

Partially Exchangeable Stochastic Block Models for Multilayer Networks

Daniele Durante^{*}, Francesco Gaffi[†], Antonio Lijoi^{*} and Igor Prünster^{*}

^{*}Bocconi Institute for Data Science and Analytics, Bocconi University

[†]Department of Mathematics, University of Maryland, College Park

Abstract

Multilayer networks generalize single-layered connectivity data in several directions. These generalizations include, among others, settings where multiple types of edges are observed among the same set of nodes (edge-colored networks) or where a single notion of connectivity is measured between nodes belonging to different pre-specified layers (node-colored networks). While progress has been made in statistical modeling of edge-colored networks, principled approaches that flexibly account for both within and across layer block-connectivity structures while incorporating layer information through a rigorous probabilistic construction are still lacking for node-colored multilayer networks. We fill this gap by introducing a novel class of partially exchangeable stochastic block models specified in terms of a hierarchical random partition prior for the allocation of nodes to groups, whose number is learned by the model. This goal is achieved without jeopardizing probabilistic coherence, uncertainty quantification and derivation of closed-form predictive within- and across-layer co-clustering probabilities. Our approach facilitates prior elicitation, the understanding of theoretical properties and the development of yet-unexplored predictive strategies for both the connections and the allocations of future incoming nodes. Posterior inference proceeds via a tractable collapsed Gibbs sampler, while performance is illustrated in simulations and in a real-world criminal network application. The notable gains achieved over competitors clarify the importance of developing general stochastic block models based on suitable node-exchangeability structures coherent with the type of multilayer network being analyzed.

Keywords: Bayesian Nonparametrics, Hierarchical Normalized Completely Random Measure, Node-Colored Network, Partial Exchangeability, Partially Exchangeable Partition Probability Function

1 Introduction

Modern network data encode elaborate connectivity information among a set of nodes. This growing complexity has motivated increasing efforts towards extending the wide statistical literature on single-layered networks to the context of multilayer networks. Recalling the comprehensive review by [Kivelä et al. \(2014\)](#) (see also Section S1.1 in the Supplementary Material), multilayer networks broadly refer to connectivity data characterizing edges among nodes in a multidimensional layering space, where each dimension denotes a feature of edges, nodes, or both. Remarkable special cases within this general class are edge-colored networks, where layers encode connections among the same, or possibly varying, set of

nodes w.r.t. different types of relationships, and node-colored networks, that encode a single notion of connectivity among nodes belonging to different pre-specified layers (Kivelä et al., 2014, Sect. 2.1–2.5).

In modeling the above data structures, attention has focused on a primary goal in network analysis, namely the identification of block-connectivity architectures based on shared patterns of edge formation among nodes (*e.g.*, Fortunato and Hric, 2016; Abbe, 2017; Lee and Wilkinson, 2019). This has led to several extensions of classical single-layered community detection algorithms (Girvan and Newman, 2002; Blondel et al., 2008), spectral clustering methods (Von Luxburg, 2007) and stochastic block models (SBMs) (Holland et al., 1983; Nowicki and Snijders, 2001) to identify grouping structures among nodes in multilayer networks. Within this context, key advances have been achieved in edge-colored settings (*e.g.*, Mucha et al., 2010; Stanley et al., 2016; Durante et al., 2017a,b; Wilson et al., 2017; Paul and Chen, 2020; Lei et al., 2020; Arroyo et al., 2021; Jing et al., 2021; Gao et al., 2022; Pensky and Wang, 2024; Noroozi and Pensky, 2024; Amini et al., 2024). Albeit relevant, these extensions do not naturally apply to node-colored networks which possess a substantially different structure. As showcased in Figure 1, these connectivity data can be reconstructed from a single-layered network whenever the nodes are labelled with some “type” defining a natural division into subpopulations. Key examples include connectivity networks among brain regions belonging to different lobes (Bullmore and Sporns, 2009), bill co-sponsorship networks between lawmakers from different parties (Briatte, 2016), social relationships among individuals divided in sociodemographic groups (Handcock et al., 2007), and co-attendances to summits of criminals belonging to different territorial units (Calderoni et al., 2017). To infer group structures among nodes in these ubiquitous networks, it is fundamental to devise rigorous models that: (i) incorporate flexible block-connectivity structures, both within and across layers, via a principled representation; (ii) automatically learn the number of nodes’ clusters from the data; (iii) provide formal uncertainty quantification; (iv) yield coherent projections as the network size grows; (v) allow for investigation of co-clustering properties; (vi) preserve computational tractability.

Despite the importance of node-colored networks, a state-of-the-art formulation capable of incorporating all these *desiderata* is still missing. Leveraging the previously-mentioned direct connection with node-labelled network data, one possibility to partially address such a gap is to treat the layer

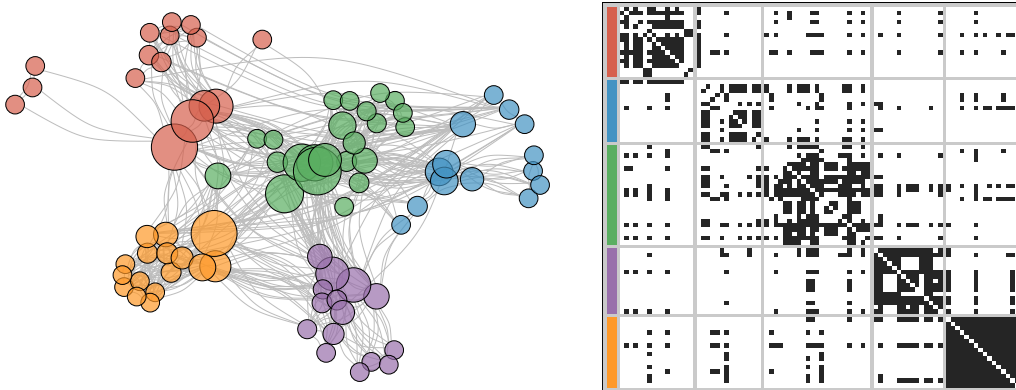


Figure 1: Left: Graphical representation of the node-colored *Infinito* network (e.g., [Calderoni et al., 2017](#)). Each node is a criminal and edges denote co-attendance to at least one of the summits of the Mafia organization. The size of each node is proportional to the corresponding betweenness, whereas the colors indicate layers which, in this context, define the membership to different *locali*, i.e., structural units administering crime in specific territories. Right: *Supra*-adjacency matrix representation of the network. Rows and columns correspond to nodes, whereas the entries measure the presence (black) or absence (white) of an edge for each pair of nodes. Side colors refer to layer (*locali*) division. See also Section 5.

membership information as a categorical attribute, and then rely on available attribute-assisted methods for single-layered networks (e.g., [Tallberg, 2004](#); [Xu et al., 2012](#); [Yang et al., 2013](#); [Sweet, 2015](#); [Newman and Clauset, 2016](#); [Zhang et al., 2016](#); [Binkiewicz et al., 2017](#); [Stanley et al., 2019](#); [Yan and Sarkar, 2021](#); [Mele et al., 2022](#); [Legramanti et al., 2022](#)). Although these solutions provide useful extensions of community detection algorithms, spectral clustering and SBMs, none of them jointly addresses the above *desiderata*. In fact, these extensions are either based on algorithmic strategies, which lack a generative model and fail to properly quantify uncertainty, or consider model-based approaches that, however, do not rely on consistent constructions, in the Kolmogorov sense, for the layer-dependent mechanism of nodes’ allocations to groups. This, in turn, undermines the development of models that preserve projectivity as the network size grows, a necessary condition for theoretically validated inference and principled prediction. As illustrated in Sections 4–5, achieving these objectives translates into remarkable empirical gains over state-of-the-art attribute-assisted models for single-layer networks.

