBSB mixture corrects this by adjusting the value of the dependence parameter.

5.2.2. Results for database 3

Insomuch as the distributions in Figure 10 are asymmetrical, once again we estimate the randomized parameter by the posterior modes. In the same figure we observe that the posterior distribution of κ for every BBSB mixture (A, B and C) favours values of κ that are bigger than 0, yet smaller than 50. Specifically, the posterior modes of κ for the BBSB models with (α, θ) fixed to $(1, 1.3)$, $(1, 0.3)$ and $(0.3, 2)$ are 12, 12 and 30, respectively. That is to say, in every case the model estimates that corresponding l.v.'s are dependent. In fact, if we insert the parameters $\alpha = 1, 1, 0.3$, $\theta = 1.3, 0.3, 2$, and the posterior mode of $\kappa = 12, 12, 30$, into Proposition 2.1 (d), we estimate the correlation coefficients of consecutive l.v.'s by 0.8992, 0.9023 and 0.9288, respectively. Notice that although the posterior modes of κ are not large, these choices affect greatly

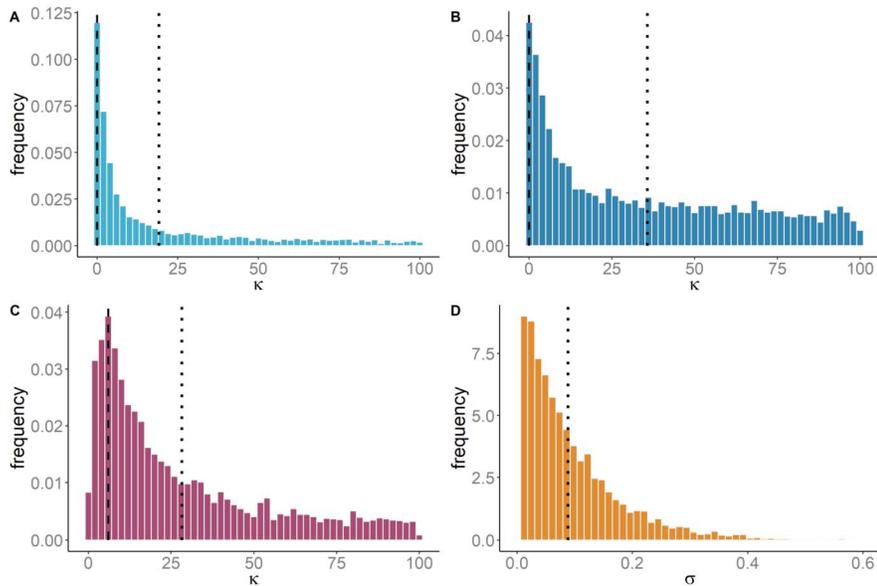


FIG 8. Posterior distributions of κ (A, B and C) for the BBSB mixtures with parameters (α, θ) fixed to $(1, 1)$, $(1, 0.3)$ and $(0.3, 2)$, respectively. D illustrates the posterior distribution of σ for the Pitman-Yor mixture with $\theta = 1$. The dotted and dashed lines indicate the posterior means and modes, respectively.

the dependence of the l.v.'s in question. In particular, for the couple of BBSB mixtures with $\alpha = 1$, this suggest the Dirichlet mixture is not the best choice. Among these two, for the one with $\theta = 1.3$, we chose this parameter so that for the Dirichlet prior $\mathbb{E}[\mathbf{K}_{200}] \approx 7$, which coincides with the number of actual modes featured in **database 3**. Even in this case, the posterior distribution of κ suggest that other BBSB models fit better than the Dirichlet mixture. As to the Pitman-Yor mixture, for which θ was also chosen as above, we see in Figure 10 (D) that the posterior distribution of σ favours values close to 0. Meaning that this model suggests that among the possibilities, the Dirichlet process is the best fit. However, if we concentrate in Figure 9 we see that the estimated densities by all three BBSB mixtures adjust well the data and recover the seven modes featuring the data set, whilst the Pitman-Yor model confuses the couple of modes in the left hand side of the figure. This suggests the class of BBSB mixtures offers a bigger capacity to adjust to the data by tuning the parameter κ , than the class of Pitman-Yor mixtures have by tuning the parameter σ .

6. Discussion

By using Beta chains as the l.v.'s of stick-breaking sequences, we were able to construct a new family of distributions over the infinite dimensional simplex, hence a new class of species sampling priors. The parameter, κ , that modulates

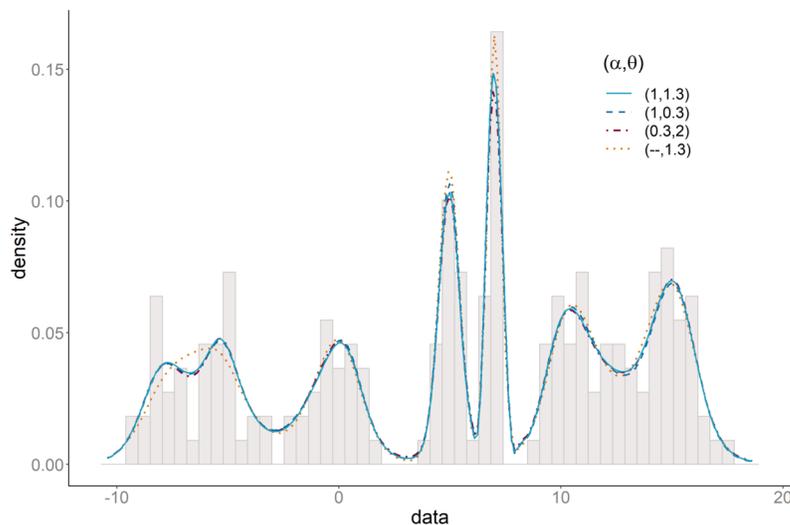


FIG 9. Estimated densities for database 3, taking into account 5000 iterations of the Gibbs sampler after a burn-in period of 3000, for three distinct BBSB mixtures with parameters (α, θ) fixed to $(1, 1.3)$, $(1, 0.3)$ and $(0.3, 2)$, and a Pitman-Yor mixture with parameter $\theta = 1.3$.

the dependence among the elements of the Beta chain, also modulates the ordering of the corresponding weights. While the choice $\kappa = 0$ and $\alpha = 1$ recovers the size-biased permutation of the weights of Dirichlet processes, as $\kappa \rightarrow \infty$, we recover the decreasing ordered weights of Geometric processes, both classes of processes being models of interest. This approach to define priors also allows the construction of random measures that are hybrids between Dirichlet and Geometric processes. Furthermore, how similar is the BBSB prior to one model or the other can also be tuned by the parameter κ . As to the prior distribution of \mathbf{K}_n , generally speaking, we found that a larger value of κ translates to a less informative prior. This in turn allows more flexible models in a density estimation context. In general the class of BBSB mixtures offers models with a great capacity to adapt to distinct data sets and models with a efficient component mixing. By endowing the parameter κ with a prior distribution, one can estimate its optimal value for a given data set, thus choose the BBSB mixture that admits the optimal balance between flexibility and efficient mixing.

