

# The Semi-Hierarchical Dirichlet Process and Its Application to Clustering Homogeneous Distributions\*

Mario Beraha<sup>†,‡</sup>, Alessandra Guglielmi<sup>§</sup>, and Fernando A. Quintana<sup>¶,||</sup>

**Abstract.** Assessing homogeneity of distributions is an old problem that has received considerable attention, especially in the nonparametric Bayesian literature. To this effect, we propose the semi-hierarchical Dirichlet process, a novel hierarchical prior that extends the hierarchical Dirichlet process of Teh et al. (2006) and that avoids the degeneracy issues of nested processes recently described by Camerlenghi et al. (2019a). We go beyond the simple yes/no answer to the homogeneity question and embed the proposed prior in a random partition model; this procedure allows us to give a more comprehensive response to the above question and in fact find groups of populations that are internally homogeneous when  $I \geq 2$  such populations are considered. We study theoretical properties of the semi-hierarchical Dirichlet process and of the Bayes factor for the homogeneity test when  $I = 2$ . Extensive simulation studies and applications to educational data are also discussed.

**Keywords:** Bayes factors, Bayesian nonparametrics, partial exchangeability, posterior consistency, homogeneity test.

**MSC2020 subject classifications:** Primary 60G57, 62G05, 62F15; secondary 62H30.

## 1 Introduction

The study and development of random probability measures in models that take into account the notion of data that are not fully exchangeable has sparked considerable interest in the Bayesian nonparametric literature. We consider here the notion of partial exchangeability in the sense of de Finetti (see de Finetti, 1938; Diaconis, 1988), which straightforwardly generalized the notion of an exchangeable sequence of random variables to the case of invariance under a restricted class of permutations. See also Camerlenghi et al. (2017) and references therein. In particular, our focus is on assessing whether two or more populations (or groups) of random variables can be con-

---

\*Fernando A. Quintana was supported by Fondecyt Grant 1180034. This work was supported by ANID – Millennium Science Initiative Program – NCN17.059.

<sup>†</sup>Department of Mathematics, Politecnico di Milano, [mario.beraha@polimi.it](mailto:mario.beraha@polimi.it)

<sup>‡</sup>Department of Computer Science, Università di Bologna

<sup>§</sup>Department of Mathematics, Politecnico di Milano, [alessandra.guglielmi@polimi.it](mailto:alessandra.guglielmi@polimi.it)

<sup>¶</sup>Department of Statistics, Pontificia Universidad Católica de Chile, [quintana@mat.uc.cl](mailto:quintana@mat.uc.cl)

<sup>||</sup>ANID – Millennium Science Initiative Program – Millennium Nucleus Center for the Discovery of Structures in Complex Data

sidered exchangeable rather than partially exchangeable, that is whether they arose from a common population/random distribution or not.

To be mathematically accurate, let us introduce partial exchangeability for a sequence of random variables. Let  $\mathbb{Y}$  denote a complete and separable metric space (i.e. a Polish space) with corresponding metric  $d$ . Let  $\mathcal{Y}$  denote the Borel  $\sigma$ -algebra of  $\mathbb{Y}$ , and  $\mathbb{P}_{\mathbb{Y}}$  denote the space of all probability measures on  $(\mathbb{Y}, \mathcal{Y})$ , with Borel  $\sigma$ -algebra  $\mathcal{P}_{\mathbb{Y}}$ . We will often skip reference to  $\sigma$ -algebras. A double sequence  $(y_{11}, y_{12}, y_{13}, \dots, y_{21}, y_{22}, y_{23}, \dots)$  of  $\mathbb{Y}$ -valued random variables, defined on a probability space  $(\Omega, \mathcal{F}, P)$  is called *partially exchangeable* if for all  $n, m \geq 1$  and all permutations  $(i(1), \dots, i(n))$  and  $(j(1), \dots, j(m))$  of  $(1, \dots, n)$  and  $(1, \dots, m)$  respectively, we have

$$\mathcal{L}(y_{11}, \dots, y_{1n}, y_{21}, \dots, y_{2m}) = \mathcal{L}(y_{1i(1)}, \dots, y_{1i(n)}, y_{2j(1)}, \dots, y_{2j(m)}).$$

Partial exchangeability can thus be conceptualized as invariance of the joint law above under the class of *all* permutations acting on the indices *within* each of the samples. Here and from now on, the distribution of a random element  $y$  is denoted by  $\mathcal{L}(y)$ .

The previous setting can be immediately extended to the case of  $I$  different populations or groups. By de Finetti's representation theorem (see the proof in Regazzini, 1991), partial exchangeability for the array of  $I$  sequences of random variables  $(y_{11}, y_{12}, \dots, y_{21}, y_{22}, \dots, y_{I1}, y_{I2}, \dots)$  is equivalent to

$$P(y_{ij} \in A_{ij}, j = 1, \dots, N_i, i = 1, \dots, I) = \int_{\mathbb{P}_{\mathbb{Y}}^I} \prod_{i=1}^I \prod_{j=1}^{N_i} p_i(A_{ij}) Q(dp_1, \dots, dp_I),$$

for any  $N_1, \dots, N_I \geq 1$  and Borel sets  $\{A_{ij}\}$  for  $j = 1, \dots, N_i$  and  $i = 1, \dots, I$ . In this case, de Finetti's measure  $Q$  is defined on the  $I$ -fold product space  $\mathbb{P}_{\mathbb{Y}}^I = \mathbb{P}_{\mathbb{Y}} \times \mathbb{P}_{\mathbb{Y}} \times \dots \times \mathbb{P}_{\mathbb{Y}}$ , and  $(p_1, p_2, \dots, p_I) \sim Q$ . The whole joint sequence of random variables is exchangeable if and only if  $Q$  gives probability 1 to the measurable set  $S = \{(p_1, p_2, \dots, p_I) \in \mathbb{P}_{\mathbb{Y}}^I : p_1 = p_2 = \dots = p_I\}$ .

Hence, partial exchangeability of data from different groups (or related studies) is a convenient context to analyze departures from exchangeability. While homogeneity of groups here amounts to full exchangeability, departures from this case may follow different directions, including independence of the population distributions  $p_1, p_2, \dots, p_I$ . However it could be interesting to investigate other types of departures from exchangeability beyond independence. The main goal here is to build a prior  $Q$  for  $(p_1, p_2, \dots, p_I)$  that is able to capture a wider range of different behaviors, not only restricting the analysis to assessing equality or independence among  $p_1, p_2, \dots, p_I$ . In the simplest case of  $I = 2$ , we just compare two distributions/populations, but we aim here at extending this notion to  $I > 2$  groups. In particular, we address the following issue: if the answer to the question of homogeneity within all these groups is negative, a natural question immediately arises, namely, can we assess the existence of homogeneity within certain populations? In other words, we would like to find clusters of internally homogeneous populations.

Vectors of dependent random distributions appeared first in Cifarelli and Regazzini (1978), but it was in MacEachern (1999) where a large class of dependent Dirichlet

processes was introduced, incorporating dependence on covariates through the atoms and/or weights of the stick-breaking representation. Following this line, De Iorio et al. (2004) proposed an ANOVA-type dependence for the atoms. These last two papers have generated an intense stream of research which is not our focus here. For a review of such constructions, see Quintana et al. (2020).

Our approach instead constructs a prior that explicitly considers a departure from exchangeability. Other authors have considered similar problems. Müller et al. (2004) and Lijoi et al. (2014a) constructed priors for the population distributions by these distributions with the addition of a common component. See also Hatjispyros et al. (2011), Hatjispyros et al. (2016) and Hatjispyros et al. (2018) for related models with increasing level of generalization. Several references where the focus is on testing homogeneity across groups of observations are available. Ma and Wong (2011) and Soriano and Ma (2017) propose the coupling optional Pólya tree prior, which jointly generates two dependent random distributions through a random-partition-and-assignment procedure similar to Pólya trees. The former paper considers both testing hypotheses from a global point of view, while the latter takes a local perspective on the two-sample hypothesis, detecting high resolution local differences. Bhattacharya and Dunson (2012) propose a Dirichlet process (DP) mixture model for testing whether there is a difference in distributions between groups of observations on a manifold. Both Chen and Hanson (2014) and Holmes et al. (2015) consider the two-sample testing problem, using a Pólya tree prior for the common distribution in the null, while the model for the alternative hypothesis assumes that the two population distributions are independent draws from the same Pólya tree prior. Their approaches differ in the way they specify the Pólya tree prior. Gutiérrez et al. (2019) consider a related problem, where a Bayesian nonparametric strategy to test for differences between a control group and several treatment regimes is proposed. Pereira et al. (2020) extend this idea to testing equality of distributions of paired samples, with a model for the joint distribution of both samples defined as a mixture of DPs with a spike-and-slab prior specification for its base measure.

Another traditional (and fruitful) approach for modeling data arising from a collection of groups or related studies involves the construction of hierarchical random prior probability measures. One of the first such examples in the Bayesian nonparametrics literature, is the well-known hierarchical DP mixtures introduced in Teh et al. (2006). Generalizations beyond the DP case are currently an active area of research, as testified by a series of recent papers dealing with various such hierarchical constructions; these include Camerlenghi et al. (2019b), Argiento et al. (2020) and Bassetti et al. (2019). See the discussion below.

Our first contribution is the introduction of a novel class of nonparametric priors that, just as discussed in Camerlenghi et al. (2019a), avoids the degeneracy issue of the nested Dirichlet process (NDP) of Rodríguez et al. (2008) that arises from the presence of shared atoms across populations. Indeed, Camerlenghi et al. (2019a) showed that under the NDP, if two populations share at least one common latent variable in the mixture model, then the model identifies the corresponding distributions as completely equal. To overcome the degeneracy issue, they resort to a latent nested construction in terms of normalized random measures that adds a shared random measure to draws

from the NDP. Instead, we use a variation of the hierarchical DP (HDP), that we term the semi-HDP, but where the baseline distribution is itself a mixture of a DP and a non-atomic measure. We will show that this procedure solves the degeneracy problem as well. While relying on a different model, Lijoi et al. (2020a) also propose to build on the HDP, combining it with the NDP, to overcome the degeneracy issue of nested processes.

Our second contribution is that the proposed model overcomes some of the practical and applied limitations of the latent nested approach by Camerlenghi et al. (2019a). As pointed out in Beraha and Guglielmi (2019), the latent nested approach becomes computationally burdensome in the case of  $I > 2$  populations. In contrast, implementing posterior inference for the semi-HDP prior does not require restrictions on  $I$ . We discuss in detail how to carry out posterior inference in the context of hierarchical models based on the semi-HDP.

A third contribution of this article is that we combine the proposed semi-HDP prior with a random partition model that allows different populations to be grouped in clusters that are internally homogeneous, i.e. arising from the same distribution. See an early discussion of this idea in the context of contingency tables in Quintana (1998). The far more general extension we aim for here is also useful from the applied viewpoint of finding out which, if any, of the  $I$  populations are internally homogeneous when homogeneity of the whole set does not hold. For the purpose of assessing global exchangeability, one may resort to discrepancy measures (Gelman et al., 1996); see also Catalano et al. (2021). In our approach, homogeneity corresponds to a point-null hypothesis about a discrete vector parameter, as we adopt a “larger” model for the alternative hypothesis within which homogeneity is nested. We discuss the specific case of adopting Bayes factors for the proposed test within the partial exchangeability framework. We show that the Bayes factor for this test is immediately available, and derives some of its theoretical properties.

The rest of this article is organized as follows. Section 2 gives some additional background that is relevant for later developments, presents the semi-HDP prior (Section 2.2) and, in particular, it describes a food court of Chinese restaurants with private and shared areas metaphor (Section 2.3). Section 3 studies several theoretical properties of the semi-HDP such as support, moments, the corresponding partially exchangeable partition probability function (in a particular case) and specially how the degeneracy issue is overcome under this setting. Section 3.3 specializes the discussion to the related issue of testing homogeneity when  $I = 2$  populations are present, and we study properties of the Bayes Factor for this test. Section 4 describes a computational strategy to implement posterior inference for the class of hierarchical models based on our proposed semi-HDP prior. Extensive simulations, with  $I = 2, 4$  and 100 populations are presented in Section 5. An application to an educational data set is discussed in Section 6. The article concludes with a discussion in Section 7. A Supplementary Materials file (Beraha et al., 2021), available online, collects the proofs for the theoretical results, together with additional formulas and figures, and a discussion on consistency for the Bayes Factor in the case of  $I = 2$  homogeneous populations. Code for posterior inference has been implemented in C++ and is available as part of the BayesMix library.<sup>1</sup>

---

<sup>1</sup><https://github.com/bayesmix-dev/bayesmix>.

## 2 Assessing Exchangeability Within a Partially Exchangeable Framework

While exchangeability can be explored in more generality, for clarity of exposition we set up our discussion in the context of continuous univariate responses, but extensions to, e.g. multivariate responses, can be straightforwardly accommodated in our framework.

### 2.1 A Common Home for Exchangeability and Partial Exchangeability

A flexible nonparametric model for each group can be constructed by assuming a mixture, where the mixing group-specific distribution  $G_i$  is a random discrete probability measure (r.p.m.), i.e.

$$y_{ij} | G_i \stackrel{\text{iid}}{\sim} p_i(\cdot) = \int_{\Theta} k(\cdot | \theta) G_i(d\theta), \quad j = 1, \dots, N_i, \quad (2.1)$$

where  $k(\cdot | \theta)$  is a density in  $\mathbb{Y}$  for any  $\theta \in \Theta$ , and  $G_i$  is, for example, a DP on  $\Theta$ . Note that, with a little abuse of notation,  $p_i$  in (2.1) and in the rest of the paper denotes the conditional population density of group  $i$  (before  $p_i$  represented the population distribution of group  $i$  in de Finetti’s theorem). In what follows, we will always assume that the parametric space is contained in  $\mathbb{R}^p$  for some positive integer  $p$ , and we will always assume the Borel  $\sigma$ -field  $\mathcal{B}(\Theta)$  of  $\Theta$ . Using the well-known alternative representation of the mixture in terms of latent variables, the previous expression is equivalent to assuming that for any  $i$ ,

$$y_{ij} | \theta_{ij} \stackrel{\text{iid}}{\sim} k(\cdot | \theta_{ij}), \quad \theta_{ij} | G_i \stackrel{\text{iid}}{\sim} G_i, \quad j = 1, \dots, N_i. \quad (2.2)$$

In this case, partial exchangeability of observations  $(y_{ij})_{ij}$  is equivalent to partial exchangeability of the latent variables  $(\theta_{ij})_{ij}$ . Hence exchangeability of observations  $(y_{ij})_{ij}$  is equivalent to the statement  $G_1 = G_2 = \dots = G_I$  with probability one.