In order to address the above *desiderata* (i)–(vi), in Section 2 we develop a novel class of partially exchangeable stochastic block models (pEx-SBM) for node-colored networks. The proposed class relies on a constructive and interpretable probabilistic representation which crucially combines stochastic equivalence structures ([Nowicki and Snijders, 2001](#)) for the edges in the *supra*-adjacency matrix of the node-colored network ([Kivelä et al., 2014](#)) with a partial exchangeability assumption ([de Finetti, 1938](#)) that allows to rigorously incorporate layers’ division into the mechanism of groups formation.

The latter assumption is motivated by the natural parallel between the division of statistical units into subpopulations, typical of a partially exchangeable framework, and the division of nodes into layers. This approach yields a Bayesian representation that assumes within- and across-layers edges to be conditionally independent realizations from Bernoulli variables whose probabilities only depend on the group allocation of the two involved nodes. The prior distribution for these allocations is characterized by a partially exchangeable partition probability function (pEPPF) (Camerlenghi et al., 2019).

The partially exchangeable regime behind pEPPFs encodes a more general distributional invariance than exchangeability: the observations are assumed to be drawn from different subpopulations, and the corresponding distribution is only invariant to permutations within the same subpopulation. Therefore, similarly to an exchangeable partition probability function (EPPF) (Pitman, 1995) that identifies random partitions induced by the ties of an exchangeable sequence, a pEPPF characterizes the random partition induced by the ties of a partially exchangeable array, for which each row stores observations from a different subpopulation. Crucially, such ties can occur also across rows, that is across different subpopulations. A concise account on partially exchangeable arrays is provided in the Supplementary Material. The parallel between subpopulations and network layers is, therefore, natural: we suppose a partially exchangeable array of latent attributes for the nodes within subpopulations corresponding to layer membership, and then rely on a broad class of discrete nonparametric priors for such an array, namely *hierarchical normalized completely random measures* (H-NRMIs) (Camerlenghi et al., 2019). Due to the discreteness of H-NRMIs, ties in the realizations of such auxiliary attribute may occur and nodes with the same attribute value are grouped together according to a probabilistic mechanism determined by a pEPPF. Moreover, the number of groups is random and learned from data. As is clear from this formulation, the auxiliary attribute is not object of inference. Rather, it represents an instrumental quantity that facilitates the constructive definition of general pEPPF priors on node allocations.

Crucially, the above construction yields a *Kolmogorov consistent* sequence of distributions $(P_V)_{V \geq 1}$, with P_V defined on the space of partitions of V nodes. This means that the distribution of the partition of $V - 1$ nodes obtained deleting uniformly at random one node from a random partition distributed as P_V is given by P_{V-1} , for any V . These sequences are known as *partition structures* (Kingman, 1978)

and guarantee theoretically-validated prediction and uncertainty quantification. We refer to the partition structures determined by a pEPPF as partially exchangeable partition (pExP) priors. Moreover, pExP priors induced by H-NRMIs allow us to analytically investigate yet-unexplored clustering and co-clustering probabilities, within and across layers. The corresponding expressions establish generalized notions of homophily that favor the creation of within-layer clusters relative to groups spanning across multiple layers. As shown in Section 3, these results also yield closed-form urn schemes which allow for the derivation of a collapsed Gibbs sampler for posterior inference, and novel k -step-ahead predictive schemes for the allocations and edges of future incoming nodes. Both are achieved preserving coherence between the updated and the full model, which is missing in state-of-the-art alternatives. In fact, principled and accurate k -step-ahead predictive schemes as the one we develop in Section 3.2 are lacking in the literature. Hence, our contribution along these lines represents an additional crucial advancement provided by pEx-SBM. To summarize, our pEx-SBM simultaneously meets the *desiderata* (i)–(vi). More generally, although the focus is on node-colored networks, our results clarify how pairing the specific structure of the multilayer network analyzed with the most suitable notion of node-exchangeability can yield key foundational, methodological and practical gains, thus opening to future research in the broader class of multilayer networks (including edge-colored ones), beyond node-colored settings.

The extensive simulation studies (Section 4) and the application to a multilayer criminal network (Section 5) showcase the practical gains in point estimation (including empirical evidence of frequentist posterior consistency as V grows), uncertainty quantification and prediction of pEx-SBMs, when compared to state-of-the-art competitors (Blondel et al., 2008; Zhang et al., 2016; Binkiewicz et al., 2017; Côme et al., 2021; Legramanti et al., 2022). Finally, as highlighted in Section 6, although the focus is on the ubiquitous binary undirected node-colored networks, suitable adaptations of the novel modeling framework underlying pEx-SBMs facilitate the inclusion of other relevant network settings. Proofs, additional results and further empirical investigations are deferred to the Supplementary Material.

2 Partially Exchangeable Stochastic Block Models (pEx-SBM)

Networks are typically represented through an adjacency matrix with rows and columns indexed by the V nodes, and binary entries taking value 1 if there is an edge between the corresponding pair of nodes

and 0 otherwise. Similarly, a node-colored network can be represented via a *supra*-adjacency (block) matrix with diagonal and off-diagonal blocks encoding connectivity structures within each layer and across pairs of layers, respectively; see Figure 1. We consider here a generic undirected node-colored network with V nodes divided in d layers. Each layer j contains V_j nodes with $V = \sum_{j=1}^d V_j$. The corresponding $V \times V$ binary and symmetric *supra*-adjacency matrix is denoted by \mathbf{Y} . Since nodes are stacked consecutively in \mathbf{Y} , there is a one-to-one correspondence between the row-column indexes of the *supra*-adjacency matrix, namely $v = 1, \dots, V$, and the pairs (j, i) , for $i = 1, \dots, V_j$ and $j = 1, \dots, d$ denoting the i -th node in the j -th layer. In the following we use v or (j, i) based on convenience.

We now derive the proposed partially exchangeable stochastic block model (pEx-SBM). The edges y_{vu} , $v = 2, \dots, V$, $u = 1, \dots, v-1$, in \mathbf{Y} are assumed to be conditionally independent realizations from a Bernoulli distribution, given probabilities depending solely on the group allocations of the two involved nodes. These allocations are, in turn, assigned a partially exchangeable partition (pExP) prior incorporating layer memberships. Notice that this construction does not imply that edges are independent from layers. In fact, marginalizing out the layer-informed pExP prior yields a direct dependence among layers and edges. Including layer information also in block probabilities is, in principle, possible. However, while in edge-colored settings this can be useful to model the different types of relationships, in node-colored networks that encode a single notion of connectivity such a choice would arguably add redundant flexibility, while complicating prior elicitation, posterior computation and interpretability.

2.1 Stochastic block models

Set $[n] := \{1, \dots, n\}$ for any generic positive integer n . If each node is identified with the index $v \in [V]$, as in the *supra*-adjacency matrix \mathbf{Y} , each partition of nodes in H clusters can be represented by disjoint sets $(\mathbb{V}_1, \dots, \mathbb{V}_H)$, where $\mathbb{V}_h \subset [V]$ and $\bigcup_{h=1}^H \mathbb{V}_h = [V]$. To make such representation unique, as customary, we label the clusters in *order of appearance*, so that \mathbb{V}_1 contains node $v = 1$, \mathbb{V}_2 contains node v^* , where $v^* = \min\{[V] \setminus \mathbb{V}_1\}$, etc. Define $\mathbf{z} = (z_1, \dots, z_V) \in [H]^V$ as the allocations vector corresponding to the node partition $(\mathbb{V}_1, \dots, \mathbb{V}_H)$, such that $z_v = h$ if and only if $v \in \mathbb{V}_h$. Recalling SBM representations for single-layered networks (Holland et al., 1983; Nowicki and Snijders, 2001), we assume that $(y_{vu} \mid z_v = h, z_u = h', \psi_{hh'}) \sim \text{Bern}(\psi_{hh'})$, independently for each $v = 2, \dots, V$

and $u = 1, \dots, v - 1$ with $\psi_{hh'} \in (0, 1)$ the probability of an edge among generic nodes in groups h and h' , for $h, h' \in [H]$. This model for the *supra*-adjacency matrix facilitates efficient inference on homogenous node groups, while allowing to incorporate a variety of heterogeneous block-interactions among groups via the $H \times H$ matrix Ψ with entries $\psi_{hh'}$ for $h, h' \in [H]$. These include any combination of community, core-periphery and disassortative structures, thus improving flexibility w.r.t. basic community detection algorithms (Fortunato and Hric, 2016; Abbe, 2017; Lee and Wilkinson, 2019).