The present work gives rise to interesting questions, such as how to characterize the distribution of \mathbf{K}_n for BBSB priors and analyse its asymptotic behaviour as $n \rightarrow \infty$, or even further study the underlying exchangeable partition probability functions. As to the orderings of the weights, it is also of interest to compute or approximate $\mathbb{P}[\mathbf{w}_j > \mathbf{w}_{j+1}]$ for a fixed value of κ , and to determine the rate at which $\mathbb{P}[\mathbf{w}_j > \mathbf{w}_{j+1}] \rightarrow 1$ as $\kappa \rightarrow \infty$. On a non-exchangeable context (e.g. Leisen and Griffin; 2017; De Iorio et al.; 2004), one could also use the Beta-Binomial transition to model dependence between two of more species sampling processes whose weights enjoy the stick-breaking decomposition. Hopefully, the

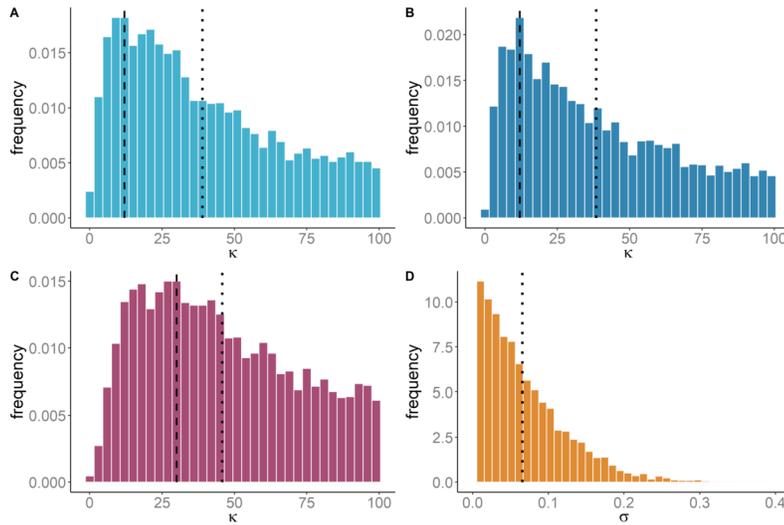


FIG 10. Posterior distributions of κ (A, B and C) for the BBSB mixtures with parameters (α, θ) fixed to $(1, 1.3)$, $(1, 0.3)$ and $(0.3, 2)$, respectively. D illustrates the posterior distribution of σ for the Pitman-Yor mixture with $\theta = 1.3$. The dotted and dashed lines indicate the posterior means and modes, respectively.

present paper motivates the study of stick-breaking sequences featuring dependent l.v.'s, that might even lead to other type of priors.

Acknowledgements

The first author gratefully thanks the support of CONACyT PhD scholarship program and CONACyT project 241195. The second author gratefully acknowledges the support of CONTEX project 2018-9B as well as the hospitality of the University of Bath, where part of the project was done, during a *Global Professor* research visit.

Appendix A.

Appendix A.1. Convergence of probability measures

To formally give the proof of the main results, we recall some topological details of measure spaces. For a Polish space S , with Borel σ -algebra $\mathcal{B}(S)$, we denote by $\mathcal{P}(S)$ the space of all probability measures over $(S, \mathcal{B}(S))$. A well-known metric on $\mathcal{P}(S)$ is the Lévy-Prokhorov metric given by

$$d_L(P, P') = \inf\{\varepsilon > 0 : P(A) \leq P'(A^\varepsilon) + \varepsilon, P'(A) \leq P(A^\varepsilon) + \varepsilon, \forall A \in \mathcal{B}(S)\}, \quad (17)$$

for any $P, P' \in \mathcal{P}(S)$, and where $A^\varepsilon = \{s \in S : d(s, A) < \varepsilon\}$, $d(s, A) = \inf\{d(a, s) : a \in A\}$ and d is some complete metric on S . For probability measures P, P_1, P_2, \dots it is said that P_n converges weakly to P , denoted by $P_n \xrightarrow{w} P$, whenever $\int_S f dP_n \rightarrow \int_S f dP$ for every continuous bounded function $f : S \rightarrow [0, \infty)$. This condition is known to be equivalent to $d_L(P_n, P) \rightarrow 0$, and to $\gamma_n \xrightarrow{d} \gamma$, whenever $\gamma_n \sim P_n$ and $\gamma \sim P$. $\mathcal{P}(S)$, equipped with the topology of weak convergence, is Polish again. Its Borel σ -field, $\mathcal{B}(\mathcal{P}(S))$, can equivalently be defined as the σ -algebra generated by all the projection maps $\{P \mapsto P(B) : B \in \mathcal{B}(S)\}$. In this sense the random probability measures (measurable mappings from a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ into $(\mathcal{P}(S), \mathcal{B}(S))$), $\boldsymbol{\mu}, \boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \dots$, are said to converge weakly, a.s. whenever $\boldsymbol{\mu}_n(\omega) \xrightarrow{w} \boldsymbol{\mu}(\omega)$, for every ω outside a \mathbb{P} -null set. Analogously, if $\int_S f d\boldsymbol{\mu}_n \xrightarrow{d} \int_S f d\boldsymbol{\mu}$ for every continuous bounded function $f : S \rightarrow [0, \infty)$, it is said that $\boldsymbol{\mu}_n$ converges weakly in distribution to $\boldsymbol{\mu}$, denoted by $\boldsymbol{\mu}_n \xrightarrow{dw} \boldsymbol{\mu}$. Evidently, $\boldsymbol{\mu}_n \xrightarrow{w} \boldsymbol{\mu}$ a.s. implies $\boldsymbol{\mu}_n \xrightarrow{dw} \boldsymbol{\mu}$, which, in turns is a necessary and sufficient condition for $\boldsymbol{\mu}_n \xrightarrow{d} \boldsymbol{\mu}$. For further details see for instance Parthasarathy (1967), Billingsley (1968) or Kallenberg (2017).