In the next subsection we develop one of the main contributions of this paper, namely, the construction of a prior distribution  $\pi(G_1, \dots, G_I)$  such that there is positive prior probability that  $G_1 = G_2 = \dots = G_I$ , but avoiding the degeneracy issues discussed in Camerlenghi et al. (2019a) and that would arise if we assumed that  $(G_1, \dots, G_I)$  were distributed as the NDP by Rodríguez et al. (2008). Briefly,  $(G_1, \dots, G_I)$  is distributed as the NDP if

$$G_i | G \stackrel{\text{iid}}{\sim} G = \sum_{\ell=1}^{\infty} \pi_{\ell} \delta_{G_{\ell}^*}, \quad i = 1, \dots, I \quad \text{and} \quad G_{\ell}^* \stackrel{\text{iid}}{\sim} Q_0 = \mathcal{D}_{\gamma G_{00}},$$

i.e., the independent atoms in  $G$  are all drawn from a DP on  $\Theta$ , specifically  $G_{\ell}^* = \sum_{h=1}^{\infty} w_{h\ell} \delta_{\theta_{h\ell}}$ , with  $\theta_{h\ell} \stackrel{\text{iid}}{\sim} G_{00}$ , a probability measure on  $\Theta$ , and  $\alpha, \gamma > 0$ . The weights  $(\pi_j)_j$  and  $(w_{h\ell})_h$ ,  $\ell = 1, 2, \dots$ , are independently obtained from the usual stick-breaking construction, with parameters  $\alpha$  and  $\gamma$ , respectively. Here  $\mathcal{D}_{\gamma G_{00}}$  denotes the Dirichlet measure, i.e. the distribution of a r.p.m. that is a DP with measure parameter

$\gamma G_{00}$ . However, nesting discrete random probability measures produces degeneracy to the exchangeable case. As mentioned in Section 1, Camerlenghi et al. (2019a) showed that the posterior distribution degenerates to the exchangeable case whenever a shared component is detected, i.e., the NDP does not allow for sharing clusters among non-homogeneous populations. The problem is shown to affect any construction that uses nesting, and not just the NDP.

To overcome the degeneracy issue, while retaining flexibility, Camerlenghi et al. (2019a) proposed the so-called Latent Nested Nonparametric priors. These models involve a shared random measure that is added to the draws from a Nested Random Measure, hence accommodating for shared atoms. See also the discussion by Beraha and Guglielmi (2019). There are two key ideas in their model: (i) nesting discrete random probability measures as in the case of the NDP, and (ii) contaminating the population distributions with a common component as in Müller et al. (2004) and also, Lijoi et al. (2014b). The contamination aspect of the model yields dependence among population-specific random probability measures, and avoids the degeneracy issue pointed out by the authors, while the former accounts for testing homogeneity in multiple-sample problems. Their approach, however, becomes computationally burdensome in the case of  $I > 2$  populations, and it is not clear how to extend their construction to allow for the desired additional analysis, i.e. assessing which, if any, of the  $I$  populations are internally homogeneous when homogeneity of the whole set does not hold.

## 2.2 The Model

We present now a hierarchical model that allows us to assess homogeneity, while avoiding the undesired degeneracy issues and which further enables us to construct a grouping of populations that are internally homogeneous. To do so we create a hierarchical representation of distributions that emulates the behavior arising from an exchangeable partition probability function (EPPF; Pitman, 2006) such as the Pólya urn. But the main difference with previous proposals to overcome degeneracy is that we now allow for different populations to arise from the same distribution, while simultaneously incorporating an additional mechanism for populations to explicitly differ from each other.

Denote  $[I] = \{1, \dots, I\}$ . A partition  $S_1, \dots, S_k$  of  $[I]$  can be described by cluster assignment indicators  $\mathbf{c} = (c_1, \dots, c_I)$  with  $c_i = \ell$  iff  $i \in S_\ell$ . Assume this partition arises from a given EPPF. We introduce the following model for the latent variables in a mixture model such as (2.2). Let  $\mathbf{y}_i := (y_{i1}, \dots, y_{iN_i})$ , for  $i = 1, \dots, I$ . We assume that  $\mathbf{y}_1, \dots, \mathbf{y}_I$ , given all the population distributions  $F_1, \dots, F_I$  are independent, and furthermore arising from

$$y_{ij} | F_1, \dots, F_I, \mathbf{c} \stackrel{\text{iid}}{\sim} \int_{\Theta} k(\cdot | \theta) F_{c_i}(d\theta), \quad j = 1, \dots, N_i, \quad \text{for all } i, \quad (2.3)$$

$$\mathbf{c} \sim \pi_{\mathbf{c}}(c_1, \dots, c_I), \quad (2.4)$$

$$F_1, \dots, F_I | \tilde{P} \stackrel{\text{iid}}{\sim} \mathcal{D}_{\alpha \tilde{P}}, \quad (2.5)$$

$$\tilde{P} = \kappa G_0 + (1 - \kappa) \tilde{G}, \quad (2.6)$$

$$\tilde{G} \sim \mathcal{D}_{\gamma G_{00}}, \tag{2.7}$$

$$\kappa \sim \text{Beta}(a_\kappa, b_\kappa), \tag{2.8}$$

where  $\alpha, \gamma > 0$ . Thus the role of the population mixing distribution  $G_i$  in (2.1) – or, equivalently, in (2.2) – is now played by  $F_{c_i}$ . Observe that  $F_1, \dots, F_I$  in (2.5) play a role similar to the cluster specific parameters in more standard mixture models. Consider for example a case where  $I = 4$  and  $\mathbf{c} = (1, 2, 3, 1)$ . Under the above setting,  $F_1, F_2, F_3$  define a model for three different distributions, so that populations 1 and 4 share a common mixing distribution, and  $F_4$  is never employed.

Equation (2.5) means that conditionally on  $\tilde{G}$  each  $F_k$  is an independent draw from a DP prior with mean parameter  $\tilde{P}$  (and total mass  $\alpha$ ), i.e.  $F_k$  is a discrete r.p.m. on  $\Theta \subset \mathbb{R}^p$  for some positive integer  $p$ , with  $F_k = \sum_{h \geq 1} w_{kh} \delta_{\theta_{kh}^*}$  where for any  $k$  the weights are independently generated from a stick-breaking process,  $\{w_{kh}\}_h \stackrel{\text{iid}}{\sim} SB(\alpha)$ , i.e.

$$w_{k1} = \beta_{k1}, \quad w_{kh} = \beta_{ih} \prod_{j=1}^{h-1} (1 - \beta_{kj}) \quad \text{for } h = 2, 3, \dots, \quad \beta_{ij} \stackrel{\text{iid}}{\sim} \text{Beta}(1, \alpha),$$

and  $\{\theta_{kh}^*\}_h, \{\beta_{kh}\}_h$  are independent, with  $\theta_{kh}^* | \tilde{P} \stackrel{\text{iid}}{\sim} \tilde{P}$ . We assume the centering measure  $\tilde{P}$  in (2.6) to be a *contaminated draw*  $\tilde{G}$  from a DP prior, with centering measure  $G_{00}$ , with a fixed probability measure  $G_0$ . Both  $G_0$  and  $G_{00}$  are assumed to be absolutely continuous (and hence non-atomic) probability measures defined on  $(\Theta, \mathcal{B}(\Theta))$ .

By (2.7),  $\tilde{G} = \sum_{h \geq 1} p_h \delta_{\tau_h}$ , where  $\{p_h\}_h \sim SB(\gamma), \tau_h \stackrel{\text{iid}}{\sim} G_{00}$  are independent weights and location points. The model definition is completed by specifying  $\pi_c(c_1, \dots, c_I)$ . We assume that the  $c_i$ 's are (conditionally) i.i.d. draws from a categorical distribution on  $[I]$  with weights  $\boldsymbol{\omega} = (\omega_1, \dots, \omega_I)$ , i.e.  $c_i | \boldsymbol{\omega} \stackrel{\text{iid}}{\sim} \text{Cat}([I]; \boldsymbol{\omega})$ , where the elements of  $\boldsymbol{\omega}$  are non-negative and constrained to add up to 1. A convenient prior for  $\boldsymbol{\omega}$  is a finite dimensional Dirichlet distribution with parameter  $\boldsymbol{\eta} = (\eta_1, \dots, \eta_I)$ . Observe that distributions  $F_{c_1}, \dots, F_{c_I}$  allow us to cluster populations, so that there are at most  $I$  clusters and consequently  $F_1, \dots, F_I$  are all of the cluster distributions that ever need to be considered.

We say that a vector of random probability measures  $(F_1, \dots, F_I)$  has the semi-hierarchical Dirichlet process (semi-HDP) distribution if (2.5)–(2.7) hold, and we write  $(F_1, \dots, F_I) \sim \text{semi-HDP}(\alpha, \gamma, \kappa, G_0, G_{00})$ . It is straightforward to prove that, conditional on  $\kappa$  and eventual hyperparameters in  $G_0$  and  $G_{00}$ , the expectation of any  $F_i$  is  $\kappa G_0 + (1 - \kappa)G_{00}$  which further reduces to  $G_{00}$  if  $G_0 = G_{00}$ . Note that  $(F_1, \dots, F_I) \sim \text{semi-HDP}(\alpha, \gamma, \kappa, G_0, G_{00})$  defines an exchangeable prior over a vector of random probability measures.

We note several immediate yet interesting properties of the model. First, note that if  $\kappa = 1$  in (2.6), then all the atoms and weights in the representation of the  $F_i$ 's are independent and different with probability one, since the beta distribution and  $G_0$  are absolutely continuous. If  $\kappa = 0$ , then our prior (2.5)–(2.7) coincides with the Hierarchical Dirichlet Process in Teh et al. (2006). Since  $\tilde{G} = \sum_{h \geq 1} p_h \delta_{\tau_h}$ , then, with

positive probability, we have  $\theta_{kh}^* = \theta_{k'm}^* = \tau_\ell$  for  $k \neq k'$ , i.e. all the  $F_k$ 's share the same atoms in the stick-breaking representation of  $\tilde{G}$ . However, even when  $\kappa = 0$ ,  $F_k \neq F_j$  with probability one, as the weights  $\{w_{kh}\}_h$  and  $\{w_{jh}\}_h$  are different, since they are built from (conditionally) independent stick-breaking priors. This is precisely the feature that allows us to circumvent the degeneracy problem.

Second, our model introduces a vector parameter  $\mathbf{c}$ , which assists selecting each population distribution from the finite set  $F_1, \dots, F_I$ , in turn assumed to arise from the semi-HDP prior (2.5)–(2.7). The former allows two different populations to have the same distribution (or mixing measure) with positive probability, while the latter allows to overcome the degeneracy issue while retaining exchangeability. Indeed, as noted above,  $F_i$  and  $F_j$  may share atoms. The atoms in common arise from the atomicity of the base measure and we let the atomic component of the base measure to be a draw from a DP. The result is a very flexible model, that on one hand is particularly well-suited for problems such as density estimation, and on the other, can be used to construct clusters of the  $I$  populations, as desired.

### 2.3 A Restaurant Representation

To better understand the cluster allocation under model (2.3)–(2.7), we rewrite (2.3) introducing the latent variables  $\{\theta_{ij}\}$  as follows

$$y_{ij} \mid F_1, \dots, F_I, \mathbf{c}, \theta_{ij} \stackrel{\text{ind}}{\sim} k(\cdot \mid \theta_{ij}), \tag{2.9}$$

$$\theta_{i1}, \dots, \theta_{iN_i} \mid F_1, \dots, F_I, \mathbf{c} \stackrel{\text{iid}}{\sim} F_{c_i} \tag{2.10}$$

and  $\{\theta_{i\ell}\}_\ell \perp \{\theta_{jm}\}_m$  for  $i \neq j$ , conditionally on  $F_1, \dots, F_I$ .

We first derive the conditional law of the  $\theta_{ij}$ 's under (2.9)–(2.10), and (2.4)–(2.6), given  $\tilde{G}$ . All customers of group  $i$  enter restaurant  $r$  (such that  $c_i = r$ ). If group  $i$  is the first group entering restaurant  $r$ , then the usual Chinese Restaurant metaphor applies. Instead, let us imagine that group  $i$  is the last group entering restaurant  $r$  among those such that  $c_m = r$ . Upon entering the restaurant, the customer is presented with the usual Chinese Restaurant Process (CRP), so that

$$\theta_{ij} \mid \mathbf{c}, \{\theta_{mk}, \forall m : c_m = c_i = r\}, \theta_{i1}, \dots, \theta_{ij-1}, \tilde{G} \sim \sum_{\ell=1}^{H_r} \frac{n_{r\ell}}{\alpha + n_r} \delta_{\theta_{r\ell}^*} + \frac{\alpha}{\alpha + n_r} \tilde{P}, \tag{2.11}$$

that is the CRP when considering all the groups entering restaurant  $r$  as a single group. Here  $H_r$  denotes the number of tables in restaurant  $r$ , and  $n_{r\ell}$  is the number of customers who entered from restaurant  $r$  and are seating at table  $\ell$ . Moreover, note that  $\theta_{r\ell}^* \mid \tilde{G} \stackrel{\text{iid}}{\sim} \tilde{P}$ , so that, as in the HDP, there might be ties among the  $\theta_{r\ell}^*$  also when keeping  $r$  fixed. This is an important observation as the fact that there might be ties for different values of  $r \neq r'$  instead, is exactly what let us avoid the degeneracy to the exchangeable case. Note that (2.11) holds also for  $\theta_{i1}$ , i.e. the first customer in group  $i$ . In the following, we will use *clusters* or *tables* interchangeably. However, note that, unlike traditional CRPs,



the number of *clusters* does not coincide with the number of unique values in a sample. This point is clarified in Argiento et al. (2020), who introduce the notion of  $\ell$ -cluster, which is essentially the *table* in our restaurant metaphor.