Consistent with the above focus on inferring grouping structures among stochastically-equivalent nodes, we follow classical SBM implementations for single-layered networks (*e.g.*, Schmidt and Morup, 2013) by assuming independent $\text{beta}(a, b)$ priors for the block-probabilities in Ψ , and then marginalize these quantities out from the model to obtain the beta-binomial likelihood

$$p(\mathbf{Y} \mid \mathbf{z}) = \prod_{h=1}^H \prod_{h'=1}^h \frac{\text{B}(a + m_{hh'}, b + \overline{m}_{hh'})}{\text{B}(a, b)}, \quad (1)$$

where m_{hh} and $\overline{m}_{hh'}$ are the total number of edges and non-edges among the nodes in groups h and h' , respectively. Although inference on Ψ is also of interest, as clarified in Section 3, treating it as a nuisance parameter facilitates computation, inference and prediction for the main object of analysis \mathbf{z} .

Importantly, the likelihood in (1) does not account for layer division. In fact, with such a distribution for the entries of the *supra*-adjacency matrix, within- and across-layers edges are not disentangled. The proposed pEx-SBMs crucially address this shortcoming by embedding the layers' division information into the prior for the allocation vector \mathbf{z} . Conversely, single-layered SBMs, which combine (1) with Dirichlet-multinomial (Nowicki and Snijders, 2001), Dirichlet process (Kemp et al., 2006), mixture-of-finite-mixture (Geng et al., 2019) or general unsupervised Gibbs-type (Legramanti et al., 2022) priors for \mathbf{z} , would be conceptually and practically suboptimal, since these priors are not designed to incorporate structure from layer division. Recalling Section 1 and Figure 1, one expects nodes in the same layer to be more likely to exhibit similar connectivity patterns, with these patterns possibly varying both within and across layers. This would translate into a prior for \mathbf{z} that reinforces the formation of one or multiple within-layer groups, while still allowing for the possibility of creating clusters that comprise nodes in different layers. The latter property is important since the exogenous layer division does not necessarily overlap with the endogenous stochastic equivalence structures within the network.

2.2 Partially exchangeable partition priors

Our goal is to embed the layers’ division information into the model via the prior on the allocation vector. This translates into completing the likelihood (1) with a pExp prior for \mathbf{z} informed by the division in layers. Set $\mathbf{V} = (V_1, \dots, V_d)$, where we recall that V_j stands for the number of nodes in the j -th layer for $j \in [d]$, and the total number of nodes is $V = \sum_{j=1}^d V_j$. We assume

$$\mathbf{z} \sim \text{pExp}(\mathbf{V}), \quad (2)$$

meaning that, given the partition structure defined by a pEPPF (Camerlenghi et al., 2019), the distribution of \mathbf{z} is determined by this pEPPF evaluated in the vector of group sizes $\mathbf{V} = (V_1, \dots, V_d)$. Let \mathbf{e}_j be the vector with all zero entries but the j -th which is 1, for $j \in [d]$. Then, in view of our construction, marginalizing $\text{pExp}(\mathbf{V} + \mathbf{e}_j)$ w.r.t. a node in the j -th group yields $\text{pExp}(\mathbf{V})$ for any $j \in [d]$.

REMARK 2.1. The direct construction of a prior for a random partition is challenging even in the exchangeable case. Operationally, it is more convenient to start from an exchangeable sequence directed by a discrete random probability measure and then obtain the partition structure by marginalizing over the labels and the probability weights. Hence, *a fortiori*, we follow this approach also in our more general setting and induce a pExp prior through a partially exchangeable array of latent auxiliary attributes of the nodes, directed by a large class of discrete random probability measures known as *hierarchical normalized completely random measures* (H-NRMIs). See Definition 2.2. As anticipated in Section 1, the array of latent auxiliary attributes is not of interest for inference. Rather, it defines a purely instrumental quantity that is useful for constructively defining the pExp prior.

2.2.1 Hierarchical construction

The pExp prior in (2) allows for an appealing generative construction employing partial exchangeability. Let \mathbf{X} denote a random array with row j having V_j entries $x_{ji} \in \mathbb{X}$ for $j \in [d]$ and $i \in [V_j]$. Each entry acts as a latent auxiliary attribute corresponding to a node $v \in [V]$, according to the previously mentioned bijection $v \leftrightarrow (j, i)$. Partial exchangeability is enforced by assuming the rows of \mathbf{X} to be directed by discrete random probability measures $\tilde{P}_1, \dots, \tilde{P}_d$. This setup is ideally suited for node-colored networks as it incorporates a generalized notion of homophily: if $(\tilde{P}_1, \dots, \tilde{P}_d)$ are

marginally identically distributed, the probability of a tie across the rows of \mathbf{X} is always less than or equal to the probability of a tie within the same row, with equality holding when the \tilde{P}_j 's coincide almost surely (Franzolini, 2022). Hence, nodes in the same layer are more likely to exhibit the same latent attribute *a priori*, and therefore to show a similar connectivity behaviour. Since our focus is on developing a projective representation whose construction and theoretical properties hold, coherently, for any, finite, network size $(V_1, \dots, V_d) \in \mathbb{N}^d$, the natural invariance condition to assume is partial exchangeability of the infinite array \mathbf{X}^∞ , as clarified by the following Definition 2.1.

DEFINITION 2.1. \mathbf{X}^∞ is partially exchangeable if, for any $(V_1, \dots, V_d) \in \mathbb{N}^d$,

$$(x_{ji} : j \in [d]; i \in [V_j]) \stackrel{d}{=} (x_{j, \pi_j(i)} : j \in [d]; i \in [V_j]), \quad (3)$$

for every family of index permutations $\{\pi_1, \dots, \pi_d\}$, where $\stackrel{d}{=}$ denotes equality in distribution.

Definition 2.1 implies that the joint distribution of the entries in \mathbf{X}^∞ is invariant w.r.t. within-layer permutations of the nodes, but not necessarily across-layers ones. When this assumption holds, de Finetti's representation theorem (de Finetti, 1938) ensures the existence of a vector of random probability measures $(\tilde{P}_1, \dots, \tilde{P}_d)$ such that, for any $j_1, \dots, j_k \in [d]$ and $i_1, \dots, i_k \geq 1$, one has

$$x_{j_1 i_1}, \dots, x_{j_k i_k} \mid (\tilde{P}_1, \dots, \tilde{P}_d) \stackrel{\text{iid}}{\sim} \tilde{P}_{j_1} \times \dots \times \tilde{P}_{j_k}, \quad (\tilde{P}_1, \dots, \tilde{P}_d) \sim Q, \quad (4)$$

for some distribution Q , named *de Finetti measure*, on the space of d -dimensional vectors of probability measures on \mathbb{X} . Note that, when all \tilde{P}_j 's coincide (or $d = 1$), (4) reduces to de Finetti's representation of an exchangeable sequence. Its extension to distributional invariance within layers, but not necessarily across, allows a probabilistically sound incorporation of the layer division in the distribution of \mathbf{X}^∞ .