Appendix A.2. Proof of Proposition 2.1

a) Using elementary properties of conditional expectation and the fact that given \mathbf{x}_i , \mathbf{v}_{i+1} is conditionally independent of \mathbf{v}_i , we obtain

$$\mathbb{E}[\mathbf{v}_{i+1}|\mathbf{v}_i] = \mathbb{E}[\mathbb{E}[\mathbf{v}_{i+1}|\mathbf{x}_i]|\mathbf{v}_i] = \mathbb{E}\left[\frac{\alpha + \mathbf{x}_i}{\alpha + \theta + \kappa} \middle| \mathbf{v}_i\right] = \frac{\alpha + \kappa\mathbf{v}_i}{\alpha + \theta + \kappa}.$$

b) Notice that $\text{Var}(\mathbf{v}_{i+1}|\mathbf{v}_i) = \mathbb{E}[\text{Var}(\mathbf{v}_{i+1}|\mathbf{x}_i)|\mathbf{v}_i] + \text{Var}(\mathbb{E}[\mathbf{v}_{i+1}|\mathbf{x}_i]|\mathbf{v}_i)$, with

$$\text{Var}(\mathbb{E}[\mathbf{v}_{i+1}|\mathbf{x}_i]|\mathbf{v}_i) = \text{Var}\left(\frac{\alpha + \mathbf{x}_i}{\alpha + \theta + \kappa} \middle| \mathbf{v}_i\right) = \frac{\mathbf{v}_i(1 - \mathbf{v}_i)\kappa}{(\alpha + \theta + \kappa)^2}.$$

Now, note that

$$\begin{aligned} \mathbb{E}[(\alpha + \mathbf{x}_i)(\theta + \kappa - \mathbf{x}_i)|\mathbf{v}_i] &= \text{Cov}(\alpha + \mathbf{x}_i, \theta + \kappa - \mathbf{x}_i|\mathbf{v}_i) \\ &\quad + \mathbb{E}[\alpha + \mathbf{x}_i|\mathbf{v}_i]\mathbb{E}[\theta + \kappa - \mathbf{x}_i|\mathbf{v}_i] \\ &= -\text{Var}(\mathbf{x}_i|\mathbf{v}_i) + (\alpha + \kappa\mathbf{v}_i)(\theta + \kappa - \kappa\mathbf{v}_i) \\ &= -\kappa\mathbf{v}_i(1 - \mathbf{v}_i) + (\alpha + \kappa\mathbf{v}_i)(\theta + \kappa(1 - \mathbf{v}_i)) \end{aligned}$$

Hence

$$\begin{aligned} \mathbb{E}[\text{Var}(\mathbf{v}_{i+1}|\mathbf{x}_i)|\mathbf{v}_i] &= \mathbb{E}\left[\frac{(\alpha + \mathbf{x}_i)(\theta + \kappa - \mathbf{x}_i)}{(\alpha + \theta + \kappa)^2(\alpha + \theta + \kappa + 1)} \middle| \mathbf{v}_i\right] \\ &= \frac{-\kappa\mathbf{v}_i(1 - \mathbf{v}_i) + (\alpha + \kappa\mathbf{v}_i)(\theta + \kappa(1 - \mathbf{v}_i))}{(\alpha + \theta + \kappa)^2(\alpha + \theta + \kappa + 1)}, \end{aligned}$$

and we can conclude the proof of b),

$$\text{Var}(\mathbf{v}_{i+1}|\mathbf{v}_i) = \frac{(\alpha + \kappa \mathbf{v}_i)(\theta + \kappa(1 - \mathbf{v}_i)) + \kappa \mathbf{v}_i(1 - \mathbf{v}_i)(\alpha + \theta + \kappa)}{(\alpha + \theta + \kappa)^2(\alpha + \theta + \kappa + 1)}.$$

c) We first note that as a consequence of the joint reversibility of the Beta-Binomial chain, $\mathbf{v}_i \sim \text{Be}(\alpha + \mathbf{x}_i, \theta + \kappa - \mathbf{x}_i)$ conditionally given \mathbf{x}_i , thus

$$\mathbb{E}[\mathbf{v}_i \mathbf{v}_{i+1}] = \mathbb{E}[\mathbb{E}[\mathbf{v}_i \mathbf{v}_{i+1} | \mathbf{x}_i]] = \mathbb{E}[\mathbb{E}[\mathbf{v}_i | \mathbf{x}_i] \mathbb{E}[\mathbf{v}_{i+1} | \mathbf{x}_i]] = \mathbb{E} \left[\left(\frac{\alpha + \mathbf{x}_i}{\alpha + \theta + \kappa} \right)^2 \right],$$

conditioning on \mathbf{v}_i , we obtain

$$\begin{aligned} \mathbb{E} \left[\left(\frac{\alpha + \mathbf{x}_i}{\alpha + \theta + \kappa} \right)^2 \right] &= \mathbb{E} \left[\mathbb{E} \left[\left(\frac{\alpha + \mathbf{x}_i}{\alpha + \theta + \kappa} \right)^2 \middle| \mathbf{v}_i \right] \right] \\ &= \mathbb{E} \left[\frac{\alpha^2 + 2\alpha \mathbb{E}[\mathbf{x}_i | \mathbf{v}_i] + \mathbb{E}[\mathbf{x}_i^2 | \mathbf{v}_i]}{(\alpha + \theta + \kappa)^2} \right] \\ &= \frac{\alpha^2 + 2\alpha \kappa \mathbb{E}[\mathbf{v}_i] + \kappa \mathbb{E}[\mathbf{v}_i] + \kappa(\kappa - 1) \mathbb{E}[\mathbf{v}_i^2]}{(\alpha + \theta + \kappa)^2} \\ &= \left[\alpha^2 + \frac{\kappa(2\alpha^2 + \alpha)}{\alpha + \theta} + \frac{\kappa(\kappa - 1)\alpha(\alpha + 1)}{(\alpha + \theta)(\alpha + \theta + 1)} \right] (\alpha + \theta + \kappa)^{-2}, \end{aligned}$$

hence

$$\text{Cov}(\mathbf{v}_i, \mathbf{v}_{i+1}) = \mathbb{E}[\mathbf{v}_i \mathbf{v}_{i+1}] - \mathbb{E}[\mathbf{v}_i] \mathbb{E}[\mathbf{v}_{i+1}] = \frac{\kappa \alpha \theta}{(\alpha + \theta)^2 (\alpha + \theta + 1) (\alpha + \theta + \kappa)}.$$

d) The correlation simplifies as follows

$$\rho_{\mathbf{v}_i, \mathbf{v}_{i+1}} = \frac{\text{Cov}(\mathbf{v}_i, \mathbf{v}_{i+1})}{\sqrt{\text{Var}(\mathbf{v}_i)} \sqrt{\text{Var}(\mathbf{v}_{i+1})}} = \frac{\kappa}{\alpha + \theta + \kappa}.$$

Appendix A.3. Proof of Proposition 2.2

To prove Proposition 2.2 we need some preliminary results.