Observe from (2.11) that when a new cluster is created, its label is sampled from  $\tilde{P}$ . In practice, we augment the parameter space with a new binary latent variable for each cluster, namely  $h_{r\ell}$ , with  $h_{r\ell} \stackrel{\text{iid}}{\sim} \text{Bernoulli}(\kappa)$ , so that

$$\theta_{r\ell}^* | h_{r\ell} = 1 \sim G_0 \quad \text{and} \quad \theta_{r\ell}^* | h_{r\ell} = 0, \tilde{G} \sim \tilde{G}.$$

Upon conditioning on  $\{h_{r\ell}\}$  it is straightforward to integrate out  $\tilde{G}$ . Indeed, we can write the joint distribution of  $\{\theta_{r\ell}^*, \forall r \forall \ell\}$ , conditional on  $\{h_{r\ell}\}$  as

$$\{\theta_{r\ell}^*\} | \{h_{r\ell}\}, \tilde{G} \sim \prod_{r,\ell} G_0(d\theta_{r\ell}^*)^{h_{r\ell}} \prod_{r,\ell} \tilde{G}(d\theta_{r\ell}^*)^{1-h_{r\ell}}.$$

Hence we see that  $\{\theta_{r\ell}^*, \forall r \forall \ell : h_{r\ell} = 0\}$  is a conditionally i.i.d sample from  $\tilde{G}$  (given all the  $h_{r\ell}$ 's and  $\tilde{G}$ ), so that we can write:

$$\theta_{r\ell}^* | h_{r\ell} = 0, \{\theta_{ij}^* : h_{ij} = 0\} \sim \sum_{k=1}^{H_0} \frac{m_{..k}}{m_{..} + \gamma} \delta_{\tau_k} + \frac{\gamma}{m_{..} + \gamma} G_{00} \tag{2.12}$$

and  $\tau_k \stackrel{\text{iid}}{\sim} G_{00}$ , where  $H_0$  denotes the number of tables in the common area in Figure 1, and  $m_{rk}$  denotes the cardinality of the set  $\{\theta_{r\ell}^* : \theta_{r\ell}^* = \tau_k\}$ . The dot subindex denotes summation over the corresponding subindex values. Hence, conditioning on all the  $(r, \ell)$  such that  $h_{r\ell} = 0$ , with  $r$  corresponding to a non-empty restaurant, we recover the Chinese Restaurant Franchise (CRF) that describes the HDP.

We can describe the previously discussed clustering structure in terms of a restaurant metaphor as the “food court of Chinese restaurants with private and shared areas”. Here, the  $\theta_{r\ell}^*$  correspond to the tables and  $\theta_{ij}$  to the customers. Moreover, a dish is associated to each table. Dishes are represented by the various  $\theta_{r\ell}^*$ 's. There is one big common area where tables are shared among all the restaurants and  $I$  additional “private” small rooms, one per restaurant, as seen in Figure 1. The common area accommodates tables arising from the HDP, i.e. those tables such that  $\tau_k \stackrel{\text{iid}}{\sim} G_{00}$ , while the small rooms host those tables associated to non empty restaurants, such that  $\theta_{r\ell}^* | h_{r\ell} = 1 \stackrel{\text{iid}}{\sim} G_0$ . All the customers of group  $i$  enter restaurant  $r$  (such that  $c_i = r$ ). Upon entering the restaurant, a customer is presented with a menu. The  $H_r$  dishes in the menu are the  $\theta_{r\ell}^*$ 's, and because  $\theta_{r\ell}^* \stackrel{\text{iid}}{\sim} \tilde{P}$ , there might be repeated dishes; see (2.11). The customer either chooses one of the dishes in the menu, with probability proportional to the number of customers who entered the same restaurant and chose that dish, or a new dish (that is not included in the menu yet) with probability proportional to  $\alpha$ ; again, see (2.11). If the latter option is chosen, with probability  $\kappa$  a new table is created in the restaurant-specific area,  $H_r$  is incremented by one and a new dish  $\theta_{rH_{r+1}}^*$  is drawn from  $G_0$ . With probability  $1 - \kappa$  instead, the customer is directed to the shared area, where (s)he chooses to seat in one of the occupied tables with a probability proportional to  $m_{..k}$ , i.e.

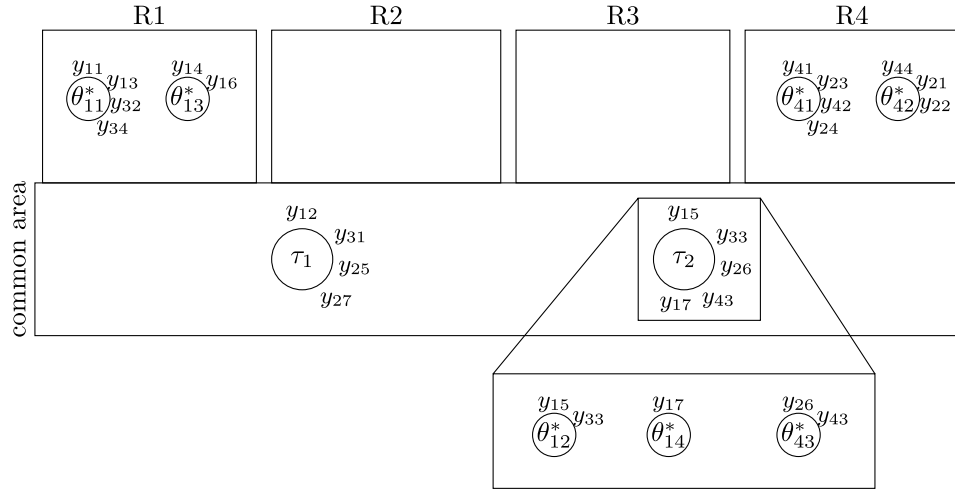


Figure 1: Restaurant representation of the semi-HDP allocation. In the image,  $\mathbf{c} = (1, 4, 1, 4)$  so that groups one and three enter in restaurant R1 while groups two and four enter in restaurant R4. In the “common area” two tables are represented,  $\tau_1$  and  $\tau_2$ . “Zooming” into  $\tau_2$  shows that there are three different  $\theta^*$ ’s associated to the value  $\tau_2$ , namely  $\theta_{12}^*$ ,  $\theta_{13}^*$  and  $\theta_{43}^*$ . The first two originate from R1, showing that it is possible to have ties among the  $\theta^*$ ’s even inside the same restaurant, while the table labeled  $\theta_{43}^*$  shows that it is possible to have ties across different restaurants.

the number of items in the menus (from all the restaurants) that are equal to dish  $\tau_k$ , or seats at a new table with a probability proportional to  $\gamma$ , as seen from (2.12). We point out that the choice of table in this case is made without any knowledge of which restaurant the dishes came from. Moreover, if the customer chooses to sit at a new table, we increment  $H_0$  by one and draw  $\tau_{H_0+1} \sim G_{00}$ ; we also increment  $H_r$  by one and set  $\theta_{rH_r+1}^* = \tau_{H_0+1}$ . Observe that in the original CRF metaphor, it is not the tables that are shared across restaurants, but rather the dishes. In our metaphor instead, we group together all the tables corresponding to the same  $\tau_h$  and place them in the shared area. This is somewhat reminiscent of the direct sampler scheme for the HDP. Nevertheless, observe that the bookkeeping of the  $m_{r,k}$ ’s is still needed. To exemplify this, in Figure 1 we report a “zoom” on a particular shared table  $\tau$ , showing that the  $\theta^*$ ’s associated to that table are still present in our metaphor, but can be collapsed into a single shared table when it is convenient.

### 3 Theoretical Properties of the Semi-HDP Prior

Here we develop additional properties of the proposed prior model. In particular, we study the topological support of the semi-HDP and show how exactly the degeneracy issue is resolved by studying the induced joint random partition model on the  $I$  populations.

### 3.1 Support and Moments

An essential requirement of nonparametric priors is that they should have large topological support; see Ferguson (1973). Let us denote by  $\pi_{\mathbf{G}}$  the probability measure on  $\mathbb{P}_{\Theta}^I$  corresponding to the prior distribution  $\pi(G_1, \dots, G_I)$  of the random vector  $(G_1, \dots, G_I)$  specified in (2.4)–(2.7), with  $G_i = F_{c_i}$ ; see (2.1). We show here that the prior probability measure  $\pi_{\mathbf{G}}$  has full weak support, i.e. given any point  $\mathbf{g} = (g_1, \dots, g_I)$  in  $\mathbb{P}_{\Theta}^I$ ,  $\pi_{\mathbf{G}}$  gives positive mass to any weak neighborhood  $\mathcal{U}(\mathbf{g}; \epsilon)$  of  $\mathbf{g}$ , of diameter  $\epsilon$ .

**Proposition 3.1** (Full Weak Support). *Let  $\pi_{\mathbf{G}}(g_1 \dots, g_I)$  be the prior probability measure on  $\mathbb{P}_{\Theta}^I$  defined by (2.4)–(2.7).*

- (a) *If  $G_0$  in (2.6) has full support on  $\Theta$  and  $0 < \kappa \leq 1$ , then  $\pi_{\mathbf{G}}(g_1 \dots, g_I)$  has full weak support.*
- (b) *If  $\kappa = 0$  and  $G_{00}$  in (2.7) has full support, then  $\pi_{\mathbf{G}}(g_1 \dots, g_I)$  has full weak support.*

*Proof.* See the Supplementary Materials file, Section 1. □

It is straightforward to show that in case where  $\pi_c(c_1, \dots, c_I)$  is exchangeable and  $P(c_i = \ell) = \omega_{\ell}$  for  $\ell = 1, \dots, I$  then (2.3)–(2.7) becomes, after marginalizing with respect to  $\mathbf{c}$ ,

$$y_{ij} | F_1, \dots, F_I \stackrel{\text{iid}}{\sim} \sum_{\mathbf{c}} \int_{\Theta} k(\cdot | \theta) F_{c_i}(d\theta) \pi_c(c_1, \dots, c_I) = \sum_{\ell=1}^I \omega_{\ell} \int_{\Theta} k(\cdot | \theta) dF_{\ell}(\theta).$$

In this case, the conditional marginal distribution of each observation can be expressed as a finite mixture of mixtures of the density  $k(\cdot | \theta)$  with respect to each of the random measures  $F_1, \dots, F_I$ , i.e. a finite mixture of Bayesian nonparametric mixtures.

We have mentioned above that in the case in which  $G_{00} = G_0$  in (2.6)–(2.7), the marginal law of  $F_i$  is  $G_0$ , and equivalently, for each  $A \in \mathcal{B}(\Theta)$ ,  $\mathbb{E}[F_i(A)] = G_0(A)$  for any  $i$ . In this case, the covariance between  $F_1$  and  $F_2$  is given by

$$\text{cov}(F_1(A), F_2(B)) = \frac{(1 - \kappa)^2}{1 + \gamma} (G_0(A \cap B) - G_0(A)G_0(B)).$$

See the Supplementary Materials file, Section 1, for the proof of these formulas. Note that, in the case of Hierarchical Normalized Completely Random Measures, and hence in the HDP, the covariance between  $F_1$  and  $F_2$  depends exclusively on the intensity of the random measure governing  $\tilde{G}$  (in the case of the DP the dependence is on  $\gamma$ ). For instance, see Argiento et al. (2020), Equation (5) in the Supplementary Material. Instead, in the Semi-HDP, an additional parameter can be used to tune such covariance: the weight  $\kappa$ . Indeed, as  $\kappa$  approaches 1, the two measures become more and more uncorrelated, the limiting case being full independence as discussed at the end of Section 2.2. In the Supplementary Materials file, Section 1, we also report an expression for the higher moments of  $F_i(A)$  for any  $i$ .

### 3.2 Degeneracy and Marginal Law

We now formalize the intuition given in Section 2.3 and show that our model, as defined in (2.3)–(2.7), does not incur in the degeneracy issue described by Camerlenghi et al. (2019a). The degeneracy of a nested nonparametric model refers to the following situation: if there are shared values (or atoms in the corresponding mixture model) across any two populations, then the posterior of these populations/random probabilities degenerates, forcing homogeneity across the corresponding samples. See also the discussion in Beraha and Guglielmi (2019).

From the food court metaphor described above, it is straightforward to see that degeneracy is avoided if two customers sit in the same table (of the common area) with positive probability, conditioning on the event that they entered from two different restaurants.