To complete representation (4), the prior Q is to be specified. Since the final goal is to induce a prior on \mathbf{z} , our Bayesian nonparametric approach focuses on priors Q selecting, almost surely, vectors of discrete probability measures. With V_1, \dots, V_d being positive integers associated with the observed network \mathbf{Y} , this choice implies that the entries $\{x_{ji} : j \in [d]; i \in [V_j]\}$ of the finite array \mathbf{X} can display ties and, hence, it is possible, and natural, to induce a prior on node partitions by interpreting these ties as a co-clustering relationship. More specifically, let $\mathbf{x}^* = (x_1^*, \dots, x_H^*)$ denote the unique values of \mathbf{X} , with H the number of clusters, which can be random. Then, the random array \mathbf{Z} of node

allocations z_{ji} , for $j \in [d]$ and $i \in [V_j]$, has entries which can be defined as

$$z_{ji} = h \quad \text{if and only if} \quad x_{ji} = x_h^*, \quad \text{for } h \in [H], \quad (5)$$

thus inducing a prior on the space of random partitions of $\{1, \dots, V\}$ from the ties in \mathbf{X} . Clearly, its clustering and co-clustering properties depend on the choice of Q . We opt for Q belonging to the broad class of H-NRMIs (Camerlenghi et al., 2019), which includes the hierarchical Dirichlet process (H-DP) (Teh et al., 2006) and the hierarchical normalized stable process (H-NSP) as noteworthy special cases. The choice of H-NRMIs is motivated by three main reasons: (i) the nodes can be clustered both within and across layers; (ii) the number of clusters in the network is random and learned from the data; (iii) mathematical and computational tractability. Extending the NRMIs construction in the Supplementary Material to the hierarchical setting, it is possible to define H-NRMIs as follows.

DEFINITION 2.2. $(\tilde{P}_1, \dots, \tilde{P}_d)$ is a vector of H-NRMIs on \mathbb{X} with parameters $(\rho, \rho_0, c, c_0, P_0)$ if

$$\tilde{P}_1, \dots, \tilde{P}_d \mid \tilde{P}_0 \stackrel{iid}{\sim} \text{NRMI}(\rho, c, \tilde{P}_0), \quad \tilde{P}_0 \sim \text{NRMI}(\rho_0, c_0, P_0), \quad (6)$$

where ρ and ρ_0 are the jump components of the underlying Lévy intensities of $\tilde{P}_j \mid \tilde{P}_0$ for any $j \in [d]$ and \tilde{P}_0 , respectively, with $c, c_0 \in \mathbb{R}^+$, while P_0 is a diffuse probability measure on \mathbb{X} .

The role of the parameters (ρ, ρ_0, c, c_0) characterizing specific H-NRMIs is described in the Supplementary Material. As desired, (4) combined with (6) generates ties among latent node attributes both within and across layers, the latter being induced by the almost sure discreteness of \tilde{P}_0 at the root of the hierarchy. The distribution of the allocation vector associated to the unique values in \mathbf{x}^* is specified as

$$\mathbf{z} \sim \text{pExp}(\mathbf{V}; \rho, \rho_0, c, c_0). \quad (7)$$

The prior induced on this grouping structure arising from \mathbf{X} is characterized by its pEPPF. This means that the prior probability mass function for \mathbf{z} (or alternatively \mathbf{Z}), can be derived from the distribution induced by the H-NRMI construction on the array of positive integers $(\mathbf{n}_1, \dots, \mathbf{n}_d)$, where each $\mathbf{n}_j = (n_{j1}, \dots, n_{jH})$ represents the allocation frequencies to H different groups of the V_j nodes in layer j , for each $j \in [d]$. As shown in Camerlenghi et al. (2019), this distribution is equal to

$$p_H^{(V)}(\mathbf{n}_1, \dots, \mathbf{n}_d) = \sum_{\boldsymbol{\ell}} \sum_{\mathbf{q}} \Phi_{H,0}^{(|\boldsymbol{\ell}|)}(\ell_{\cdot 1}, \dots, \ell_{\cdot H}) \prod_{j=1}^d \prod_{h=1}^H \frac{1}{\ell_{jh}!} \binom{n_{jh}}{q_{jh1}, \dots, q_{jh\ell_{jh}}} \Phi_{\ell_{j\cdot}, j}^{(V_j)}(\mathbf{q}_{j1}, \dots, \mathbf{q}_{jH}), \quad (8)$$

where n_{jh} is the total number of nodes in layer j allocated to group h , $\boldsymbol{\ell}$ is an array with generic element

$\ell_{jh} \in [n_{jh}]$, for $j \in [d]$ and $h \in [H]$, whose meaning will be clarified in detail later, while $\ell_{j\cdot} = \sum_{h=1}^H \ell_{jh}$, $\ell_{\cdot h} = \sum_{j=1}^d \ell_{jh}$ and $|\mathcal{L}| = \sum_{j=1}^d \ell_{j\cdot}$. Finally, $\mathbf{q}_{jh} = (q_{jh1}, \dots, q_{jh\ell_{jh}})$ denotes a vector of integers such that $\sum_{t=1}^{\ell_{jh}} q_{jht} = n_{jh}$, while the functions $\Phi_{\ell_{j\cdot}, j}^{(V_j)}$ and $\Phi_{H, 0}^{(|\mathcal{L}|)}$ are the EPPFs associated to NRMs with parameters (c, ρ) and (c_0, ρ_0) , respectively. Such EPPFs characterize the probability distribution of the exchangeable random partitions of V_j nodes in $\ell_{j\cdot}$ groups and $|\mathcal{L}|$ elements in H clusters, respectively; see [James et al. \(2009\)](#) for the closed-form expressions of such functions.

To help intuition with the pExp prior and the quantities involved in (8), we recast the *Chinese restaurant franchise* (CRF) metaphor ([Teh et al., 2006](#)) within the node-colored network setting for the whole class of H-NRMs. According to this metaphor, in each layer the nodes are first allocated to within-layer subgroups, and then each subgroup is assigned a sociability profile from a single list which is common across layers. Nodes in subgroups with the same sociability profile are, therefore, naturally characterized by similar connectivity patterns in the multilayer network and, hence, have the same group allocation in \mathbf{z} . Notice that the within-layer division in subgroups is only a latent quantity which is not of interest for inference, but rather provides an intermediate clustering that leads, under the hierarchical construction, to the formation of the grouping structure among nodes which we aim to infer (i.e., the one defined by the sociability profiles). Nonetheless, such a within-layer partition nested in the sociability profiles further clarifies how layer information is effectively incorporated in the formation of node groups. In fact, two nodes in the same layer have equal group indicator if they are allocated to the same subgroup or, when they belong to different subgroups, the same sociability profile is assigned to these two subgroups. Conversely, nodes in different layers belong to a same final group only if the corresponding subgroups have the same sociability profile. This clarifies the role of layer division in reinforcing the formation of within-layer groups without ruling out across-layer clusters.

Consistent with the above metaphor, each \tilde{P}_j in (6) allocates nodes to subgroups within each layer, whereas \tilde{P}_0 assigns to subgroups the sociability profiles from a list shared across layers. Due to the discreteness of \tilde{P}_0 , the same sociability profile can be assigned to multiple subgroups both within and across layers. This facilitates also the interpretation of the arrays \mathcal{L} and \mathbf{q} in (8). More specifically, ℓ_{jh} is the number of subgroups in layer j with sociability profile h and q_{jht} is the number of nodes

in layer j assigned to the t -th subgroup with sociability profile h ; consequently, ℓ_j is the number of subgroups in layer j , ℓ_h is the total number of subgroups with sociability profile h , and n_{jh} denotes the number of nodes in layer j with profile h . The metaphor is also useful to exemplify the role of the EPPFs $\Phi_{\ell_j, j}^{(V_j)}$ and $\Phi_{H, 0}^{(|\ell|)}$ in the sampling of the two nested partitions. The former regulates, for any $j \in [d]$, the distribution of the division in subgroups within each layer, whereas the latter drives the sociability profile assignment to the subgroups in the different layers.

The above discussion also clarifies the hierarchical process through which the auxiliary attribute values in \mathbf{X} are generated. As already discussed, \mathbf{X} is treated in our construction as a latent quantity useful to define, study and manage the pExp prior for \mathbf{z} . While \mathbf{X} is not of interest for inference, its generative process is, however, useful for posterior sampling and prediction under the pEx-SBM.