Lemma A.1 (Continuous mappings). *Let S and T be Polish spaces. Let $\boldsymbol{\eta}, \boldsymbol{\eta}_1, \dots$ be random elements taking values in S , with $\boldsymbol{\eta}_n \xrightarrow{d} \boldsymbol{\eta}$, and consider some measurable mappings f, f_1, f_2, \dots from S into T satisfying $f_n(s_n) \rightarrow f(s)$, for every $s_n \rightarrow s$ in S . Then $f_n(\boldsymbol{\eta}_n) \xrightarrow{d} f(\boldsymbol{\eta})$.*

Lemma A.2. *Let $\boldsymbol{\gamma}^n = (\gamma_1^n, \gamma_2^n, \dots)$, $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \dots)$ be random sequences taking values in a Polish space S . Then $\boldsymbol{\gamma}^n \xrightarrow{d} \boldsymbol{\gamma}$ if and only if*

$$(\gamma_1^n, \dots, \gamma_i^n) \xrightarrow{d} (\gamma_1, \dots, \gamma_i), \text{ for every } i \geq 1.$$

Lemmas A.2 and A.1 are well-known result in probability theory, see for instance Theorems 4.27 and 4.29, respectively, in Kallenberg (2002).

Lemma A.3. *Let S and T be Polish spaces. Consider some random elements $\gamma, \gamma_1, \gamma_2, \dots$ and $\eta, \eta_1, \eta_2, \dots$ taking values in S and T , respectively. Let ρ be the distribution of γ and ρ_n the distribution of γ_n , also consider some regular versions, $\pi(\cdot|\gamma)$ and $\pi_n(\cdot|\gamma_n)$, of $\mathbb{P}[\eta \in \cdot|\gamma]$ and $\mathbb{P}[\eta_n \in \cdot|\gamma_n]$ respectively. If $\rho_n \xrightarrow{w} \rho$ and for every $s_n \rightarrow s$ in S we have that $\pi_n(\cdot|s_n) \xrightarrow{w} \pi(\cdot|s)$, then $(\gamma_n, \eta_n) \xrightarrow{d} (\gamma, \eta)$.*

Proof. Let $g : S \times T \rightarrow \mathbb{R}$ be a continuous and bounded function. Define $f, f_1, f_2, \dots : S \rightarrow \mathbb{R}$ by

$$f_n(s) = \int g(s, t)\pi_n(dt|s) \quad \text{and} \quad f(s) = \int g(s, t)\pi(dt|s)$$

The first thing we will prove is that

$$f_n(s_n) \rightarrow f(s) \quad \text{as} \quad s_n \rightarrow s. \quad (18)$$

So let $s_n \rightarrow s$. Choose some random elements $\zeta, \zeta_1, \zeta_2, \dots$ with $\zeta_n \sim \pi_n(\cdot|s_n)$ and $\zeta \sim \pi(\cdot|s)$, this way, $\zeta_n \xrightarrow{d} \zeta$ by hypothesis. Define $h, h_1, h_2, \dots : T \rightarrow \mathbb{R}$ by $h_n(t) = g(s_n, t)$ and $h(t) = g(s, t)$. As g is continuous, we have that $h_n(t_n) = g(s_n, t_n) \rightarrow g(s, t) = h(t)$, for every $t_n \rightarrow t$ in T . By Lemma A.1 we obtain $h_n(\zeta_n) \xrightarrow{d} h(\zeta)$, which in turn implies

$$\begin{aligned} \int g(s_n, t)\pi_n(dt|s_n) &= \mathbb{E}[g(s_n, \zeta_n)] = \mathbb{E}[h_n(\zeta_n)] \\ &\rightarrow \mathbb{E}[h(\zeta)] = \mathbb{E}[g(s, \zeta)] = \int g(s, t)\pi(dt|s). \end{aligned}$$

Since $s_n \rightarrow s$ was arbitrary, this proves equation (18), which together with the hypothesis and by Lemma A.1 show that $f_n(\gamma_n) \xrightarrow{d} f(\gamma)$. Particularly,

$$\begin{aligned} \int \left(\int g(s, t)\pi_n(dt|s) \right) \rho_n(ds) &= \mathbb{E}[f_n(\gamma_n)] \\ &\rightarrow \mathbb{E}[f(\gamma)] = \int \left(\int g(s, t)\pi(dt|s) \right) \rho(ds). \end{aligned} \quad (19)$$

Note that the double integral in the left side of equation (19) coincides with $\mathbb{E}[g(\gamma_n, \eta_n)]$, whilst the one at the right side coincides with $\mathbb{E}[g(\gamma, \eta)]$. That is, we have proven that $\mathbb{E}[g(\gamma_n, \eta_n)] \rightarrow \mathbb{E}[g(\gamma, \eta)]$, for every continuous and bounded function $g : S \times T \rightarrow \mathbb{R}$. Or equivalently $(\gamma_n, \eta_n) \xrightarrow{d} (\gamma, \eta)$. \square

Lemma A.4. *Let $(\mathbf{x}_n)_{n \geq 1}$ be a sequence of random variables such that $\mathbf{x}_n \sim \text{Bin}(n, p_n)$ for every $n \geq 1$ and where $p_n \rightarrow p$ in $[0, 1]$. Then*

$$\frac{\mathbf{x}_n}{n} \xrightarrow{\mathcal{L}_2} p.$$

Proof. For $n \geq 1$,

$$\begin{aligned} \mathbb{E} \left[\left(\frac{\mathbf{x}_n}{n} - p \right)^2 \right] &= \frac{1}{n^2} \mathbb{E} [\mathbf{x}_n^2] - \frac{2p}{n} \mathbb{E}[\mathbf{x}_n] + p^2 \\ &= \frac{p_n(1-p_n)}{n} + (p_n - p)^2. \end{aligned} \quad (20)$$

By taking limits as $n \rightarrow \infty$ in (20) we obtain

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\left(\frac{\mathbf{x}_n}{n} - p \right)^2 \right] = 0.$$

□

Proof of Proposition 2.2.

(i) Insomuch as the corresponding spaces are Borel, we may construct on some probability space $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}})$ a Beta-Binomial chain $(\hat{\mathbf{V}}, \hat{\mathbf{X}})$ with parameters $(0, \alpha, \theta)$. Now, the elements of $\hat{\mathbf{V}}$ are conditionally independent given $\hat{\mathbf{X}}$, and given that $\kappa = 0$, $\hat{\mathbf{X}} \stackrel{a.s.}{=} (0, 0, \dots)$, so we may think of $\hat{\mathbf{X}}$ as if it was deterministic, which implies that the elements of $\hat{\mathbf{V}}$ must be independent and $\text{Be}(\alpha, \theta)$ distributed.