To see that this is so for the proposed model, let us consider the case  $I = 2$  and  $\theta_{i1} | F_1, F_2, \mathbf{c} = (1, 2) \sim F_i$ , for  $i = 1, 2$ . Marginalizing out  $(F_1, F_2)$ , this is equivalent to  $\theta_{11}, \theta_{21} | \tilde{G}, \mathbf{c} = (1, 2) \stackrel{\text{iid}}{\sim} wG_0 + (1 - w)\tilde{G}$ . Now, since  $G_0$  is absolutely continuous,  $\{\theta_{11} = \theta_{21}\}$  if and only if (i)  $\theta_{11}$  and  $\theta_{21}$  are sampled i.i.d. from  $\tilde{G}$ ; and (ii) we have a tie (which arises from the Pólya-urn scheme), i.e.  $\theta_{21} = \tau_1 = \theta_{11}$  and  $\tau_1 \sim G_{00}$ . This means that  $\theta_{11}$ , the first customer, sits in a table of the common area, an event that happens with probability  $1 - \kappa$  since she is the first one in the whole system, and  $\theta_{21}$  decides to sit in the common area (with probability  $1 - \kappa$ ) and subsequently decides to sit at the same table of  $\theta_{11}$  (which happens with probability  $\frac{1}{\gamma+1}$ ). Summing up we have that  $p(\theta_{11} = \theta_{21} | \mathbf{c} = (1, 2)) = (1 - \kappa)^2 / (1 + \gamma)$  which is strictly positive if  $\kappa < 1$ . Hence, by Bayes' rule, we have that

$$P(c_1 \neq c_2 | \theta_{11} = \theta_{21}) = \frac{P(\theta_{11} = \theta_{21} | c_1 \neq c_2)P(c_1 \neq c_2)}{\sum_{i,j} P(\theta_{11} = \theta_{21} | \mathbf{c} = (i, j))P(\mathbf{c} = (i, j))} > 0.$$

Moreover, when  $\kappa = 1$  we find the same degeneracy issue described in Camerlenghi et al. (2019a), as proved in Proposition 3.2 below.

To get a more in-depth look at these issues, we follow Camerlenghi et al. (2019a) and study properties of the partially exchangeable partition probability function (pEPPF) induced by our model, which we define in the special case of  $I = 2$ . Consider a sample  $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$  of size  $N = N_1 + N_2$  from model (2.10), together with (2.4)–(2.7) for  $I = 2$  populations; let  $k = k_1 + k_2 + k_0$  the number of unique values in the samples, with  $k_1$  ( $k_2$ ) unique values specific to group 1 (2) and  $k_0$  shared between the groups. Call  $\mathbf{n}_i$  the frequencies of the  $k_i$  unique values in group  $i$  and  $\mathbf{q}_i$  the frequencies of the  $k_0$  shared values in group  $i$ ; this is the same notation as in Camerlenghi et al. (2019a), Section 2.2. The pEPPF is defined as

$$\begin{aligned} \Pi_k^N(\mathbf{n}_1, \mathbf{n}_2, \mathbf{q}_1, \mathbf{q}_2 | \mathbf{c} = (\ell, m)) \\ = \int_{\Theta^k} \mathbb{E} \left[ \prod_{j=1}^{k_1} F_\ell^{n_{1j}}(d\theta_{1j}^*) \prod_{j=1}^{k_2} F_m^{n_{2j}}(d\theta_{2j}^*) \prod_{j=1}^{k_0} F_\ell^{q_{1j}}(d\tau_j) F_m^{q_{2j}}(d\tau_j) \right] \end{aligned}$$

**Proposition 3.2.** *Let  $\kappa$  in (2.6) be equal to 1, let  $\pi_1 = P(c_1 = c_2)$ , then the pEPPF  $\Pi_k^{(N)}(\mathbf{n}_1, \mathbf{n}_2, \mathbf{q}_1, \mathbf{q}_2)$  can be expressed as:*

$$\begin{aligned} \Pi_k^{(N)}(\mathbf{n}_1, \mathbf{n}_2, \mathbf{q}_1, \mathbf{q}_2) &= \pi_1 \Phi_k^{(N)}(\mathbf{n}_1, \mathbf{n}_2, \mathbf{q}_1 + \mathbf{q}_2) \\ &\quad + (1 - \pi_1) \Phi_{k_0+k_1}^{(N_1)}(\mathbf{n}_1, \mathbf{q}_1) \Phi_{k_0+k_1}^{(N_2)}(\mathbf{n}_2, \mathbf{q}_2) I(k_0 = 0), \end{aligned} \quad (3.1)$$

where

$$\Phi_k^{(N)}(\mathbf{n}_1, \mathbf{n}_2, \mathbf{q}_1 + \mathbf{q}_2) = \frac{\alpha^{k_1+k_2+k_0} \Gamma(\alpha)}{\Gamma(\alpha + N)} \prod_{j=1}^{k_1} \Gamma(n_{1j}) \prod_{j=1}^{k_2} \Gamma(n_{2j}) \prod_{j=1}^{k_0} \Gamma(q_{1j} + q_{2j})$$

is the EPPF of the fully exchangeable case, and

$$\Phi_{k_0+k_i}^{(N_i)}(\mathbf{n}_i, \mathbf{q}_i) = \frac{\alpha^{k_i+k_0} \Gamma(\alpha)}{\Gamma(\alpha + N_i)} \prod_{j=1}^{k_i} \Gamma(n_{ij}) \prod_{j=1}^{k_0} \Gamma(q_{ij}), \quad i = 1, 2$$

is the marginal EPPF for the individual group  $i$ .

*Proof.* See the Supplementary Materials file, Section 1. □

This result shows that a suitable prior for  $\kappa$  requires assigning zero probability to the event  $\kappa = 1$ . The assumption in (2.8) trivially satisfies this requirement.

Finally, we consider the marginal law of a sequence of vectors  $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_I)$ ,  $\boldsymbol{\theta}_\ell = (\theta_{\ell 1}, \dots, \theta_{\ell N_\ell})$  from model (2.3)–(2.7). Let us first derive the marginal law conditioning on  $\mathbf{c}$ , as the full marginal law will be the mixture of these conditional laws over all the possible values of  $\mathbf{c}$ .

**Proposition 3.3.** *The marginal law of a sequence of vectors  $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_I)$ ,  $\boldsymbol{\theta}_\ell = (\theta_{\ell 1}, \dots, \theta_{\ell N_\ell})$  from model (2.3)–(2.7), conditional to  $\mathbf{c}$  is*

$$\prod_{i=1}^{R(\mathbf{c})} \text{epff}(\mathbf{n}_{r_i}; \alpha) \sum_{\mathbf{h} \in \{0,1\}^L} p(\mathbf{h}) \prod_{\ell=1}^L G_0(d\theta_\ell^*)^{h_\ell} \times \text{epff}(\mathbf{m}_{r_i} | \mathbf{h}; \gamma) \prod_{k=1}^M G_{00}(d\theta_k^{**}). \quad (3.2)$$

Here,  $\{\theta_\ell^*\}_{\ell=1}^L = \{\theta_{11}^*, \dots, \theta_{IH_I}^*\}$  is a sequence representing all the tables in the process, obtained by concatenating the tables in each restaurant. Moreover,  $R(\mathbf{c})$  is the number of unique values in  $\mathbf{c}$ , i.e. the number of non-empty restaurants,  $\mathbf{n}_{r_i}$  is the vector of  $\ell$ -cluster sizes for restaurant  $r_i$ ,  $\mathbf{m}_{r_i}$  is the vector of the cluster sizes of the  $\theta_\ell^*$  such that  $h_\ell = 0$  and  $\theta_k^{**}$  are the unique values among such  $\theta_\ell^*$ , where “epff” denotes the distribution of the partition induced by the table assignment procedure in the food court of Chinese restaurants described in Section 2.3.

*Proof.* See the Supplementary Materials file, Section 1. □

The marginal law of  $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_I)$  is then

$$\mathcal{L}(d\boldsymbol{\theta}_1, \dots, d\boldsymbol{\theta}_I) = \sum_{\mathbf{c}} \mathcal{L}(d\boldsymbol{\theta}_1, \dots, d\boldsymbol{\theta}_I | \mathbf{c}) \pi(\mathbf{c}),$$

where  $\mathcal{L}(d\boldsymbol{\theta}_1, \dots, d\boldsymbol{\theta}_I | \mathbf{c})$  is given in (3.2).

Observe that in Proposition 3.2 we denoted by  $\Phi$  the EPPF, while in (3.2) we use notation “*eppf*”. This is to remark that these objects are inherently different:  $\Phi$  is the EPPF of the partition of *unique* values in the sample, while *eppf* here is the EPPF of the tables, or  $\ell$ -clusters, induced by the table assignment procedure described in Section 2.3. Hence, from a sample  $\boldsymbol{\theta}$  one can recover  $\mathbf{n}_1, \mathbf{n}_2, \mathbf{q}_1, \mathbf{q}_2$  in (3.1) but not  $\mathbf{n}_{r_i}$  in (3.2).

### 3.3 Some Results on the Bayes Factor for Testing Homogeneity

We consider now testing for homogeneity within the proposed partial exchangeability framework. As a byproduct of the assumed model, the corresponding Bayes factor is immediately available. For example, if one wanted to test whether populations  $i$  and  $j$  were homogeneous, it would suffice to compute the Bayes factor for the test

$$H_0 : c_i = c_j \quad \text{vs.} \quad H_1 : c_i \neq c_j \quad (3.3)$$

which can be straightforwardly estimated from the output of the posterior simulation algorithm that will be presented later on. Note that these “pairwise” homogeneity tests are not the only object of interest that we can tackle within our framework. Indeed it is possible to test any possible combination of  $\mathbf{c}$  against an alternative.

These tests admit an equivalent representation in terms of a model selection problem; for example in the case of  $I = 2$  populations, we can rewrite (3.3), for  $i = 1$  and  $j = 2$ , as a model selection test for  $M_1$  against  $M_2$ , where

$$M_1 : y_{11}, \dots, y_{1N_1}, y_{21}, \dots, y_{2N_2} | F_1 \stackrel{\text{iid}}{\sim} \int_{\Theta} k(\cdot | \theta) F(d\theta), \quad F_1 \sim \text{semi-HDP}(\alpha, \gamma, \kappa, G_0, G_{00})$$

and

$$M_2 : y_{i1}, \dots, y_{iN_i}, | F_i \stackrel{\text{iid}}{\sim} \int_{\Theta} k(\cdot | \theta) F_i(d\theta), \quad i = 1, 2, \quad F_1, F_2 \sim \text{semi-HDP}(\alpha, \gamma, \kappa, G_0, G_{00}).$$

In this case

$$BF_{12} := BF_{12}(y_{11}, \dots, y_{1N_1}, y_{21}, \dots, y_{2N_2}) = \frac{m_{M_1}(y_{11}, \dots, y_{1N_1}, y_{21}, \dots, y_{2N_2})}{m_{M_2}(y_{11}, \dots, y_{1N_1}, y_{21}, \dots, y_{2N_2})},$$

where  $m_{M_i}$  denotes the marginal law of the data under model  $M_i$ ,  $i = 1, 2$ , defined above. Asymptotic properties of Bayes factors have been discussed by several authors. We refer to Walker et al. (2004), Ghosal et al. (2008) for a more detailed discussion and to Chib and Kuffner (2016) for a recent survey on the topic. Chatterjee et al. (2020) is

a recent and solid contribution to the almost sure convergence of Bayes factor in the general set-up that includes dependent data, i.e. beyond the usual i.i.d. context.

In words, our approach can be described as follows. When the data are assumed to be exchangeable, we assume that both samples are generated i.i.d from a distribution  $P_0$  with density  $p_0$ . If the data are instead assumed to be partially exchangeable, then we consider the first population to be generated i.i.d from a certain  $P_0$  with density  $p_0$ , while the second one is generated from  $Q_0$  with density  $q_0$ , with  $P_0 \neq Q_0$  and independence holds across populations. The Bayes factor for comparing  $M_1$  against  $M_2$  is thus consistent if:

(i)  $BF_{12} \rightarrow +\infty$   $P_0^\infty$ -a.s. when  $N_1, N_2 \rightarrow +\infty$  if the groups are truly homogeneous,

and

(ii)  $BF_{12} \rightarrow 0$   $(P_0 \otimes Q_0)^\infty$ -a.s. when  $N_1, N_2 \rightarrow +\infty$  if the groups are not homogeneous.

The two scenarios must be checked separately. In the latter case, consistency of the Bayes factor can be proved by arguing that only model  $M_2$  satisfies the so-called Kullback-Leibler property, so that consistency is ensured by the theory in Walker et al. (2004). We summarize this result in the following proposition.

**Proposition 3.4.** *Assume that  $y_{11}, \dots, y_{1N_1} \stackrel{iid}{\sim} P_0$ ,  $y_{21}, \dots, y_{2N_2} \stackrel{iid}{\sim} Q_0$ ,  $P_0 \neq Q_0$ , and that  $\{y_{1i}\}$  and  $\{y_{2j}\}$  are independent. Assume that  $P_0$  and  $Q_0$  are absolutely continuous measures with probability density functions  $p_0$  and  $q_0$  respectively. Then, under conditions B1-B9 in Wu and Ghosal (2008),  $BF_{12} \rightarrow 0$  as  $N_1, N_2 \rightarrow +\infty$ .*

*Proof.* See the Supplementary Materials file, Section 1. □

Observe that, out of the nine conditions B1-B9, we have that B1 – B3, B7 and B9 involve regularity conditions of the kernel  $k(\cdot|\theta)$ . These are satisfied if the kernel is, for example, univariate Gaussian with parameters  $\theta = (\mu, \sigma^2)$ . Conditions B4 – B6 involve regularity of the true data generating density, which are usually satisfied in practice. Condition B8 requires that the mixing measure has full weak support, already proved in Proposition 3.1.

On the other hand, when  $p_0 = q_0$ , consistency of the Bayes factor would require  $BF_{12} \rightarrow +\infty$ . This is a result we have not been able to prove so far. The Supplementary Materials file, Section 2 discusses the relevant issues arising when trying to prove the consistency in this setting; we just report here that the key missing condition is an upper bound of the prior mass of  $M_2$ . The lack of such bounds for general nonparametric models is well known in the literature, and not specific to our case, as it is shared, for instance, by Bhattacharya and Dunson (2012) and Tokdar and Martin (2019). In both cases, the authors were able to prove the consistency under the alternative hypothesis but not under the null. For a discussion on the “necessity” of these bounds in nonparametric models, see Tokdar and Martin (2019).