1. For each layer $j \in [d]$, sample the vector $\mathbf{t}_j = (t_{j1}, \dots, t_{jV_j})$ of the subgroup labels as $t_{j1}, \dots, t_{jV_j} \mid \tilde{Q}_j \stackrel{\text{iid}}{\sim} \tilde{Q}_j$, with $\tilde{Q}_j \sim \text{NRMI}(\rho, c, G_j)$, for some diffuse probability measure G_j .
2. For each layer $j \in [d]$, allocate the nodes to subgroups according to the partition induced by the ties in \mathbf{t}_j . This yields the subgroup allocation array \mathbf{W} with entries

$$w_{ji} = \tau \quad \text{if and only if} \quad t_{ji} = t_{j\tau}^*, \quad (9)$$

for $\tau \in [\ell_j]$ with $\mathbf{t}_j^* = (t_{j1}^*, \dots, t_{j\ell_j}^*)$ the unique values of \mathbf{t}_j in order of occurrence.

3. Sample the array of sociability profile labels \mathbf{S} , which has an entry $s_{j\tau}$ for each subgroup, as

$$s_{j\tau} \mid \tilde{P}_0 \stackrel{\text{iid}}{\sim} \tilde{P}_0, \quad (10)$$

for each $\tau \in [\ell_j]$ and $j \in [d]$, with $\tilde{P}_0 \sim \text{NRMI}(c_0, \rho_0, P_0)$.

4. Obtain \mathbf{X} by assigning to each node the sociability profile label of its subgroup, that is $x_{ji} := s_{j\tau}$ whenever $w_{ji} = \tau$.

Recalling (5), the ties in \mathbf{X} yield the final clustering structure encoded in \mathbf{z} (or alternatively \mathbf{Z}), which comprises the sociability profiles of the V nodes. As \mathbf{Z} takes values in $[H]$, the j -th row of \mathbf{W} takes values in $[\ell_j]$, for any $j \in [d]$. We will represent sociability profiles and subgroups with these indices. Note that, since we are considering ties only for clustering, G_j in step 1. and P_0 in step 3. are arbitrary, as long as they are diffuse. As discussed previously, the subgroup allocation array \mathbf{W} should

be interpreted as a variable augmentation scheme. Such a quantity is not object of inference, but is crucial to recover tractable full conditionals for posterior computation and prediction on \mathbf{z} and \mathbf{Y} .

A crucial feature of pExP priors is that Kolmogorov consistency holds by construction. This implies that the joint distribution of any finite-dimensional sequence of node allocations can be obtained by marginalizing one of higher dimension defined through the same generative scheme. Consequently, our formulation remains coherent for every network size and preserves projectivity to growing number of nodes. This property, which is not met, e.g., by the supervised approach of [Legramanti et al. \(2022\)](#), is not only crucial in deriving novel and principled predictive strategies, but is also conceptually desirable since, in practice, not all the nodes in a network are observed. Hence, it is important that the model postulated for a subset of the whole network remains coherent also at the population level.

REMARK 2.2. Despite the potential for multilayer network modeling of pExP priors induced by HN-RMIs, to the best of our knowledge, the only contribution moving in a related direction is the one by [Amini et al. \(2024\)](#) which, however, focuses on edge-colored (multiplex) networks, rather than node-colored ones, and considers only the special case of H-DP priors. Such a contribution can be recovered as a special case of our framework by imposing a block-diagonal *supra*-adjacency matrix and selecting the H-DP from the general H-NRMIs class. However, although [Amini et al. \(2024\)](#) provide a valuable contribution in the substantially-different edge-colored setting, the partial exchangeability assumption behind the H-DP and its induced clustering and co-clustering mechanisms do not align naturally with the intrinsic structure of edge-colored networks. In fact, when the same set of nodes is replicated across layers, as in edge-colored settings, such an assumption forces the model to ignore this key node-identity information across layers. As a result, it is also unclear, from a network perspective, how to interpret groups shared across multiple layers, which may include copies of the same nodes. Similar comments apply to the recent contribution by [Josephs et al. \(2023\)](#) which relies on the nested Dirichlet process (also a partially exchangeable prior) and still focuses mainly on edge-colored, rather than node-colored, networks. As highlighted in Section 6, an assumption of separate exchangeability (e.g., [Rebaudo et al., 2024](#)) would be more coherent with the structure of edge-colored networks.

Finally, notice that the use of specific discrete hierarchical structures can be found also in [Dempsey](#)

et al. (2022), but with a focus on edge-exchangeable, rather than node-exchangeable, models. This perspective is desirable in situations where edge-specific structures are of interest, regardless of the specific identities of the nodes among which these edges exist. This prevents from inferring group structures among nodes via SBMs, possibly informed by layer partitions. When this is the focus of inference, one has to identify nodes (and not edges) as statistical units and, hence, enforce notions of node exchangeability. In fact, in Dempsey et al. (2022, Remark 3.1), the hierarchical structure is used to generate, and infer, the interaction process itself, rather than as a prior on nodes' partitions.

2.2.2 Clustering and co-clustering properties

Clustering and co-clustering properties play a crucial role in stochastic block models. Nonetheless analytical results are missing even in the more general H-NRMI literature (Camerlenghi et al., 2019). Here we fill this gap by deriving novel explicit expressions for clustering and co-clustering probabilities from a predictive perspective. In addition to providing insights into the model's behaviour, such results are at the basis of the sampling algorithm and the prediction schemes developed in Section 3.

In the sequel we use the vectorized forms \mathbf{z} and \mathbf{w} for the sociability profiles and the subgroups allocations of \mathbf{Z} and \mathbf{W} , respectively, defined as in (5) and (9). With \mathbf{z}^{-v} and \mathbf{w}^{-v} we instead indicate the $(V-1)$ -dimensional vectors obtained from \mathbf{z} and \mathbf{w} respectively, upon removing node v and suitably rearranging the labels in order of appearance. Similarly, ℓ^{-v} and \mathbf{q}^{-v} also stand for the arrays with node v removed from the corresponding counts. Moreover, define the set

$$\mathbb{T}_{jh}^{-v} := \{\tau \in [\ell_j^{-v}] \mid s_{j\tau}^{-v} = x_h^{*-v}\}, \quad (11)$$

where $s_{j\tau}$ is an entry of the sociability array \mathbf{S}^{-v} obtained as in (10), and \mathbf{x}^{*-v} displays the unique sociability labels, as in (5), but disregarding node v . The indices in \mathbb{T}_{jh}^{-v} are associated to the subgroups in layer j with profile h . Indeed, $|\mathbb{T}_{jh}^{-v}| = \ell_{jh}^{-v}$. With these settings we can state the following result.

PROPOSITION 2.3. *Let \mathbf{z} denote a random allocation vector such that $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \rho, \rho_0, c, c_0)$ as in (7), and j be the layer of the generic node $v \in [V]$. Then, with \mathbb{T}_{jh}^{-v} as in (11) and $\tau_{\text{new}} := \ell_j^{-v} + 1$ standing for a new subgroup in layer j , we have*

$$\mathbb{P}(z_v = h \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \mathbb{P}(v \rightsquigarrow \mathbb{T}_{jh}^{-v}) + \mathbb{P}(v \rightsquigarrow \tau_{\text{new}} \mid \tau_{\text{new}} \leftrightsquigarrow h) \cdot \mathbb{P}(\tau_{\text{new}} \leftrightsquigarrow h), \quad (12)$$

for any $h \in [H^{-v} + 1]$, with \rightsquigarrow and \leftharpoonup meaning “assigned to”, and

$$\begin{aligned}\mathbb{P}(v \rightsquigarrow \mathbb{T}_{jh}^{-v}) &= [\sum_{t=1}^{\ell_{j\cdot}^{-v}} \Phi_{\ell_{j\cdot}^{-v}, j}^{(V_j)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jh}^{-v} + \mathbf{e}_t, \dots, \mathbf{q}_{jH^{-v}}^{-v})] / \Phi_{\ell_{j\cdot}^{-v}, j}^{(V_j-1)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jH^{-v}}^{-v}), \\ \mathbb{P}(v \rightsquigarrow \tau_{new} \mid \tau_{new} \leftharpoonup h) &= \Phi_{\ell_{j\cdot}^{-v}+1, j}^{(V_j)}(\mathbf{q}_{j1}^{-v}, \dots, (\mathbf{q}_{jh}^{-v}, 1), \dots, \mathbf{q}_{jH_h^{-v}}^{-v}) / \Phi_{\ell_{j\cdot}^{-v}, j}^{(V_j-1)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jH^{-v}}^{-v}), \\ \mathbb{P}(\tau_{new} \leftharpoonup h) &= \Phi_{H_h^{-v}, 0}^{(|\ell^{-v}|+1)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot h}^{-v} + 1, \dots, \ell_{\cdot H_h^{-v}}^{-v}) / \Phi_{H^{-v}, 0}^{(|\ell^{-v}|)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot H^{-v}}^{-v}),\end{aligned}\tag{13}$$

where \mathbf{e}_t is the t -th vector of the $\ell_{j\cdot}^{-v}$ -dimensional canonical basis, and $H_h^{-v} := h \vee H^{-v}$.