(ii) For every $\kappa \geq 1$, let $\mathbf{V}^{(\kappa)} = \left(\mathbf{v}_i^{(\kappa)} \right)_{i \geq 1}$ be a Beta chain with parameters (κ, α, θ) , and let $\pi_\kappa \left(\cdot | \mathbf{v}_i^{(\kappa)} \right)$ be some regular version of $\mathbb{P} \left[\mathbf{v}_{i+1}^{(\kappa)} \in \cdot | \mathbf{v}_i^{(\kappa)} \right]$ (which clearly does not depends on i). Further let $\boldsymbol{\lambda} \sim \text{Be}(\alpha, \theta)$ and fix $\pi(\cdot | \boldsymbol{\lambda}) = \delta_{\boldsymbol{\lambda}}$. The first thing we are interested in proving is that for every $p_\kappa \rightarrow p$ in $[0, 1]$ we have that

$$\pi_\kappa(\cdot | p_\kappa) \xrightarrow{w} \pi(\cdot | p). \quad (21)$$

So, let $p_\kappa \rightarrow p$ in $[0, 1]$, by Lemma A.4 and given that all the corresponding spaces are Borel, we may construct on a probability space $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}})$, with expectations $\hat{\mathbb{E}}[\cdot]$, some pairs of r.v.'s $(\hat{\mathbf{x}}_\kappa, \hat{\mathbf{v}}_\kappa)_{\kappa \geq 1}$ such that $\hat{\mathbf{x}}_\kappa \sim \text{Bin}(\kappa, p_\kappa)$, $\{\hat{\mathbf{v}}_\kappa | \hat{\mathbf{x}}_\kappa\} \sim \text{Be}(\alpha + \hat{\mathbf{x}}_\kappa, \theta + \kappa - \hat{\mathbf{x}}_\kappa)$, and $\hat{\mathbf{x}}_\kappa / \kappa \xrightarrow{a.s.} p$. Note that marginally $\hat{\mathbf{v}}_\kappa \sim \pi_\kappa(\cdot | p_\kappa)$ so to prove equation (21), it suffices to show $\hat{\mathbf{v}}_\kappa \xrightarrow{d} p$.

Conditionally given $\hat{\mathbf{x}}_\kappa$ the moment generator function of $\hat{\mathbf{v}}_\kappa$ is

$$\hat{\mathbb{E}} \left[e^{t \hat{\mathbf{v}}_\kappa} | \hat{\mathbf{x}}_\kappa \right] = 1 + \sum_{k=1}^{\infty} \left(\prod_{r=0}^{k-1} \frac{\alpha + \hat{\mathbf{x}}_\kappa + r}{\alpha + \theta + \kappa + r} \right) \frac{t^k}{k!}, \quad t \in \mathbb{R}. \quad (22)$$

By construction we have that $\hat{\mathbf{x}}_\kappa / \kappa \xrightarrow{a.s.} p$, which means that for every $r \geq 0$,

$$\frac{\alpha + \hat{\mathbf{x}}_\kappa + r}{\alpha + \theta + \kappa + r} = \left(\frac{\alpha + r}{\kappa} + \frac{\hat{\mathbf{x}}_\kappa}{\kappa} \right) \left(\frac{\alpha + \theta + r}{\kappa} + 1 \right)^{-1} \xrightarrow{a.s.} p, \quad (23)$$

as $\kappa \rightarrow \infty$. Hence by the tower property of conditional expectation, equations (22) and (23), and Lebesgue dominated convergence theorem (the corresponding

functions are dominated by e^t) we obtain

$$\begin{aligned} \lim_{\kappa \rightarrow \infty} \hat{\mathbb{E}} [e^{t\hat{\mathbf{v}}_\kappa}] &= \lim_{\kappa \rightarrow \infty} \hat{\mathbb{E}} \left[\hat{\mathbb{E}} [e^{t\hat{\mathbf{v}}_\kappa} | \hat{\mathbf{x}}_\kappa] \right] \\ &= \hat{\mathbb{E}} \left[1 + \sum_{k=1}^{\infty} \left(\prod_{r=0}^{k-1} \lim_{\kappa \rightarrow \infty} \frac{\alpha + \hat{\mathbf{x}}_\kappa + r}{\alpha + \theta + \kappa + r} \right) \frac{t^k}{k!} \right] \\ &= \hat{\mathbb{E}} \left[1 + \sum_{k=1}^{\infty} \frac{(pt)^k}{k!} \right] \\ &= e^{tp}, \end{aligned}$$

which proves altogether $\hat{\mathbf{v}}_\kappa \xrightarrow{d} p$ and equation (21).

Returning to the original Beta chains, we have that $\mathbf{v}_1^{(\kappa)} \stackrel{d}{=} \boldsymbol{\lambda}$ for every $\kappa \geq 1$, so trivially, $\mathbf{v}_1^{(\kappa)} \xrightarrow{d} \boldsymbol{\lambda}$, this together with equation (21) and the recursive application of Lemma A.3 allows us to obtain

$$\left(\mathbf{v}_1^{(\kappa)}, \dots, \mathbf{v}_i^{(\kappa)} \right) \xrightarrow{d} (\boldsymbol{\lambda}, \dots, \boldsymbol{\lambda}), \quad i \geq 1,$$

and by Lemma A.2 we can conclude $\mathbf{V}^{(\kappa)} = \left(\mathbf{v}_i^{(\kappa)} \right)_{i \geq 1} \xrightarrow{d} (\boldsymbol{\lambda}, \boldsymbol{\lambda}, \dots)$. □

Appendix A.4. Proof of Proposition 3.1

For sequences that enjoy the decomposition (3) we may equivalently prove that

$$\left(1 - \sum_{i=1}^j \mathbf{w}_i \right) = \prod_{i=1}^j (1 - \mathbf{v}_i) \xrightarrow{a.s.} 0,$$

as $j \rightarrow \infty$ (see for instance Ghosal and van der Vaart; 2017). Further, these r.v.'s are non-negative and bounded by 1, thus it is enough to show that

$$\lim_{j \rightarrow \infty} \mathbb{E} \left[\prod_{i=1}^j (1 - \mathbf{v}_i) \right] = 0. \tag{24}$$