In light of the previous consistency result for the non-homogeneous case, our recommendation to carry out the homogeneity test is to decide in favor of  $H_0$  whenever the posterior of  $c_i, c_j$  does not strongly concentrate on  $c_i \neq c_j$ . As Section 5 shows, in our simulated data experiments this choice consistently identifies the right structure of homogeneity among populations. See also the discussion later in Section 7.

## 4 Posterior Simulation

We illustrate a Markov chain Monte Carlo (MCMC) sampler based on the restaurant representation derived in Section 2.3. The random measures  $\{F_i\}_i$  and  $\tilde{G}$  are marginalized out for all the updates except for the case of  $\mathbf{c}$ , for which we use a result from Pitman (1996) to sample from the full conditional of each  $F_i$ , truncating the infinite stick-breaking sum adaptively; see below. We refer to this algorithm as *marginal*. We also note that, by a prior truncation of all the stick-breaking infinite sums to a fixed number of atoms, we can derive a blocked Gibbs sampler as in Ishwaran and James (2001). However, in our applications the blocked Gibbs sampler was significantly slower both in reaching convergence to the stationary distribution and to complete one single iteration of the MCMC update. Hence, we will describe and use only the marginal algorithm.

We follow the notation introduced in Section 2.3. The state of our MCMC sampler consists of the restaurant tables  $\{\theta_{rh}^*\}$ , the tables in the common area  $\{\tau_h\}$ , a set of binary variables  $\{h_{rj}\}$ , indicating if each table is “located” in the restaurant-specific or in the common area, a set of discrete shared table allocation variables  $t_{r\ell}$ , one for each  $\theta_{r\ell}^*$  such that  $\theta_{r\ell}^* = \tau_k$  iff  $t_{r\ell} = k$  and  $h_{r\ell} = 0$ , the categorical variables  $c_i$ , indicating the restaurant for each population,  $\kappa \in (0, 1)$ , and the table allocation variable  $s_{ij}$ : for each observation such that  $\theta_{ij} = \theta_{rh}^*$  iff  $c_i = r$  and  $s_{ij} = h$ . We also denote by  $H_0$  and  $H_r$  the number of tables occupied in the shared area and in restaurant  $r$  respectively,  $m_{rk}$  indicates the number of customers in the common area entered from restaurant  $r$  seating at table  $k$ .

We use the dot notation for marginal counts, for example  $n_{r\cdot}$  indicates all the customers entered in restaurant  $r$ . We summarize the Gibbs sampling scheme next.

- Sample the cluster allocation variables using the Chinese Restaurant Process,

$$p(s_{ij} = s \mid c_i = r, rest) \propto \begin{cases} n_{r\ell}^{-ij} k(y_{ij} \mid \theta_{rh}^*) & \text{if } s \text{ previously used,} \\ \alpha p(y_{ij} \mid \mathbf{s}^{-ij}, rest) & \text{if } s = s^{new}, \end{cases} \quad (4.1)$$

where

$$p(y_{ij} \mid \mathbf{s}^{-ij}, rest) = \kappa \int k(y_{ij} \mid \theta) G_0(d\theta) + (1 - \kappa) \left( \sum_{k=1}^{H_0} \frac{m_{\cdot k}^{-ij}}{m_{\cdot\cdot}^{-ij} + \gamma} k(y_{ij} \mid \tau_k) + \frac{\gamma}{m_{\cdot\cdot}^{-ij} + \gamma} \int k(y_{ij} \mid \theta) G_{00}(d\theta) \right), \quad (4.2)$$



and where the notation  $x^{-ij}$  means that observation  $y_{ij}$  is removed from the calculations involving the variable  $x$ .

If  $s = s^{new}$ , a new table is created. The associated value  $\theta_{rs^{new}}^*$  is sampled from  $G_0$  with probability  $\kappa$  or from  $\tilde{G}$  with probability  $1 - \kappa$ , as described in Section 2.3. The corresponding latent variables  $h_{rs^{new}}$  and  $t_{rs^{new}}$  are set accordingly. When sampling from (2.12) a new table in the shared area might be created. In that case,  $t_{rs^{new}}$  is set to  $H_0 + 1$ .

- Sample the table allocation variables  $t_{r\ell}$  as in the HDP:

$$p(t_{r\ell} = k | rest) \propto \begin{cases} m_{\cdot k}^{-r\ell} \prod_{(i,j):c_i=r,s_{i,j}=\ell} k(y_{ij} | \tau_k) & \text{if } k \text{ previously used,} \\ \gamma \int \prod_{(i,j):c_i=r,s_{i,j}=\ell} k(y_{ij} | \tau) G_{00}(d\tau) & \text{if } k = k^{new}, \end{cases} \quad (4.3)$$

where the notation  $x^{-r\ell}$  means that table  $\theta_{r\ell}^*$ , including all the associated observations, is entirely removed from the calculations involving variable  $x$ . If  $k = k^{new}$  a new table is created in the shared area, the allocation variables  $s_{ij}$  are left unchanged.

- Sample the cluster values from

$$\mathcal{L}(\theta_{r\ell}^* | h_{r\ell} = 1, rest) \propto G_0(\theta_{r\ell}^*) \prod_{(i,j):c_i=r,s_{i,j}=\ell} k(y_{ij} | \theta_{r\ell}^*)$$

and

$$\mathcal{L}(\tau_k | rest) \propto G_{00}(\tau_k) \prod_{(i,j) \in (*)} k(y_{ij} | \tau_k),$$

where the product  $(*)$  is over all the index couples such that  $c_i = r, s_{ij} = \ell, h_{r\ell} = 0$  and  $\theta_{r\ell}^* = \tau_k$ . Observe that, when  $h_{r\ell} = 0$ , it means that  $\theta_{r\ell}^* = \tau_k$  for some  $k$ . Hence, in this case,  $\theta_{r\ell}^*$  is purely symbolic and we do not need to sample a value for it.

- Sample each  $h_{r\ell}$  independently from

$$p(h_{r\ell} = 1 | rest) \propto \kappa G_0(\theta_{r\ell}^*),$$

$$p(h_{r\ell} = 0 | rest) \propto (1 - \kappa) \left( \sum_{k=1}^{H_0} \frac{m_{\cdot k}^{-r\ell}}{m_{\cdot\cdot}^{-r\ell} + \gamma} \delta_{\tau_k}(\theta_{r\ell}^*) + \frac{\gamma}{m_{\cdot\cdot}^{-r\ell} + \gamma} G_{00}(\theta_{r\ell}^*) \right),$$

where, as in (4.3), the notation  $x^{-r\ell}$  means that table  $\theta_{r\ell}^*$ , including all its associated observations, is removed from the calculations involving variable  $x$ . Observe that, while in the update of the cluster values all the  $\theta_{r\ell}^*$  referring to the same  $\tau_k$  were updated at once, here we move the tables one by one.

- Sample  $\kappa$  from  $\mathcal{L}(\kappa | rest) \sim \text{Beta}(a_\kappa + \sum_{i,j} h_{ij}, b_\kappa + \sum_{i,j} (1 - h_{ij}))$ .

- Sample  $\omega$  from

$$\omega \mid \text{rest} \sim \text{Dirichlet}\left(\eta_1 + \sum_{i=1}^I \mathbb{I}[c_i = 1], \dots, \eta_I + \sum_{i=1}^I \mathbb{I}[c_i = I]\right),$$

where  $\mathbb{I}[\cdot]$  denotes the indicator function.

- Sample each  $c_i$  in  $\mathbf{c} = (c_1, \dots, c_I)$  independently from

$$P(c_i = r \mid F_1, \dots, F_I, \omega, \mathbf{y}_i) \propto \omega_r \prod_{j=1}^{N_i} \int k(y_{ij} \mid \theta) F_r(d\theta). \quad (4.4)$$

If the new value of  $c_i$  differs from the previous one, then following (4.1), all the observations  $y_{i1}, \dots, y_{iN_i}$  are reallocated to the new restaurant.

Note that the update in (4.4) involves the previously marginalized random probability measures  $F_1, \dots, F_I$ . Thus, before performing this update, we need to draw the  $F_i$ 's from their corresponding full conditional distributions. It follows from Corollary 20 in Pitman (1996) that the conditional distribution of  $F_r$  given  $\mathbf{c}$ ,  $\mathbf{n}_r$ ,  $\boldsymbol{\theta}_r^*$ ,  $\kappa$ , and  $\tilde{G}$  coincides with the distribution of  $\pi_{r0} F_r' + \sum_{h=1}^{H_r} \pi_{rh} \delta_{\theta_{rh}^*}$ , where  $(\pi_{r0}, \pi_{r1}, \dots, \pi_{rH_r}) \sim \text{Dirichlet}(\alpha, n_{r1}, \dots, n_{rH_r})$  and  $F_r' \mid \tilde{G} \sim \mathcal{D}_{\alpha \tilde{P}}$ . This result was employed in Taddy et al. (2012) to quantify posterior uncertainty of functionals of a Dirichlet process, and also in Canale et al. (2019) to derive an alternative MCMC scheme for mixture models. It follows from the usual stick breaking representation that  $F_r' = \sum_{h=1}^{\infty} w'_{rh} \delta_{\theta'_{rh}}$  with  $\{w'_{rh}\}_h \sim SB(\alpha)$  and  $\theta'_{rh} \mid \kappa, \tilde{G} \stackrel{\text{iid}}{\sim} \kappa G_0 + (1 - \kappa) \tilde{G}$ . Similarly, the conditional distribution of  $\tilde{G}$  given  $\boldsymbol{\tau}$  and  $\mathbf{m}$  coincides with the distribution of  $v_0 \tilde{G}' + \sum_{k=1}^{H_0} v_k \delta_{\tau_k}$ , where  $(v_0, v_1, \dots, v_{H_0}) \sim \text{Dirichlet}(\gamma, m_{\cdot 1}, \dots, m_{\cdot H_0})$  and  $\tilde{G}' \sim \mathcal{D}_{\gamma G_{00}}$ .

In practice, we draw each  $F_r'$  by truncating the infinite sum. Note that we do not need to set a priori the truncation level. Instead, we can specify an upper bound for the error introduced by the truncation and set the level adaptively. In fact, as a straightforward consequence of Theorem 1 in Ishwaran and James (2002) we have that the total variation distance between  $F_r'$  and its approximation with  $M$  atoms, say  $F_r^{M'}$ , is bounded by  $\varepsilon_M = 1 - \sum_{h=1}^M w'_{rh}$  (see also Theorem 2 in Lijoi et al., 2020b). The error induced on  $F_r$  is then bounded by  $\pi_{r0} \varepsilon_M$ . Note that simulation of the atoms  $\theta'_{rh}$  involves the discrete measure  $\tilde{G}'$ . However, we only need to draw a finite number of samples from it, and not its full trajectory, so that no truncation is necessary for  $\tilde{G}'$ . For ease of bookkeeping, we employ retrospective sampling (Papaspiliopoulos and Roberts, 2008) to simulate the atoms. Alternatively, the classical CRP representation can be used. In our experiments, because  $\sum_{h=1}^{H_r} n_{rh} \gg \alpha$  we have  $\pi_{r0} \ll \sum_{h=1}^{H_r} \pi_{rh} \approx 1$ . Thus, choosing a truncation level  $M = 10$  always produces an error on  $F_i$  lower than  $10^{-4}$  (henceforth fixed as the truncation error threshold). Furthermore, we are often not even required to draw samples from  $\tilde{G}'$ .

Of the aforementioned steps, the bottleneck is the update of  $\mathbf{c}$  because for each  $c_i$  we are required to evaluate the densities of  $N_i$  points in  $I$  mixtures. If  $N_i = N$  for all  $i$ ,

the computational cost of this step is  $O(NI^2)$ , which can be extremely demanding for large values of  $I$ . We can mitigate the computational burden by replacing this Gibbs step with a Metropolis-within-Gibbs step, in the same spirit of the *Metropolised Carlin and Chib* algorithm proposed in Dellaportas et al. (2002). At each step we propose a move from  $c_i^{(\ell)} = r$  to  $c_i^{(\ell+1)} = m$  with a certain probability  $p_i(m|r)$ . The transition is then accepted with the usual Metropolis-Hastings rule, i.e. the new update becomes:

- Propose a candidate  $m$  by sampling  $p_i(m|r)$
- Accept the move with probability  $q$ , where

$$q = \min \left[ 1, \frac{P(c_i = m) \prod_{j=1}^{N_i} \int k(y_{ij}|\theta) F_m(d\theta) p_i(r|m)}{P(c_i = r) \prod_{j=1}^{N_i} \int k(y_{ij}|\theta) F_r(d\theta) p_i(m|r)} \right].$$

We call this alternative sampling scheme the Metropolised sampler. The key point is that if evaluating the proposal  $p_i(\cdot|r)$  has a negligible cost, the computational cost of this step will be  $O(2NI)$  as for each data point we need to evaluate only two mixtures: the one corresponding to the current state  $F_r$  and the one corresponding to the proposed state  $F_m$ . Of course, the efficiency and mixing of the Markov chain will depend on a suitable choice of the transition probabilities  $p_i(\cdot|r)$ ; some possible alternatives are discussed in Section 5.

When, at the end of an iteration, a cluster is left unallocated (or empty), the probability of assigning an observation to that cluster will be zero for all subsequent steps. As in standard literature, we employ a relabeling step that gets rid of all the unused clusters. However, this relabeling step is slightly more complicated since there are two different types of clusters: one arising from  $G_0$  and ones arising from  $\tilde{G}$ . Details of the relabeling procedure are discussed in the Supplementary Materials file, Section 3.

## 4.1 Use of Pseudopriors

The above mentioned sampling scheme presents a major issue that could severely impact the mixing. Consider as an example the case when  $I = 2$ ; if, at iteration  $k$ , the state jumps to  $c_1 = c_2 = 1$ , then all the tables of the second restaurant would be erased from the state, because no observation is assigned to them anymore. Switching back to  $c_1 \neq c_2$  would then require that the approximation of  $F_2$  sampled from its prior distribution gives sufficiently high likelihood to either  $\mathbf{y}_1$  or  $\mathbf{y}_2$ , an extremely unlikely event in practice.