REMARK 2.4. The expression for $\mathbb{P}(z_v = h \mid \mathbf{z}^{-v}, \mathbf{w}^{-v})$ in Proposition 2.3 can be naturally interpreted in terms of the previously-introduced metaphor. The first summand is the probability of node v being allocated to any of the already-occupied subgroups in j with sociability profile h , while the second is the probability of being allocated to a new subgroup, which has been assigned sociability profile h . If profile h is not present in layer j , then $\ell_{jh}^{-v} = 0$, the first summand disappears and h can be assigned to v only by creating a new subgroup. Moreover, since \mathbf{z}^{-v} is arranged according to the order of occurrence in \mathbf{x}^{-v} , $h = H^{-v} + 1$ represents the case of a sociability profile new to the whole network.

Corollary 2.5 specializes Proposition 2.3 to the popular H-DP case, thus re-obtaining its known urn scheme (Teh et al., 2006). An analogous result for the H-NSP is given in the Supplementary Material.

COROLLARY 2.5. Let $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \theta, \theta_0)$, where $\text{pExp}(\mathbf{V}; \theta, \theta_0)$ is the partition structure induced by a H-DP with concentration parameters θ and θ_0 . Then, for each $h \in [H^{-v} + 1]$,

$$\begin{aligned}\mathbb{P}(z_v = h \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) &= \mathbb{1}_{\{\ell_{\cdot h}^{-v} = 0\}} \frac{\theta_0}{(\theta_0 + |\ell^{-v}|)} \frac{\theta}{(\theta + V_j - 1)} + \mathbb{1}_{\{\ell_{\cdot h}^{-v} \neq 0\}} \left[\frac{\ell_{jh}^{-v}}{(\theta_0 + |\ell^{-v}|)} \frac{\theta}{(\theta + V_j - 1)} + \frac{n_{jh}^{-v}}{\theta + V_j - 1} \right].\end{aligned}\tag{14}$$

In Corollary 2.5 the two scenarios are more explicit. If h is a new sociability profile for the whole network, $\ell_{\cdot h}^{-v} = 0$ and the probability of node v receiving h coincides with the probability of being assigned a new profile at a new subgroup. Conversely, if $\ell_{\cdot h}^{-v} \neq 0$, then h is not new to the network and we sum two terms, namely the probability of assigning h to a new subgroup and the probability of being allocated to an already-occupied subgroup among those with profile h . The latter is 0 if the sociability profile h is new for the layer of node v , since $n_{jh}^{-v} = 0$.

Besides clarifying the generative nature of \mathbf{z} , as highlighted in Section 3, the results in Proposition 2.3 and Corollary 2.5 are a key to develop tractable collapsed Gibbs sampling schemes for inference

on the posterior distribution $p(\mathbf{z}|\mathbf{Y}) \propto p(\mathbf{z})p(\mathbf{Y}|\mathbf{z})$ induced by the model defined in (1) and (7). The co-clustering probabilities derived in Theorem 2.6 below are instead crucial for prior elicitation and to devise rigorous predictive strategies for both the allocations and edges of future incoming nodes. The apex $^{-vu}$ denotes all the previously-defined quantities, evaluated disregarding nodes $v, u \in [V]$.

THEOREM 2.6. *Let \mathbf{z} be a random allocation vector such that $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \rho, \rho_0, c, c_0)$. Moreover, let j and j' be the layers of any two distinct nodes $v, u \in [V]$, respectively. Then, with τ_{new} and τ'_{new} indicating new subgroups in layers j and j' , respectively, we have:*

(1) *If $j = j'$,*

$$\begin{aligned} \mathbb{P}(z_v = z_u \mid \mathbf{z}^{-vu}, \mathbf{w}^{-vu}) &= \sum_{h=1}^{H^{-vu}+1} \left[\mathbb{P}(\{v, u\} \rightsquigarrow \mathbb{T}_{jh}^{-vu}) \right. \\ &\quad + \mathbb{P}(\tau_{new} \leftarrow h) \cdot [\mathbb{P}(\{v, u\} \rightsquigarrow \tau_{new} \mid \tau_{new} \leftarrow h) + 2\mathbb{P}(v \rightsquigarrow \tau_{new}, u \rightsquigarrow \mathbb{T}_{jh}^{-vu} \mid \tau_{new} \leftarrow h)] \\ &\quad \left. + \mathbb{P}(\tau_{new} \leftarrow h, \tau'_{new} \leftarrow h) \cdot \mathbb{P}(v \rightsquigarrow \tau_{new}, u \rightsquigarrow \tau'_{new} \mid \tau_{new} \leftarrow h, \tau'_{new} \leftarrow h) \right], \end{aligned} \quad (15)$$

with

$$\begin{aligned} \mathbb{P}(\{v, u\} \rightsquigarrow \mathbb{T}_{jh}^{-vu}) &= \frac{\sum_{t=1}^{\ell_{jh}^{-vu}} \Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jh}^{-vu} + 2\mathbf{e}_t, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}{\Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j-2)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})} \\ &\quad + \frac{\sum_{A \in C_{\ell_{jh}^{-vu}, 2}} \Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jh}^{-vu} + \mathbf{e}_A, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}{\Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j-2)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}, \\ \mathbb{P}(\tau_{new} \leftarrow h) &= \frac{\Phi_{H_h^{-vu}, 0}^{(|\ell^{-vu}|+1)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot h}^{-vu} + 1, \dots, \ell_{\cdot H^{-vu}}^{-vu})}{\Phi_{H^{-vu}, 0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot H^{-vu}}^{-vu})}, \\ \mathbb{P}(\{v, u\} \rightsquigarrow \tau_{new} \mid \tau_{new} \leftarrow h) &= \frac{\Phi_{\ell_{j\cdot}^{-vu}+1, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, (\mathbf{q}_{jh}^{-vu}, 2), \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}{\Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j-2)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}, \\ \mathbb{P}(v \rightsquigarrow \tau_{new}, u \rightsquigarrow \mathbb{T}_{jh}^{-vu} \mid \tau_{new} \leftarrow h) &= \frac{\sum_{t=1}^{\ell_{jh}^{-vu}} \Phi_{\ell_{j\cdot}^{-vu}+1, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, (\mathbf{q}_{jh}^{-vu} + \mathbf{e}_t, 1), \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}{\Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j-2)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}, \\ \mathbb{P}(\tau_{new} \leftarrow h, \tau'_{new} \leftarrow h) &= \frac{\Phi_{H_h^{-vu}, 0}^{(|\ell^{-vu}|+2)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot h}^{-vu} + 2, \dots, \ell_{\cdot H^{-vu}}^{-vu})}{\Phi_{H^{-vu}, 0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot H^{-vu}}^{-vu})}, \\ \mathbb{P}(v \rightsquigarrow \tau_{new}, u \rightsquigarrow \tau'_{new} \mid \tau_{new} \leftarrow h, \tau'_{new} \leftarrow h) &= \frac{\Phi_{\ell_{j\cdot}^{-vu}+2, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, (\mathbf{q}_{jh}^{-vu}, 1, 1), \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}{\Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j-2)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}, \end{aligned} \quad (16)$$

for any $h \in [H^{-vu} + 1]$, where $C_{\ell, 2}$ denotes the set of 2-combinations of indices in $[\ell]$, whereas, for $A \in C_{\ell, 2}$, \mathbf{e}_A is a ℓ -dimensional vector with $\mathbf{e}_i = 1$ if $i \in A$ and 0 otherwise. Finally $H_h^{-vu} = H^{-vu} \vee h$.