As the corresponding spaces are Borel, (after possibly enlarging the original probability space) it is possible to construct a Binomial chain \mathbf{X} such that (\mathbf{V}, \mathbf{X}) defines a Beta-Binomial chain. Conditionally given $\mathbf{X} = \{\mathbf{x}_i\}_{i \geq 1}$, the elements of $\mathbf{V} = \{\mathbf{v}_i\}_{i \geq 1}$ are independent with, $\{\mathbf{v}_1 | \mathbf{x}_1\} \sim \text{Be}(\alpha + \mathbf{x}_1, \theta + \kappa - \mathbf{x}_1)$ and $\{\mathbf{v}_{i+1} | \mathbf{x}_i, \mathbf{x}_{i+1}\} \sim \text{Be}(\alpha + \mathbf{x}_i + \mathbf{x}_{i+1}, \theta + 2\kappa - \mathbf{x}_i - \mathbf{x}_{i+1})$, for $i \geq 1$. Hence

$$\begin{aligned} \mathbb{E} \left[\prod_{i=1}^j (1 - \mathbf{v}_i) \right] &= \mathbb{E} \left[\mathbb{E} \left[\prod_{i=1}^j (1 - \mathbf{v}_i) \middle| \mathbf{X} \right] \right] \\ &= \mathbb{E} \left[\mathbb{E}[(1 - \mathbf{v}_1) | \mathbf{x}_1] \prod_{i=2}^j \mathbb{E}[(1 - \mathbf{v}_i) | \mathbf{x}_{i-1}, \mathbf{x}_i] \right] \end{aligned}$$

$$= \mathbb{E} \left[\frac{\theta + \kappa - \mathbf{x}_1}{\alpha + \theta + \kappa} \prod_{i=2}^j \frac{\theta + 2\kappa - \mathbf{x}_i - \mathbf{x}_{i-1}}{\alpha + \theta + 2\kappa} \right].$$

Recalling that $0 \leq \mathbf{x}_i \leq \kappa$ a.s. we obtain

$$\frac{\theta}{\alpha + \theta + \kappa} \left(\frac{\theta}{\alpha + \theta + 2\kappa} \right)^{j-1} \leq \mathbb{E} \left[\prod_{i=1}^j (1 - \mathbf{v}_i) \right] \leq \frac{\theta + \kappa}{\alpha + \theta} \left(\frac{\theta + 2\kappa}{\alpha + \theta + 2\kappa} \right)^{j-1},$$

for every $j \geq 1$. Finally by taking limits as $j \rightarrow \infty$ in the last equation, (24) follows.

Appendix A.5. Proof of Theorem 3.2

To prove Theorem 3.2 we will first prove a couple of elementary results.

Lemma A.5. *Let S be a Polish space and fix some distinct $s_1, s_2, \dots \in S$, let $p = (p_1, p_2, \dots)$ and $q = (q_1, q_2, \dots)$ be elements of Δ_∞ and define $P = \sum_{j \geq 1} p_j \delta_{s_j}$ and $Q = \sum_{j \geq 1} q_j \delta_{s_j}$. Then for d_L as in equation (17)*

$$d_L(P, Q) \leq \sum_{j \geq 1} |p_j - q_j|.$$

Proof. Define $\varepsilon(p, q) = \sum_{j \geq 1} |p_j - q_j|$, by definition of d_L , it suffices to prove for all $A \in \mathcal{B}(S)$

$$P(A) \leq Q(A^{\varepsilon(p,q)}) + \varepsilon(p, q), \quad \text{and} \quad Q(A) \leq P(A^{\varepsilon(p,q)}) + \varepsilon(p, q). \quad (25)$$

So let $A \in \mathcal{B}(S)$ and set $M_A = \{j \geq 1 : s_j \in A\}$, then

$$\begin{aligned} P(A) &= \sum_{j \in M_A} P(\{s_j\}) = \sum_{j \in M_A} p_j \leq \sum_{j \in M_A} q_j + \sum_{j \in M_A} |p_j - q_j| \\ &\leq Q(A) + \varepsilon(p, q) \\ &\leq Q(A^{\varepsilon(p,q)}) + \varepsilon(p, q). \end{aligned}$$

Analogously, we have that $Q(A) \leq P(A^{\varepsilon(p,q)}) + \varepsilon(p, q)$. \square

Lemma A.6. *For fixed and distinct elements $s_1, s_2, \dots \in S$, the mapping,*

$$(w_1, w_2, \dots) \mapsto \sum_{j \geq 1} w_j \delta_{s_j},$$

from Δ_∞ into $\mathcal{P}(S)$ is continuous with respect to the weak topology.

Proof. Let $w^{(n)} = (w_1^{(n)}, w_2^{(n)}, \dots)$ and $w = (w_1, w_2, \dots)$ be any elements of Δ_∞ such that $w_j^{(n)} \rightarrow w_j$, for every $j \geq 1$. Define $P^{(n)} = \sum_{j \geq 1} w_j^{(n)} \delta_{s_j}$ and

$P = \sum_{j \geq 1} w_j \delta_{s_j}$. By Lemma A.5

$$d_L(P^{(n)}, P) \leq \sum_{j \geq 1} |w_j^{(n)} - w_j| \leq \sum_{j \geq 1} w_j^{(n)} + \sum_{j \geq 1} w_j = 2,$$

and by the general Lebesgue dominated convergence theorem we obtain

$$\lim_{n \rightarrow \infty} d_L(P^{(n)}, P) = \lim_{n \rightarrow \infty} \sum_{j \geq 1} |w_j^{(n)} - w_j| = \sum_{j \geq 1} \lim_{n \rightarrow \infty} |w_j^{(n)} - w_j| = 0,$$

which means that the mapping $(w_1, w_2, \dots) \mapsto \sum_{j \geq 1} w_j \delta_{s_j}$ is continuous. \square

Remark A.7. *Despite the choice of the metric, ρ , in Δ_∞ , as long as ρ generates the Borel σ -algebra, $\rho(w^{(n)}, w) \rightarrow 0$ implies $|w_j^{(n)} - w_j| \rightarrow 0$, for every $j \geq 1$. For this reason, in the above proof we did not discuss the details on the metric, of Δ_∞ , that is being used.*