To overcome this issue, we make use of pseudopriors as in Carlin and Chib (1995), that is, whenever a random measure  $F_r$  in  $(F_1, \dots, F_I)$  is not associated with any group, we sample the part of the state corresponding to that measure (the atoms  $\{\theta_{r\ell}^*\}$  and number of customers  $\{n_{r\ell}\}$  in each restaurant) from its pseudoprior. From the computational point of view, this is accomplished by running first a preliminary MCMC simulation where the  $c_i$ 's are fixed as  $c_i = i$ , and collecting the samples. Then, in the actual MCMC simulation, whenever restaurant  $r$  is empty we change the state by choosing at random one of the previous samples obtained with fixed  $c_i$ 's. Note that this use

of pseudopriors does not alter the stationary distribution of the MCMC chain. Furthermore, the way pseudopriors are collected and sampled from is completely arbitrary, and our proposed solution works well in practice. Other valid options include approximations based on preliminary chain runs, as discussed in Carlin and Chib (1995).

Section 5 below contains extensive simulation studies that show that the proposed model can be used to efficiently estimate densities for each population. We also tried the case of a large number of populations, e.g.  $I = 100$  without any significant loss of performance.

## 5 Simulation Study

In this section we investigate the ability of our model to estimate dependent random densities. We fix the kernel  $k(\cdot|\theta)$  in (2.2) to be the univariate Gaussian density with parameter  $\theta = (\mu, \sigma^2)$  (mean and variance, respectively). Both base measures  $G_0$  and  $G_{00}$  are chosen to be

$$\mathcal{N}(\mu|0, 10\sigma^2) \times \text{inv-gamma}(\sigma^2|1, 1),$$

and unless otherwise stated, with hyperparameters  $\alpha, \gamma$  fixed to 1,  $a_\kappa = b_\kappa = 2$ , and  $\boldsymbol{\eta} = (1/I, \dots, 1/I)$ . Chains were run for 100,000 iterations after discarding the first 10,000 iterations as burn-in, keeping one every ten iterations, resulting in a final sample size of 10,000 MCMC draws.

### 5.1 Two Populations

We first focus on the special case of  $I = 2$  populations. Consider generating data as follows

$$\begin{aligned} y_{1j} &\stackrel{\text{iid}}{\sim} w_1 \mathcal{N}(\mu_1, \sigma_1) + (1 - w_1) \mathcal{N}(\mu_2, \sigma_2) & j = 1, \dots, N_1, \\ y_{2j} &\stackrel{\text{iid}}{\sim} w_2 \mathcal{N}(\mu_3, \sigma_3) + (1 - w_2) \mathcal{N}(\mu_4, \sigma_4) & j = 1, \dots, N_2, \end{aligned} \quad (5.1)$$

that is each population is a mixture of two normal components. This is the same example considered in Camerlenghi et al. (2019a). Table 1 summarizes the parameters used to generate the data. Note that these three scenarios cover either the full exchangeability case across both populations (Scenario I), as well as the partial exchangeability between the two populations (scenarios II and III). For each case, we simulated  $N_1 = N_2 = 100$  observations for each group (independently).

Table 2 reports the posterior probabilities of the two population being identified as equal for the three scenarios. We can see that our model recovers the ground truth.

	$(\mu_1, \sigma_1)$	$(\mu_2, \sigma_2)$	$(\mu_3, \sigma_3)$	$(\mu_4, \sigma_4)$	$w_1$	$w_2$
Scenario I	(0.0, 1.0)	(5.0, 1.0)	(0.0, 1.0)	(5.0, 1.0)	0.5	0.5
Scenario II	(5.0, 0.6)	(10.0, 0.6)	(5.0, 0.6)	(0.0, 0.6)	0.9	0.1
Scenario III	(0.0, 1.0)	(5.0, 1.0)	(0.0, 1.0)	(5.0, 1.0)	0.8	0.2

Table 1: Parameters of the simulated datasets.

	$P(c_1 = c_2   data)$	$BF_{01}$
Scenario I	0.99	98.9
Scenario II	0.0	0.0
Scenario III	0.0	0.0

Table 2: Posterior inference.




Figure 2: Posterior distribution of the number of shared unique values and unique values specific to first and second group in Scenario II.

Moreover Figure 3 in the Supplementary Materials file shows the density estimates, i.e. the posterior mean of the density evaluated on a fixed grid of points, together with pointwise 95% posterior credible intervals at each point  $x$  in the grid, obtained by our MCMC for scenarios I and III. Here, densities are estimated from the corresponding posterior mean evaluated on a fixed grid of points, while credible intervals are obtained by approximating the  $F_i$ 's as discussed in Section 4. We can see that in both the cases, locations and scales of the populations are recovered perfectly, while it seems that the weights of the mixture components are slightly more precise in Scenario I than in Scenario III.

Comparing the Bayes Factors shown in Table 2 with the ones in Camerlenghi et al. (2019a) (5.86, 0.0 and 0.54 for the three scenarios, respectively), we see that both models are able to correctly assess homogeneity. However, the Bayes Factors obtained under our model tend to assume more extreme than those from Camerlenghi et al. (2019a). Figure 2 shows the posterior distribution of the number of shared and private unique values (reconstructed from the cluster allocation variables  $s_{ij}$  and the table allocation variables  $t_{r\ell}$ ) in Scenario II, when either  $\kappa \sim Beta(2, 2)$  or  $\kappa = 1$ . Also in the latter case  $P(c_1 = c_2 | data) = 0$ , but the shared component between groups one and two is not recovered, due to the degeneracy issue described in Proposition 3.2.

As the central point of our model is to allow for different random measures to share at least one atom, we test more in detail this scenario. To do so, we simulate 50 different datasets from (5.1), by selecting  $\mu_1, \mu_2, \mu_4 \stackrel{iid}{\sim} \mathcal{N}(0, 10)$  and  $\sigma_1^2, \sigma_2^2, \sigma_4^2 \stackrel{iid}{\sim} inv - gamma(2, 2)$ ,  $w_1 \sim Beta(1, 1)$  and setting  $\mu_3 = \mu_1, \sigma_3^2 = \sigma_1^2, w_2 = w_1$ . In this way we create 50 independent scenarios where the two population share exactly one component and give the same weight to this component. Figure 4 in the Supplementary Materials file reports the scatter plot of the estimated posterior probabilities of  $c_1 = c_2$  obtained from the MCMC samples. It is clear that our model recovers the right scenario

most of the times. Out of 50 examples, only in four of them  $P(c_1 = c_2 | data)$  is greater than 0.5, by a visual analysis we see from the plot of the true densities that in those cases the two populations were really similar.

## 5.2 More than Two Populations

We extend now the simulation study to scenarios with more than two populations. We consider three simulated datasets with four populations each and different clustering structures at the population level. In particular, we use the same scenarios as in Gutiérrez et al. (2019), and simulate  $N_i = 100$  points for each population  $i = 1, 2, 3, 4$  as follows

- Scenario IV

$$y_{1j}, y_{2k}, y_{3\ell} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1) \quad y_{4n} \stackrel{\text{iid}}{\sim} SN(0, 1, 1) \quad j, k, \ell, n = 1, \dots, 100.$$

- Scenario V

$$y_{1j}, y_{4n} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1) \quad y_{2k} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 2.25) \quad y_{3\ell} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 0.25) \quad j, k, \ell, n = 1, \dots, 100.$$

- Scenario VI

$$\begin{aligned} y_{1j}, y_{2k} &\stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(0, 1) + 0.5\mathcal{N}(5, 1) \quad j, k = 1, \dots, 100, \\ y_{3\ell} &\stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(0, 1) + 0.5\mathcal{N}(-5, 1) \quad \ell = 1, \dots, 100, \\ y_{4n} &\stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(-5, 1) + 0.5\mathcal{N}(5, 1) \quad n = 1, \dots, 100. \end{aligned}$$

Hence, the true clusters of the label set of the populations,  $\{1, 2, 3, 4\}$ , are:  $\boldsymbol{\rho}_4^{true} = \{\{1, 2, 3\}, \{4\}\}$ ,  $\boldsymbol{\rho}_5^{true} = \{\{1, 4\}, \{2\}, \{3\}\}$  and  $\boldsymbol{\rho}_6^{true} = \{\{1, 2\}, \{3\}, \{4\}\}$  for the three scenarios under investigation respectively. By  $SN(\xi, \omega, \alpha)$  in Scenario IV we mean the skew-normal distribution with location  $\xi$ , scale  $\omega$  and shape  $\alpha$ ; in this case, the mean of the distribution is equal to

$$\xi + \omega \frac{\alpha}{1 + \alpha^2} \sqrt{\frac{2}{\pi}}.$$

Note that we focus on a different problem than what Gutiérrez et al. (2019) discussed, as they considered testing for multiple treatments against a control. In particular they were concerned about testing the hypothesis of equality in distribution between data coming from different treatments  $\mathbf{y}_j$  ( $j = 2, 3, 4$  in these scenarios), and data coming from a control group  $\mathbf{y}_1$ . Instead our goal is to cluster these populations based on their distributions.

Observe how the prior chosen for  $\mathbf{c}$  does not translate directly into a distribution on the partition  $\boldsymbol{\rho}$ , as it is affected by the so called label switching. Thus, in order to summarize our inference, we post-process our chains and transform the samples  $\mathbf{c}^{(1)}, \dots, \mathbf{c}^{(M)}$  from  $\mathbf{c}$  to samples  $\boldsymbol{\rho}^{(1)}, \dots, \boldsymbol{\rho}^{(M)}$  from  $\boldsymbol{\rho}$ . For example we have that  $\mathbf{c}^{(i)} = (1, 1, 1, 3)$  and  $\mathbf{c}^{(j)} = (2, 2, 2, 4)$  both get transformed into  $\boldsymbol{\rho}^{(i)} = \boldsymbol{\rho}^{(j)} = \{\{1, 2, 3\}, \{4\}\}$ .




Figure 3: Posterior probability of  $\rho$  for Scenario IV.




Figure 4: Density estimates and pointwise 95% posterior credible intervals for Scenario IV.

The posterior probabilities of the true clusters  $P(\rho_i = \rho_i^{true} | data)$  are estimated using the transformed (as described above) MCMC samples and equal 0.75, 0.95 and 0.99 for the three scenarios respectively. Figure 3 shows the posterior distribution of  $\rho$ , and Figure 4 reports the density estimation of each group, for Scenario IV. Observe how the posterior mode is in  $\rho_4^{true}$  but significant mass is given also to the case  $\{\{1\}, \{2, 3\}, \{4\}\}$ . We believe that this behavior is mainly due to our use of pseudopriors, as it makes the transition between these three states fairly smooth. On the other hand, in Scenario V, where the posterior mass on the true cluster is close to 1, it is clear that such transitions happen very rarely, as the posterior distribution, not shown here, is completely concentrated on  $\rho_5^{true}$ . Our insight is that the pseudopriors make a transition between two states, say  $\mathbf{c}^{(j)} = (1, 1, 3, 4)$  and  $\mathbf{c}^{(j+1)} = (1, 2, 3, 4)$  (or viceversa), more likely when the mixing distributions of population one and two are the same.

We compared the performance of the Metropolised algorithm against the full Gibbs move for the update of  $\mathbf{c}$ , computing the effective sample size (ESS) of the number of population level clusters (i.e. the number of unique values in  $\mathbf{c}$ ) over CPU time. We consider two choices for the proposal distribution  $p_i(r|m)$ , namely, the discrete uniform over  $\{1, \dots, I\}$  and another discrete alternative, with weights given by

$$p_i(r|m) \propto 1 + (1 + d^2(F_r, F_m))^{-1}, \quad (5.2)$$

where  $d^2(F_r, F_m)$  is the squared  $L^2$  distance between the Gaussian mixture represented by  $F_r$  and that represented by  $F_m$ , which are sampled as discussed in Section 4. This distance is available in closed form and the formula is reported in Section 4 of the Supplementary Materials file.

Results for data as in Scenario IV show that the best efficiency is obtained using the full Gibbs update, with an ESS per second of 57.1. The Metropolised sampler with proposal as in (5.2) comes second, yielding an ESS per second of 34.1 while the Metropolised sampler with uniform proposal is the worst performer with an ESS per second of 12.8. Hence, even when the number of groups is not enormous, the good performance of the Metropolised sampler is clear. Preliminary analysis showed how the Metropolised sampler outperforms the full Gibbs one as the number of groups increases.

Finally, we test how our algorithm performs when the number of populations increases significantly. We do so by generating 100 populations in Scenario VII as follows:

$$\begin{aligned} y_{ij} &\stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(-5, 1) + 0.5\mathcal{N}(5, 1) & i = 1, \dots, 20, \\ y_{ij} &\stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(-5, 1) + 0.5\mathcal{N}(0, 1) & i = 21, \dots, 40, \\ y_{ij} &\stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(0, 1) + 0.5\mathcal{N}(5, 0.1) & i = 41, \dots, 60, \\ y_{ij} &\stackrel{\text{iid}}{\sim} 0.5\mathcal{N}(-10, 1) + 0.5\mathcal{N}(0, 1) & i = 61, \dots, 80, \\ y_{ij} &\stackrel{\text{iid}}{\sim} 0.1\mathcal{N}(-10, 1) + 0.9\mathcal{N}(0, 1) & i = 81, \dots, 100. \end{aligned}$$

Thus, full exchangeability holds within populations  $\{1, \dots, 20\}$ ,  $\{21, \dots, 40\}$ ,  $\{41, \dots, 60\}$ ,  $\{61, \dots, 80\}$  and  $\{81, \dots, 100\}$  but not between these five groups. For each population  $i$ , 100 datapoints were sampled independently.