(2) If $j \neq j'$,

$$\begin{aligned}
\mathbb{P}(z_v = z_u \mid \mathbf{z}^{-vu}, \mathbf{w}^{-vu}) &= \sum_{h=1}^{H^{-vu}+1} \left[\mathbb{P}(v \rightsquigarrow \mathbb{T}_{jh}^{-v}) \cdot \mathbb{P}(u \rightsquigarrow \mathbb{T}_{j'h}^{-u}) \right. \\
&\quad + \mathbb{P}(\tau_{new} \leftarrow h) \cdot \mathbb{P}(v \rightsquigarrow \tau_{new} \mid \tau_{new} \leftarrow h) \cdot \mathbb{P}(u \rightsquigarrow \mathbb{T}_{j'h}^{-u}) \\
&\quad + \mathbb{P}(\tau'_{new} \leftarrow h) \cdot \mathbb{P}(u \rightsquigarrow \tau'_{new} \mid \tau'_{new} \leftarrow h) \cdot \mathbb{P}(v \rightsquigarrow \mathbb{T}_{jh}^{-v}) \\
&\quad \left. + \mathbb{P}(\tau_{new} \leftarrow h, \tau'_{new} \leftarrow h) \cdot \mathbb{P}(v \rightsquigarrow \tau_{new} \mid \tau_{new} \leftarrow h) \cdot \mathbb{P}(u \rightsquigarrow \tau'_{new} \mid \tau'_{new} \leftarrow h) \right],
\end{aligned} \tag{17}$$

where

$$\begin{aligned}
\mathbb{P}(v \rightsquigarrow \mathbb{T}_{jh}^{-v}) &= \frac{\sum_{t=1}^{\ell_{jh}^{-vu}} \Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jh}^{-vu} + \mathbf{e}_t, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}{\Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j-1)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}, \\
\mathbb{P}(u \rightsquigarrow \mathbb{T}_{j'h}^{-u}) &= \frac{\sum_{t=1}^{\ell_{j'h}^{-vu}} \Phi_{\ell_{j'\cdot}^{-vu}, j'}^{(V_{j'})}(\mathbf{q}_{j'1}^{-vu}, \dots, \mathbf{q}_{j'h}^{-vu} + \mathbf{e}_t, \dots, \mathbf{q}_{j'H^{-vu}}^{-vu})}{\Phi_{\ell_{j'\cdot}^{-vu}, j'}^{(V_{j'}-1)}(\mathbf{q}_{j'1}^{-vu}, \dots, \mathbf{q}_{j'H^{-vu}}^{-vu})}, \\
\mathbb{P}(\tau_{new} \leftarrow h) &= \mathbb{P}(\tau'_{new} \leftarrow h) = \frac{\Phi_{H_h^{-vu}, 0}^{(|\ell^{-vu}|+1)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot h}^{-vu} + 1, \dots, \ell_{\cdot H_h^{-vu}}^{-vu})}{\Phi_{H^{-vu}, 0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot H^{-vu}}^{-vu})}, \\
\mathbb{P}(v \rightsquigarrow \tau_{new} \mid \tau_{new} \leftarrow h) &= \frac{\Phi_{\ell_{j\cdot}^{-vu}+1, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, (\mathbf{q}_{jh}^{-vu}, 1), \dots, \mathbf{q}_{jH_h^{-vu}}^{-vu})}{\Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j-1)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu})}, \\
\mathbb{P}(u \rightsquigarrow \tau'_{new} \mid \tau'_{new} \leftarrow h) &= \frac{\Phi_{\ell_{j'\cdot}^{-vu}+1, j'}^{(V_{j'})}(\mathbf{q}_{j'1}^{-vu}, \dots, (\mathbf{q}_{j'h}^{-vu}, 1), \dots, \mathbf{q}_{j'H_h^{-vu}}^{-vu})}{\Phi_{\ell_{j'\cdot}^{-vu}, j'}^{(V_{j'}-1)}(\mathbf{q}_{j'1}^{-vu}, \dots, \mathbf{q}_{j'H^{-vu}}^{-vu})}, \\
\mathbb{P}(\tau_{new} \leftarrow h, \tau'_{new} \leftarrow h) &= \frac{\Phi_{H_h^{-vu}, 0}^{(|\ell^{-vu}|+2)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot h}^{-vu} + 2, \dots, \ell_{\cdot H_h^{-vu}}^{-vu})}{\Phi_{H^{-vu}, 0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot H^{-vu}}^{-vu})},
\end{aligned} \tag{18}$$

for any $h \in [H^{-vu} + 1]$.

REMARK 2.7. As for Proposition 2.3, the components in (15) and (17) can be interpreted in terms of our metaphor. More specifically, in (15) we have the sum of: (i) the probability of both nodes v and u being allocated to an already-occupied subgroup, either the same or two different ones but still with same sociability profile (these two scenarios are accounted for by the two summands in (16)); (ii) the probability of creating a new subgroup with either both nodes assigned to that subgroup or one node assigned to the new subgroup and the other to a previously-occupied subgroup with the same sociability profile; (iii) the probability of being allocated to two new subgroups with equal sociability profile. In (17), since v and u are in different layers, we have the sum of: (i) the probability of being allocated to

two different already-occupied subgroups having the same sociability profile; (ii) the probabilities of creating, and occupying, a new subgroup in one of the two layers, while the other node is assigned to an already-occupied subgroup with the same sociability profile; (iii) the probability of being allocated to a new subgroup in both layers, with each new subgroup having the same sociability profile.

If one specializes the pEPPF to the H-DP case, the following novel result is obtained (an analogous one for the H-NSP can be found in the Supplementary Material).

COROLLARY 2.8. *Let $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \theta, \theta_0)$. Then, if both v and u are in the same layer j , we have*

$$\begin{aligned} \mathbb{P}(z_v = z_u \mid \mathbf{z}^{-vu}, \mathbf{w}^{-vu}) \\ = \frac{1}{(\theta + V_j - 2)(\theta + V_j - 1)} \left[\sum_{h=1}^{H^{-vu}} n_{jh}^{-vu} (n_{jh}^{-vu} + 1) + \theta \left(1 + \frac{2}{\theta_0 + |\ell^{-vu}|} \sum_{h=1}^{H^{-vu}} n_{jh}^{-vu} \ell_{.h}^{-vu} \right) + \right. \\ \left. + \frac{\theta^2}{(\theta_0 + |\ell^{-vu}|)(\theta_0 + |\ell^{-vu}| + 1)} \left(\sum_{h=1}^{H^{-vu}} \ell_{.h}^{-vu} (\ell_{.h}^{-vu} + 1) + \theta_0 \right) \right], \end{aligned} \quad (19)$$

whereas, if node v is in layer j and node u is in layer j' , with $j \neq j'$, it follows that

$$\begin{aligned} \mathbb{P}(z_v = z_u \mid \mathbf{z}^{-vu}, \mathbf{w}^{-vu}) \\ = \frac{1}{(\theta + V_j - 1)(\theta + V_{j'} - 1)} \left[\sum_{h=1}^{H^{-vu}} n_{jh}^{-vu} n_{j'h}^{-vu} + \frac{\theta}{\theta_0 + |\ell^{-vu}|} \sum_{h=1}^{H^{-vu}} \ell_{.h}^{-vu} (n_{jh}^{-vu} + n_{j'h}^{-vu}) \right. \\ \left. + \frac{\theta^2}{(\theta_0 + |\ell^{-vu}|)(\theta_0 + |\ell^{-vu}| + 1)} \left(\sum_{h=1}^{H^{-vu}} \ell_{.h}^{-vu} (\ell_{.h}^{-vu} + 1) + \theta_0 \right) \right]. \end{aligned} \quad (20)$$