Proof of Theorem 3.2. The proof of (i) follows directly from Proposition 2.2 (i). To prove (ii), note that by Proposition 2.2 (ii) and given that all the corresponding spaces are Borel, we can construct on a probability space $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}})$, Beta chains $\hat{\mathbf{V}}^{(\kappa)} = (\hat{v}_i^{(\kappa)})_{i \geq 1}$ with parameters (κ, α, θ) and a $\hat{\lambda} \sim \text{Be}(\alpha, \theta)$ such that $\hat{v}_i^{(\kappa)} \xrightarrow{\text{a.s.}} \hat{\lambda}$, as $\kappa \rightarrow \infty$, for every $i \geq 1$. Define also an independent sequence, $\hat{\Xi} = (\hat{\xi}_j)_{j \geq 1}$, with $\hat{\xi}_j \stackrel{\text{iid}}{\sim} P_0$. Now, for $\kappa \geq 1$ set

$$\hat{w}_j^{(\kappa)} = \hat{v}_j^{(\kappa)} \prod_{i=1}^{j-1} (1 - \hat{v}_i^{(\kappa)}), \quad j \geq 1, \quad \text{and} \quad \hat{\mu}^{(\kappa)} = \sum_{j \geq 1} \hat{w}_j^{(\kappa)} \delta_{\hat{\xi}_j},$$

with the empty product equating to 1, also set $\hat{\mu} = \sum_{j \geq 1} \hat{\lambda} (1 - \hat{\lambda})^{j-1} \delta_{\hat{\xi}_j}$, so

$$\hat{\mu}^{(\kappa)} \stackrel{d}{=} \hat{\mu}^{(\kappa)}, \quad \kappa \geq 1 \quad \text{and} \quad \hat{\mu} \stackrel{d}{=} \hat{\mu}. \tag{26}$$

As the mapping $(\hat{v}_1^{(\kappa)}, \dots, \hat{v}_j^{(\kappa)}) \mapsto \hat{w}_j^{(\kappa)}$ is continuous, we have that

$$\hat{w}_j^{(\kappa)} \xrightarrow{\text{a.s.}} \hat{\lambda} (1 - \hat{\lambda})^{j-1}, \quad j \geq 1.$$

For the sequence $\hat{\Xi}$, the diffuseness of P_0 implies that for $i \neq j$, $\hat{\xi}_i \neq \hat{\xi}_j$ a.s., since we are dealing with a countable number of random variables, there exist some $B \in \hat{F}$ such that $\hat{\mathbb{P}}[B] = 1$ and for every $\omega \in B$

$$\hat{w}_j^{(\kappa)}(\omega) \rightarrow \hat{\lambda}(\omega) (1 - \hat{\lambda}(\omega))^{j-1}, \quad j \geq 1, \quad \text{and} \quad \hat{\xi}_j(\omega) \neq \hat{\xi}_i(\omega), \quad i \neq j$$

By Lemma A.6

$$\sum_{j \geq 1} \hat{w}_j^{(\kappa)}(\omega) \delta_{\hat{\xi}_j(\omega)} \xrightarrow{w} \sum_{j \geq 1} \hat{\lambda}(\omega) (1 - \hat{\lambda}(\omega))^{j-1} \delta_{\hat{\xi}_j(\omega)}, \quad \omega \in B$$

that is, $\hat{\mu}^{(\kappa)} \xrightarrow{w} \hat{\mu}$ a.s., implying $\hat{\mu}^{(\kappa)} \stackrel{d}{\rightarrow} \hat{\mu}$. Finally, by equation (26), the result follows. \square

Appendix A.6. Proof of Corollary 3.3

The proof of (i) can be found in Theorem 1 by Pitman (1996a). To prove (ii) note that we may write

$$\mathbf{w}_1^{(k)} = \mathbf{v}_1^{(k)}, \quad \mathbf{w}_{j+1}^{(k)} = \frac{\mathbf{v}_{j+1}^{(k)} (1 - \mathbf{v}_j^{(k)})}{\mathbf{v}_j^{(k)}} \mathbf{w}_j^{(k)}, \quad j \geq 1,$$

hence

$$\mathbb{P} \left[\mathbf{w}_{j+1}^{(k)} < \mathbf{w}_j^{(k)} \right] = \mathbb{P} \left[\mathbf{v}_{j+1}^{(k)} (1 - \mathbf{v}_j^{(k)}) < \mathbf{v}_j^{(k)} \right].$$

By the second part of Proposition 2.2 and as the corresponding spaces are Borel, we may construct on some probability space, $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}})$, with expectations $\hat{\mathbb{E}}[\cdot]$, Beta chains, $(\hat{\mathbf{v}}_i^{(\kappa)})_{i \geq 1}$, with parameters (κ, α, θ) , and a $\hat{\boldsymbol{\lambda}} \sim \text{Be}(\alpha, \theta)$ satisfying

$$(\hat{\mathbf{v}}_i^{(\kappa)})_{i \geq 1} \rightarrow (\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\lambda}}, \dots) \quad \text{a.s.}$$

Then for $j \geq 1$, there exist $A \in \hat{\mathcal{F}}$ with $\hat{\mathbb{P}}[A] = 1$ and such that for every $\omega \in A$, $\hat{\mathbf{v}}_j^{(\kappa)}(\omega) \rightarrow \hat{\boldsymbol{\lambda}}(\omega)$ and $\hat{\mathbf{v}}_{j+1}^{(\kappa)}(\omega) \rightarrow \hat{\boldsymbol{\lambda}}(\omega)$. Fix $\omega \in A$, since $\hat{\boldsymbol{\lambda}}(\omega)(1 - \hat{\boldsymbol{\lambda}}(\omega)) < \hat{\boldsymbol{\lambda}}(\omega)$, we may choose κ' such that for every $\kappa > \kappa'$, $\hat{\mathbf{v}}_{j+1}^{(\kappa)}(\omega)(1 - \hat{\mathbf{v}}_j^{(\kappa)}(\omega)) < \hat{\mathbf{v}}_j^{(\kappa)}(\omega)$. As ω was chosen arbitrarily in A we have that $\mathbf{1} \left\{ \hat{\mathbf{v}}_{j+1}^{(\kappa)}(1 - \hat{\mathbf{v}}_j^{(\kappa)}) < \hat{\mathbf{v}}_j^{(\kappa)} \right\} \rightarrow 1$ a.s., as $\kappa \rightarrow \infty$. Finally, by Lebesgue dominated convergence theorem we obtain

$$\begin{aligned} \lim_{\kappa \rightarrow \infty} \mathbb{P} \left[\mathbf{w}_{j+1}^{(\kappa)} < \mathbf{w}_j^{(\kappa)} \right] &= \lim_{\kappa \rightarrow \infty} \mathbb{E} \left[\mathbf{1} \left\{ \mathbf{v}_{j+1}^{(\kappa)} (1 - \mathbf{v}_j^{(\kappa)}) < \mathbf{v}_j^{(\kappa)} \right\} \right] \\ &= \lim_{\kappa \rightarrow \infty} \hat{\mathbb{E}} \left[\mathbf{1} \left\{ \hat{\mathbf{v}}_{j+1}^{(\kappa)} (1 - \hat{\mathbf{v}}_j^{(\kappa)}) < \hat{\mathbf{v}}_j^{(\kappa)} \right\} \right] \\ &= \hat{\mathbb{E}} \left[\lim_{\kappa \rightarrow \infty} \mathbf{1} \left\{ \hat{\mathbf{v}}_{j+1}^{(\kappa)} (1 - \hat{\mathbf{v}}_j^{(\kappa)}) < \hat{\mathbf{v}}_j^{(\kappa)} \right\} \right] \\ &= 1. \end{aligned}$$