To compute posterior inference, we run the Metropolised sampler with proposal (5.2). To get a rough idea of the computational costs associated to this large simulated dataset, we report that running the full Gibbs sampler would have required more than 24 hours on a 32-core machine (having parallelized all the computations which can be safely parallelized), while the Metropolised sampler ran in less than 3 hours on a 6-core laptop.

As a summary of the posterior distribution of the random partition  $\boldsymbol{\rho}_{100}$ , we compute the posterior similarity matrix  $[P(c_i = c_j | data)]_{i,j=1}^I$ . Estimates of these probabilities are straightforward to obtain using the output of the MCMC algorithm. Figure 5 shows the posterior similarity matrix as well as the density estimates of five different populations. It is clear that the clustering structure of the populations is recovered perfectly and that the density estimates are coherent with the true ones.



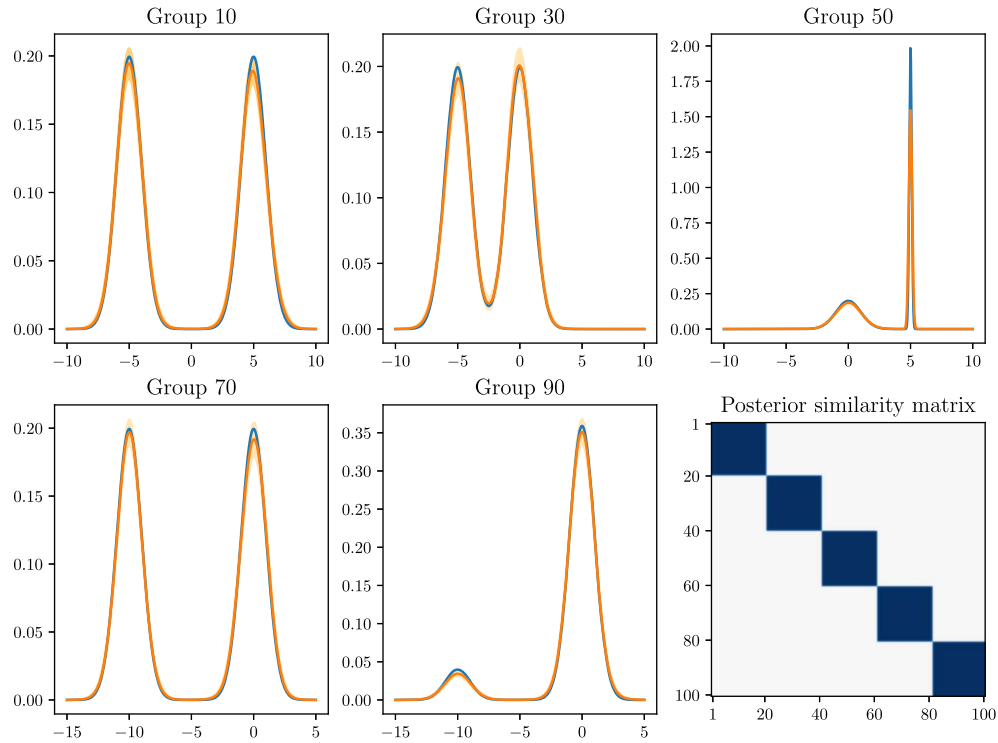


Figure 5: Density estimates (orange line), pointwise 95% posterior credible intervals (orange bands), true data generating densities (blue line) for groups 10, 30, 50, 70 and 90 and posterior similarity matrix (bottom right, white corresponds to 0.0 and dark blue to 1.0) in Scenario VII.

## 6 Chilean Grades Dataset

The School of Mathematics at Pontificia Universidad Católica de Chile teaches many undergraduate courses to students from virtually all fields. When the number of students exceeds a certain maximum pre-established quota, several sections are formed, and courses are taught in parallel. There is a high degree of preparation in such cases, so as to guarantee that courses cover the same material and are coordinated to function as virtual copies of each other. In such cases, only the instructor changes across sections, but all materials related to the courses are the same, including exams, homework, assignments, projects, etc., and there is a shared team of graders that are common to all the parallel sections. According to the rules, every student gets a final grade on a scale from 1.0 to 7.0, using one decimal place, where 4.0 is the minimum passing grade. We consider here the specific case of a version of Calculus II, taught in parallel to three different sections (A, B and C) in a recent semester. Our main goal here is to examine the instructor effectiveness, by comparing the distributions of the final grades obtained




Figure 6: Density estimates and pointwise 95% posterior credible intervals for the three groups (left); posterior distribution of the clusters (right).

by each of the three populations (sections). The sample sizes of these populations are 76, 65 and 50 respectively.

A possible way to model these data could be to employ a truncated normal distribution as the kernel in (2.2). However since our primary interest is to investigate the homogeneity of the underlying distributions and not to perform density estimates, we decided to first add a small amount of zero-mean Gaussian noise, with variance 0.1 to the data (i.e. “jittering”) and then proceeded to standardize the whole dataset, by letting  $y_{ij}^{new} = (y_{ij} - \bar{y})/s_y$ , where  $\bar{y} = (\sum_{ij} y_{ij})/(\sum_i N_i)$  and  $s_y^2 = (\sum_{ij} (y_{ij} - \bar{y})^2)/(\sum_i N_i - 1)$  are the global sample mean and variance, respectively. In the sequel, index  $i = 1, 2, 3$  denotes sections A, B and C, respectively, as described above.

Figure 6 reports density estimates in all groups (i.e. posterior density means evaluated on a fixed grid of points and pointwise 95% posterior credible intervals at each point  $x$  in the grid), as well as the posterior distribution of the random partition  $\rho$ , obtained from the posterior distribution of  $\mathbf{c}$ , getting rid of the label switching in a post-processing step (see also Section 5.2). From Figure 6 we see that the posterior distribution of  $\rho$  gives high probability to the case of the three groups being all different as well as to the case when the first and third groups are homogeneous but different from the second one. This is in accordance with a visual analysis of the observed and estimated densities.

We considered several functionals of the random population distribution  $F_{c_i}$  (see (2.3)) for  $i = 1, 2, 3$ . Recall that, according to notation in (2.1),  $F_{c_i} = G_i$ . First of all, we consider the mean and variance functionals of the random density  $p_i(y) = \int_{\Theta} k(y|\theta)F_{c_i}(d\theta) = \int_{\Theta} k(y|\theta)G_i(d\theta)$ , for each  $i = 1, 2, 3$ . Observe how they are functionals of the random probability  $F_{c_i} = G_i$ . Moreover, since Figure 6 seems to suggest that the three groups differ mainly due to their different asymmetries, we considered two more functionals of  $G_i$ , i.e. two indicators of skewness: Pearson’s moment coefficient

Section	$\mu_i$	$\sigma_i^2$	sk $_i$	$\gamma_{Mi}$	$P_{4i}$
A	-0.264	0.671	120.84	-0.01	0.53
B	0.438	1.428	-64.86	0.292	0.71
C	-0.171	0.943	55.60	-0.01	0.56

Table 3: Posterior means of functionals  $\mu_i, \dots, P_{4i}$  of the population density  $p_i$  for each Section A ( $i = 1$ ), B ( $i = 2$ ) and C ( $i = 3$ ) in the Chilean grades dataset. All the functionals refer to standardized data  $\{y_{ij}^{new} = (y_{ij} - \bar{y})/s_y\}$ .

of skewness sk and the measure of skewness with respect to the mode  $\gamma_M$  proposed by Arnold and Groeneveld (1995). Pearson’s moment coefficient of skewness of the random variable  $T$  is defined as  $sk = \mathbb{E}[(T - \mathbb{E}(T))/\sqrt{\text{Var}(T)}]^3$ , while the measure of skewness with respect to the mode as  $\gamma_M = 1 - 2F_T(M_T)$ , where  $M_T$  is the mode of  $T$  and  $F_T$  denotes its distribution function. The last functional of  $G_i$  we consider is the probability, under the density  $p_i(y) = \int_{\Theta} k(y|\theta)G_i(d\theta)$  of getting a passing grade ( $\geq 4.0$  before normalization), that is

$$P_{4i} = \int_{(4-\bar{y})/s_y}^{+\infty} p_i(y)dy.$$

Table 3 shows the posterior mean of the functionals  $\mu_i, \sigma_i^2$  (mean and variance functionals),  $sk_i, \gamma_{Mi}$  and  $P_{4i}$  of  $p_i$ , for  $i = 1, 2, 3$ . To be clear, the posterior mean of the mean functional  $\mu_1$  is computed as

$$\frac{1}{M} \sum_{\ell=1}^M \mu_1^{(\ell)} = \frac{1}{M} \sum_{\ell=1}^M \mathbb{E}[y | G_1^{(\ell)}] = \frac{1}{M} \sum_{\ell=1}^M \left( \int_{\mathbb{R}} yp_1^{(\ell)}(y)dy \right),$$

where  $M$  is the MCMC sample size, and the superscript  $(\ell)$  attached to a random variable denotes its value at the  $\ell$ -th MCMC iteration.

In agreement with the posterior distribution of the partition  $\rho$ , for all the functionals considered we observed close values for sections A and C, while both differ significantly from the values for section B. In summary, we conclude that section B presents a heavier right tail than sections A and C, hence it is characterized by a higher mean (positive) and also more spread across the range. Section B shows a larger (estimated) value for  $P_4$ , i.e. students in section B are more likely to pass the exam than their colleagues from the other sections. This seems to suggest that a higher concentration of good students (with high grades) was present in Section B, compared to A and C, possibly combined with a higher effectiveness of the instructor in this Section.

We also computed the pairwise  $L^1$  distances between the estimated densities in the populations. If  $\tilde{p}_i$  denotes the estimated density (posterior mean of  $p_i$  evaluated in a grid of points) for each population, we found  $d(\tilde{p}_A, \tilde{p}_B) = 0.56$ ,  $d(\tilde{p}_A, \tilde{p}_C) = 0.15$  and  $d(\tilde{p}_B, \tilde{p}_C) = 0.44$ . This confirms once again that the estimated densities for section A and C are closer than when comparing sections A and B and sections B and C.

To end the analysis, we show in Figure 7 estimated couples of densities  $(p_i, p_\ell)$ ,  $i \neq \ell$ ,  $i, \ell = 1, 2, 3$ , i.e. the posterior mean of  $(p_i, p_\ell)$ , evaluated on a fixed grid in  $\mathbb{R}^2$ .

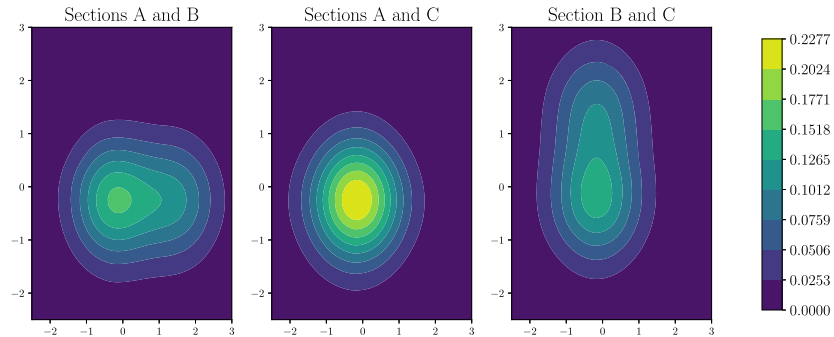


Figure 7: Posterior means of  $(p_i, p_\ell)$ ,  $i \neq \ell$ ,  $i, \ell = A, B, C$ , evaluated on a fixed grid in  $\mathbb{R}^2$  for the Chilean grades dataset.

While sections A and C look independent (central panel in Figure 7), the (posterior) propensity of section B to get higher grades is confirmed in the left and right panels in Figure 7.

## 7 Discussion

Motivated by the traditional problem of testing homogeneity across  $I$  different groups or populations, we have presented a model that is able to not only address the problem but also to perform a cluster analysis of the groups. The model is built on a prior for the population distributions that we termed the semi-hierarchical Dirichlet process, and it was shown to have good properties and also to perform well in synthetic and real data examples, also in case of  $I = 100$  groups. One of the driving features of our proposal was to solve the degeneracy limitation of nested constructions that has been pointed out by Camerlenghi et al. (2019a). The crucial aspect of the semi-HDP that solves this problem was described using the metaphor of a food court of Chinese restaurants with common and private dining area. The hierarchical construction introduces a random partition at the population level, which allows for identifying possible clusters of internally homogeneous groups.

Our examples focus on unidimensional data, though extensions to multivariate responses can be straightforwardly accommodated in our framework. However, scaling with respect to data dimension is not a property we claim to have. In fact, this is a situation shared with any type of hierarchical mixture models.

We studied support properties of the semi-HDP and also the posterior asymptotic behavior of the Bayes factor for the homogeneity test when  $I = 2$ , as posed within the proposed hierarchical construction. We showed that the Bayes factor has the appropriate asymptotic behavior under the alternative hypothesis of partial exchangeability, but a final answer under the assumption of truly exchangeable data is still pending. The lack of asymptotic guarantees is not at all specific to our case. In fact, this situation is rather common to all model selection problems when the hypothesis is not well separated and

at least one of the two models under comparison is “truly” nonparametric, as, for instance, in Bhattacharya and Dunson (2012) and Tokdar and Martin (2019). Indeed, as discussed in Tokdar and Martin (2019), it is not even clear if in such cases the need for an upper bound on the prior mass under the more complex model is a natural requirement or rather a technical one. More generally, intuition about BFs (at least in parametric cases) is that they tend to favor the more parsimonious model. In the particular context described in Section 3.3, model  $M_1$  can be regarded as a degenerate case of model  $M_2$ , even though they are “equally complicated”. In this case, the above intuition evaporates, since technically, embedding one model in the other is still one infinite-dimensional model contained in another infinite-dimensional model, and it is probably meaningless to ask which model is “simpler”. Under this scenario exploratory use of discrepancy measures, such as those discussed in Gelman et al. (1996), may offer some guidance.

In the simulation studies presented, our model always recovers the true latent clustering among groups, thus providing empirical evidence in favor of our model to perform homogeneity tests. We provide some practical suggestions when the actual interest is on making this decision. Our insight is that in order to prove asymptotic consistency of the Bayes factor, one should introduce explicit separation between the competing hypotheses. One possible way to accomplish this goal is, for example, by introducing some kind of repulsion among the mixing measures  $F_i$ 's in the model. This point will be focus of further study.

## Supplementary Material

Supplementary material for “The semi-hierarchical Dirichlet Process and its application to clustering homogeneous distributions” (DOI: [10.1214/21-BA1278SUPP](https://doi.org/10.1214/21-BA1278SUPP); .pdf).

## References

- Argiento, R., Cremaschi, A., and Vannucci, M. (2020). “Hierarchical Normalized Completely Random Measures to Cluster Grouped Data.” *Journal of the American Statistical Association*, 115(529): 318–333. MR4078466. doi: <https://doi.org/10.1080/01621459.2019.1594833>. 1189, 1195, 1197
- Arnold, B. C. and Groeneveld, R. A. (1995). “Measuring skewness with respect to the mode.” *The American Statistician*, 49(1): 34–38. MR1341197. doi: <https://doi.org/10.2307/2684808>. 1213
- Bassetti, F., Casarin, R., and Rossini, L. (2019). “Hierarchical species sampling models.” *Bayesian Analysis*, Advance publication. MR4132651. doi: <https://doi.org/10.1214/19-BA1168>. 1189
- Beraha, M. and Guglielmi, A. (2019). “Discussion on ‘Latent nested nonparametric priors’ by Camerlenghi, Dunson, Lijoi, Prünster and Rodríguez.” *Bayesian Analysis*, 14(4): 1326–1332. MR4044854. doi: <https://doi.org/10.1214/19-BA1169>. 1190, 1192, 1198

- Beraha, M., Guglielmi, A., and Quintana, F. A. (2021). “Supplementary material for “The semi-hierarchical Dirichlet Process and its application to clustering homogeneous distributions”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/21-BA1278SUPP>. 1190
- Bhattacharya, A. and Dunson, D. (2012). “Nonparametric Bayes classification and hypothesis testing on manifolds.” *Journal of multivariate analysis*, 111: 1–19. MR2944402. doi: <https://doi.org/10.1016/j.jmva.2012.02.020>. 1189, 1201, 1215
- Camerlenghi, F., Dunson, D. B., Lijoi, A., Prünster, I., and Rodríguez, A. (2019a). “Latent Nested Nonparametric Priors (with Discussion).” *Bayesian Analysis*, 14(4): 1303–1356. MR4044854. doi: <https://doi.org/10.1214/19-BA1169>. 1187, 1189, 1190, 1191, 1192, 1198, 1206, 1207, 1214
- Camerlenghi, F., Lijoi, A., Orbanz, P., and Prünster, I. (2019b). “Distribution theory for hierarchical processes.” *The Annals of Statistics*, 47(1): 67–92. MR3909927. doi: <https://doi.org/10.1214/17-AOS1678>. 1189
- Camerlenghi, F., Lijoi, A., and Prünster, I. (2017). “Bayesian prediction with multiple-samples information.” *Journal of Multivariate Analysis*, 156: 18–28. MR3624682. doi: <https://doi.org/10.1016/j.jmva.2017.01.010>. 1187
- Canale, A., Corradin, R., and Nipoti, B. (2019). “Importance conditional sampling for Pitman-Yor mixtures.” *arXiv preprint arXiv:1906.08147*. 1204
- Carlin, B. P. and Chib, S. (1995). “Bayesian model choice via Markov chain Monte Carlo methods.” *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(3): 473–484. 1205, 1206
- Catalano, M., Lijoi, A., and Prünster, I. (2021). “Measuring dependence in the Wasserstein distance for Bayesian nonparametric models.” *The Annals of Statistics*, forthcoming. doi: <https://doi.org/10.1214/21-AOS2065>. 1190
- Chatterjee, D., Maitra, T., and Bhattacharya, S. (2020). “A short note on almost sure convergence of Bayes factors in the general set-up.” *The American Statistician*, 74(1): 17–20. MR4055352. doi: <https://doi.org/10.1080/00031305.2017.1397548>. 1200
- Chen, Y. and Hanson, T. E. (2014). “Bayesian nonparametric  $k$ -sample tests for censored and uncensored data.” *Computational Statistics and Data Analysis*, 71: 335–346. MR3131974. doi: <https://doi.org/10.1016/j.csda.2012.11.003>. 1189
- Chib, S. and Kuffner, T. A. (2016). “Bayes factor consistency.” *arXiv preprint arXiv:1607.00292*. 1200
- Cifarelli, D. and Regazzini, E. (1978). “Problemi statistici non parametrici in condizioni di scambiabilità parziale e impiego di medie associative.” Technical report, Tech. rep., Quaderni Istituto Matematica Finanziaria dell’Università di Torino. 1188
- de Finetti, B. (1938). “Sur la condition de” equivalence partielle”, Colloque consacré

- la théorie des probabilités, Vol.” VI, *Université de Geneve, Hermann et C. ie, Paris*. MR0000909. 1187
- De Iorio, M., Müller, P., Rosner, G., and MacEachern, S. (2004). “An ANOVA Model for Dependent Random Measures.” *Journal of the American Statistical Association*, 99: 205–215. MR2054299. doi: <https://doi.org/10.1198/016214504000000205>. 1189
- Dellaportas, P., Forster, J. J., and Ntzoufras, I. (2002). “On Bayesian model and variable selection using MCMC.” *Statistics and Computing*, 12(1): 27–36. MR1877577. doi: <https://doi.org/10.1023/A:1013164120801>. 1205
- Diaconis, P. (1988). “Recent progress on de Finetti’s notions of exchangeability.” In J. M. Bernardo, M. H. DeGroot and A. Smith (eds.), *Bayesian statistics 3*, 111–125. Oxford: Oxford University Press. MR1008047. 1187
- Ferguson, T. S. (1973). “A Bayesian analysis of some nonparametric problems.” *The Annals of Statistics*, 1: 209–230. MR0350949. 1197
- Gelman, A., Meng, X.-L., and Stern, H. (1996). “Posterior predictive assessment of model fitness via realized discrepancies.” *Statistica sinica*, 6: 733–760. MR1422404. 1190, 1215
- Ghosal, S., Lember, J., and Van Der Vaart, A. (2008). “Nonparametric Bayesian model selection and averaging.” *Electronic Journal of Statistics*, 2: 63–89. MR2386086. doi: <https://doi.org/10.1214/07-EJS090>. 1200
- Gutiérrez, L., Barrientos, A. F., González, J., and Taylor-Rodríguez, D. (2019). “A Bayesian Nonparametric Multiple Testing Procedure for Comparing Several Treatments Against a Control.” *Bayesian Analysis*, 14(2): 649–675. MR3959876. doi: <https://doi.org/10.1214/18-BA1122>. 1189, 1208
- Hatjispyros, S. J., Merktas, C., Nicolieris, T., and Walker, S. G. (2018). “Dependent mixtures of geometric weights priors.” *Computational Statistics & Data Analysis*, 119: 1–18. MR3729519. doi: <https://doi.org/10.1016/j.csda.2017.09.006>. 1189
- Hatjispyros, S. J., Nicolieris, T., and Walker, S. G. (2011). “Dependent mixtures of Dirichlet processes.” *Computational Statistics & Data Analysis*, 55(6). MR2785111. doi: <https://doi.org/10.1016/j.csda.2010.12.005>. 1189
- Hatjispyros, S. J., Nicolieris, T., and Walker, S. G. (2016). “Random density functions with common atoms and pairwise dependence.” *Computational Statistics & Data Analysis*, 101: 236–249. MR3504848. doi: <https://doi.org/10.1016/j.csda.2016.03.008>. 1189
- Holmes, C. C., Caron, F., Griffin, J. E., Stephens, D. A., et al. (2015). “Two-sample Bayesian nonparametric hypothesis testing.” *Bayesian Analysis*, 10(2): 297–320. MR3420884. doi: <https://doi.org/10.1214/14-BA914>. 1189
- 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. doi: <https://doi.org/10.1198/016214501750332758>. 1202

- Ishwaran, H. and James, L. F. (2002). “Approximate Dirichlet process computing in finite normal mixtures: smoothing and prior information.” *Journal of Computational and Graphical statistics*, 11(3): 508–532. MR1938445. doi: <https://doi.org/10.1198/106186002411>. 1204
- Lijoi, A., Nipoti, B., and Prünster, I. (2014a). “Dependent mixture models: clustering and borrowing information.” *Computational Statistics & Data Analysis*, 71: 417–433. MR3131980. doi: <https://doi.org/10.1016/j.csda.2013.06.015>. 1189
- Lijoi, A., Nipoti, B., Prünster, I., et al. (2014b). “Bayesian inference with dependent normalized completely random measures.” *Bernoulli*, 20(3): 1260–1291. MR3217444. doi: <https://doi.org/10.3150/13-BEJ521>. 1192
- Lijoi, A., Prünster, I., and Rebaudo, G. (2020a). “Flexible clustering via hidden hierarchical Dirichlet priors.” *Collegio Carlo Alberto Notebooks*, (634). 1190
- Lijoi, A., Prünster, I., and Rigon, T. (2020b). “Sampling hierarchies of discrete random structures.” *Statistics and Computing*, 30(6): 1591–1607. MR4156338. doi: <https://doi.org/10.1007/s11222-020-09961-7>. 1204
- Ma, L. and Wong, W. H. (2011). “Coupling optional Pólya trees and the two sample problem.” *Journal of the American Statistical Association*, 106(496): 1553–1565. MR2896856. doi: <https://doi.org/10.1198/jasa.2011.tm10003>. 1189
- MacEachern, S. N. (1999). “Dependent nonparametric processes.” In *ASA proceedings of the section on Bayesian statistical science*, volume 1, 50–55. Alexandria, Virginia. Virginia: American Statistical Association; 1999. 1188
- Müller, P., Quintana, F., and Rosner, G. (2004). “A method for combining inference across related nonparametric Bayesian models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3): 735–749. MR2088779. doi: <https://doi.org/10.1111/j.1467-9868.2004.05564.x>. 1189, 1192
- Papaspiliopoulos, O. and Roberts, G. O. (2008). “Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models.” *Biometrika*, 95(1): 169–186. MR2409721. doi: <https://doi.org/10.1093/biomet/asm086>. 1204
- Pereira, L. A., Taylor-Rodríguez, D., and Gutiérrez, L. (2020). “A Bayesian nonparametric testing procedure for paired samples.” *Biometrics*, 76(4): 1133–1146. MR4186830. doi: <https://doi.org/10.1111/biom.13234>. 1189
- Pitman, J. (1996). “Some developments of the Blackwell-MacQueen urn scheme.” *Lecture Notes-Monograph Series*, 245–267. MR1481784. doi: <https://doi.org/10.1214/lnms/1215453576>. 1202, 1204
- Pitman, J. (2006). *Combinatorial Stochastic Processes: Ecole d’Eté de Probabilités de Saint-Flour XXXII-2002*. Springer. 1192
- Quintana, F. A. (1998). “Nonparametric Bayesian Analysis for Assessing Homogeneity in  $k \times l$  Contingency Tables with Fixed Right Margin Totals.” *Journal of the American Statistical Association*, 93(443): 1140–1149. MR1649208. doi: <https://doi.org/10.2307/2669857>. 1190



- Quintana, F. A., Müller, P., Jara, A., and MacEachern, S. N. (2020). “The Dependent Dirichlet Process and Related Models.” Technical report, Departamento de Estadística, Pontificia Universidad Católica de Chile. [1189](#)
- Regazzini, E. (1991). “Coherence, exchangeability and statistical models (de Finetti’s stance revisited).” In *Atti del Convegno “Sviluppi metodologici nei diversi approcci all’inferenza statistica”*, 101–137. Bologna: Pitagora. [1188](#)
- Rodríguez, A., Dunson, D. B., and Gelfand, A. E. (2008). “The nested Dirichlet process.” *Journal of the American Statistical Association*, 103(483): 1131–1154. [MR2528831](#). doi: <https://doi.org/10.1198/01621450800000553>. [1189](#), [1191](#)
- Soriano, J. and Ma, L. (2017). “Probabilistic multi-resolution scanning for two-sample differences.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(2): 547–572. [MR3611759](#). doi: <https://doi.org/10.1111/rssb.12180>. [1189](#)
- Taddy, M. A., Kottas, A., et al. (2012). “Mixture modeling for marked Poisson processes.” *Bayesian Analysis*, 7(2): 335–362. [MR2934954](#). doi: <https://doi.org/10.1214/12-BA711>. [1204](#)
- Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. (2006). “Hierarchical Dirichlet Processes.” *Journal of the American Statistical Association*, 101(476): 1566–1581. [MR2279480](#). doi: <https://doi.org/10.1198/016214506000000302>. [1187](#), [1189](#), [1193](#)
- Tokdar, S. T. and Martin, R. (2019). “Bayesian test of normality versus a Dirichlet process mixture alternative.” *Sankhya B*, 1–31. [MR4256311](#). doi: <https://doi.org/10.1007/s13571-019-00210-0>. [1201](#), [1215](#)
- Walker, S., Damien, P., and Lenk, P. (2004). “On priors with a Kullback–Leibler property.” *Journal of the American Statistical Association*, 99(466): 404–408. [MR2062826](#). doi: <https://doi.org/10.1198/016214504000000386>. [1200](#), [1201](#)
- Wu, Y. and Ghosal, S. (2008). “Kullback Leibler property of kernel mixture priors in Bayesian density estimation.” *Electronic Journal of Statistics*, 2: 298–331. [MR2399197](#). doi: <https://doi.org/10.1214/07-EJS130>. [1201](#)

### Acknowledgments

We are thankful to the Editor, Associate Editor and two anonymous referees for their constructive comments that helped us to significantly improve and clarify this manuscript.