References

- Billingsley, P. (1968). *Convergence of Probability Measures*, Wiley series in probability and statistics, John Wiley and Sons Inc. [MR2893652](#)
- Bissiri, P. and Ongaro, A. (2014). On the topological support of species sampling priors, *Electronic Journal of Statistics* **8**(1): 861–882. [MR3229100](#)
- Blackwell, D. and MacQueen, J. (1973). Ferguson distributions via Pólya urn schemes, *The Annals of Statistics* **1**: 353–355. [MR0362614](#)
- De Iorio, M. D., Müller, P., Rosner, G. L. and MacEachern, S. N. (2004). An ANOVA model for dependent random measures, *Journal of the American Statistical Association* **99**: 205–215. [MR2054299](#)

- Favaro, S., Lijoi, A., Nava, C., Nipoti, B., Prünster, I. and Teh, Y. (2016). On the stick-breaking representation for homogeneous NRMIs, *Bayesian Analysis* **11**: 697–724. [MR3498043](#)
- Favaro, S., Lijoi, A. and Prünster (2012). On the stick-breaking representation of normalized inverse Gaussian priors, *Biometrika* **99**: 663–674. [MR2966776](#)
- Ferguson, T. (1973). A Bayesian analysis of some nonparametric problems, **1**(2): 209–230. [MR0350949](#)
- Fuentes-García, R., Mena, R. H. and Walker, S. G. (2009). A nonparametric dependent process for Bayesian regression, *Statistics & Probability Letters* **79**(4): 1112–1119. [MR2510777](#)
- Fuentes-García, R., Mena, R. H. and Walker, S. G. (2010). A new Bayesian nonparametric mixture model, *Communications in Statistics – Simulation and Computation* **39**(4): 669–682. [MR2785657](#)
- Ghosal, S. and van der Vaart, A. (2017). *Fundamentals of Nonparametric Bayesian Inference*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press. [MR3587782](#)
- Hatjispyros, J., Merktas, C., Nicolieris, T. and Walker, S. (2018). Dependent mixtures of geometric weights priors, *Computational Statistics and Data Analysis* **119**: 1–18. [MR3729519](#)
- Hjort, N., Holmes, C., Müller, P. and Walker, S. G. (2010). *Bayesian Non-parametrics*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press. [MR2722988](#)
- Ishwaran, H. and James, L. F. (2001). Gibbs sampling methods for stick-breaking priors, *Journal of the American Statistical Association* **96**(453): 161–173. [MR1952729](#)
- Ishwaran, H. and Zarepour, M. (2002). Exact and approximate sum representations for the dirichlet process, *Canadian Journal of Statistics* **30**(2): 269–283. [MR1926065](#)
- James, L. F., Lijoi, A. and Prünster, I. (2009). Posterior analysis for normalized random measures with independent increments, *Scandinavian Journal of Statistics* **36**(1): 76–97. [MR2508332](#)
- Kallenberg, O. (2002). *Foundations of Modern Probability*, second edn, Springer, New York. [MR1876169](#)
- Kallenberg, O. (2017). *Random Measures, Theory and Applications*, Vol. 77, Springer. [MR3642325](#)
- Leisen, F. and Griffin, J. (2017). Compound random measures and their use in Bayesian non-parametrics, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **79**(2): 525–545. [MR3611758](#)
- Lijoi, A., Mena, R. and Prünster, I. (2007). Controlling the reinforcement in Bayesian non-parametric mixture models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **69**(4): 715–740. [MR2370077](#)
- McCloskey, T. (1965). A model for the distribution of individuals by species in an environment, *Technical report*, Michigan State University Department of Statistics.
- Mena, R., Ruggiero, M. and Walker, S. G. (2011). Geometric stick-breaking processes for continuous-time Bayesian nonparametric modelling, *Journal of*

- Statistical Planning and Inference* **141**: 3217–3230. [MR2796026](#)
- Mena, R. and Walker, S. G. (2009). On a construction of Markov models in continuous time, *METRON – International Journal of Statistics* **LXVII**: 303–323.
- Mena, R. and Walker, S. G. (2015). On the Bayesian mixture model and identifiability, *Journal of Computational and Graphical Statistics* **24**: 1155–1169. [MR3432934](#)
- Nieto-Barajas, L. E. and Walker, S. G. (2002). Markov Beta and Gamma processes for modelling hazard rates, *Scandinavian Journal of Statistics* **29**(3): 413–424. [MR1925567](#)
- Parthasarathy, K. R. (1967). *Probability measures on metric spaces*, Academic press, New York. [MR0226684](#)
- Perman, M., Pitman, J. and Yor, M. (1992). Size-biased sampling of Poisson point processes and excursions, *Probability Theory and Related Fields* **92**(1): 21–39. [MR1156448](#)
- Pitman, J. (1996a). Random discrete distributions invariant under size-biased permutation, *Advances in Applied Probability* **28**(2): 525–539. [MR1387889](#)
- Pitman, J. (1996b). Some developments of the Blackwell-MacQueen urn scheme, in T. S. Ferguson, et al. (ed.), *Statistics, Probability and Game Theory; Papers in honor of David Blackwell*, Vol. 30 of *Lecture Notes-Monograph Series*, Institute of Mathematical Statistics, Hayward, California, pp. 245–267. [MR1481784](#)
- Pitman, J. (2006). *Combinatorial stochastic processes*, Vol. 1875 of *École d'été de probabilités de Saint-Flour*, first edn, Springer-Verlag Berlin Heidelberg, New York. [MR2245368](#)
- Pitman, J. and Yor, M. (1992). Arcsine laws and interval partitions derived from a stable subordinator, *Proceedings of the London Mathematical Society* **s3-65**(2): 326–356. [MR1168191](#)
- Pitt, M., Chatfield, C. and Walker, S. G. (2002). Constructing first order stationary autoregressive models via latent processes, *Scandinavian Journal of Statistics* **29**: 657–663. [MR1988417](#)
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors, *Statistica Sinica* **4**: 639–650. [MR1309433](#)
- Walker, S. G. (2007). Sampling the Dirichlet mixture model with slices, *Communications in Statistics-Simulation and Computation* **36**(1): 45–54. [MR2370888](#)