The previous expressions are also useful for prior elicitation. According to our extended notion of homophily one expects the probability of a node being assigned to a sociability profile (*i.e.*, a group) already present in its layer to be larger than that of being allocated to a sociability profile new to its layer. As it can be argued from, *e.g.*, (14), the probability of allocating v to a group comprising nodes from its layer j or in a cluster having only nodes from other layers depends on the proportion $p_v = (1/|\ell^{-v}|) \sum_{h \in \mathbf{H}_j^{-v}} \ell_{.h}^{-v}$ where \mathbf{H}_j^{-v} is the set of unique sociability profiles for the nodes in layer j upon removing node v . In fact, allocation of nodes to within-layer groups are favored. These occur by either assigning a node to an already-occupied subgroup or by creating a new one with a sociability profile already present in the layer; conversely, only the latter option is possible to cluster strictly across layers. Nonetheless, if sociability profiles new to layer j are popular (in the sense of several subgroups in other

layers displaying them), then p_v , which measures the popularity of the sociability profiles already observed in v 's layer, is reduced. Consequently, clustering strictly across layers becomes increasingly probable. Albeit possible in some networks, this situation points in an opposite direction relative to the concept of homophily and, hence, it is natural to elicit priors which exclude this possibility. Leveraging the previously-derived results, Proposition 2.9 provides conditions on the H-DP hyperparameters which eliminate the dependence on the proportion p_v , thus enforcing the general notion of homophily. As clarified in the Supplementary Material, an analogous result holds for the H-NSP case.

PROPOSITION 2.9. *Let $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \theta, \theta_0)$. Then, for any generic node v in layer j , we have*

$$\theta \leq (V_j - 1) (\theta_0/|\ell^{-v}| + 1) \implies \mathbb{P}(z_v \in \mathbf{H}_j^{-v} | \mathbf{z}^{-v}, \mathbf{w}^{-v}) \geq \mathbb{P}(z_v \in \mathbf{H}^{-v} \setminus \mathbf{H}_j^{-v} | \mathbf{z}^{-v}, \mathbf{w}^{-v}), \quad (21)$$

for every $j \in [d]$, where \mathbf{H}^{-v} is the set of already-observed sociability profiles in all layers. Moreover

$$\theta \leq V_j - 1 \implies \mathbb{P}(z_v \in \mathbf{H}_j^{-v} | \mathbf{z}^{-v}, \mathbf{w}^{-v}) \geq \mathbb{P}(z_v \notin \mathbf{H}_j^{-v} | \mathbf{z}^{-v}, \mathbf{w}^{-v}), \quad (22)$$

for every node $v \in [V]$ and corresponding layer $j \in [d]$.

REMARK 2.10. Both conditions in Proposition 2.9 force the grouping structure to be more adherent, a priori, to the layer division, by making within-layer clusters more probable than across-layer ones. This generalized notion of homophily is sensible, since one would naturally expect that, in practice, it is more likely to share connectivity behaviour (*i.e.*, stochastic equivalence) within layers than across layers. Condition (22) implies (21), since $\{z_v \notin \mathbf{H}_j^{-v}\}$ includes z_v being assigned to a new sociability profile (*i.e.*, a new cluster). For large enough networks (or in large enough layers), (22) is always attained, whereas for smaller networks (or in layers with few nodes), our results have the merit to suggest that this key property can be still enforced by tuning the prior parameters in such a way that (22) is satisfied.

3 Inference and prediction

The pEx-SBM allows for tractable posterior inference on \mathbf{z} given \mathbf{Y} , and prediction for the allocation and edges of future nodes. In Section 3.1 we accomplish the former via a tractable collapsed Gibbs sampler which exploits the urn scheme of the pExp prior in (7). Innovative predictive schemes that leverage the projectivity properties and analytic results for the proposed pExp prior are derived in Section 3.2.

3.1 Collapsed Gibbs sampler

Unlike classical MCMC strategies for single-layer SBMs (*e.g.*, [Schmidt and Morup, 2013](#)), our pExP prior in the pEx-SBM relies on a two-level grouping structure which first assigns nodes to layer-specific subgroups and then clusters these subgroups w.r.t. common sociability profiles that form the final grouping structure in \mathbf{z} ; see the generative process for the auxiliary node attributes \mathbf{X} in Section 2.2.1. To this end, rather than devising an MCMC scheme which targets the posterior $p(\mathbf{z} \mid \mathbf{Y})$, we propose a Gibbs sampler targeting the augmented posterior $p(\mathbf{z}, \mathbf{w} \mid \mathbf{Y})$. This perspective facilitates the derivation of tractable full-conditional distributions $p(z_v, w_v \mid \mathbf{Y}, \mathbf{z}^{-v}, \mathbf{w}^{-v})$ for each node $v \in [V]$, while still allowing inference on $p(\mathbf{z} \mid \mathbf{Y})$ by simply retaining only the samples for \mathbf{z} . Neither slice-sampling steps nor truncations are required for inferring the total number of occupied groups, in contrast to, *e.g.*, [Amini et al. \(2024\)](#). Moreover, our Gibbs sampler covers the whole H-NRMI class.

By noting that (1) implies $(\mathbf{Y} \mid \mathbf{z}) \perp\!\!\!\perp \mathbf{w}$, a direct application of the Bayes rule yields the following full-conditional distributions for the generic node v

$$\mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{Y}, \mathbf{z}^{-v}, \mathbf{w}^{-v}) \propto \mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) \frac{p(\mathbf{Y} \mid z_v = h, \mathbf{z}^{-v})}{p(\mathbf{Y}^{-v} \mid \mathbf{z}^{-v})}, \quad (23)$$

for any $h \in [H^{-v} + 1]$ and $\tau \in [\ell_j^{-v} + 1]$, where j is the layer of node v , whereas \mathbf{Y}^{-v} denotes the $(V - 1) \times (V - 1)$ *supra*-adjacency matrix without the rows and columns corresponding to node v . Recalling *e.g.*, [Schmidt and Morup \(2013\)](#) and [Legramanti et al. \(2022\)](#), under the collapsed likelihood in (1) the term $p(\mathbf{Y} \mid z_v = h, \mathbf{z}^{-v})/p(\mathbf{Y}^{-v} \mid \mathbf{z}^{-v})$ in (23) can be evaluated in closed-form as

$$\frac{p(\mathbf{Y} \mid z_v = h, \mathbf{z}^{-v})}{p(\mathbf{Y}^{-v} \mid \mathbf{z}^{-v})} = \prod_{h'=1}^{H^{-v}} \frac{B(a + m_{hh'}^{-v} + r_{vh'}, b + \bar{m}_{hh'}^{-v} + \bar{r}_{vh'})}{B(a + m_{hh'}^{-v}, b + \bar{m}_{hh'}^{-v})}, \quad (24)$$

where $m_{hh'}^{-v}$ and $\bar{m}_{hh'}^{-v}$ are the number of edges and non-edges between groups h and h' , after excluding node v , while $r_{vh'}$ and $\bar{r}_{vh'}$ denote the number edges and non-edges among v and the nodes in h' . As for the prior factor, a closed-form and tractable expression for $\mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v})$ in (23) is given in Proposition 3.1.

PROPOSITION 3.1. *Let \mathbf{z} be a random allocation vector such that $\mathbf{z} \sim \text{pExP}(\mathbf{V}; \rho, \rho_0, c, c_0)$. Then for any node v in layer j we have the following:*

(1) If $h \in \mathbf{H}_j^{-v}$ and $t \in [\ell_{jh}^{-v}]$:

$$\mathbb{P}(z_v = h, w_v = \tau_{jht}^{-v} \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \frac{\Phi_{\ell_{j\cdot}^{-v}, j}^{(V_j)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jh}^{-v} + \mathbf{e}_t, \dots, \mathbf{q}_{jH^{-v}}^{-v})}{\Phi_{\ell_{j\cdot}^{-v}, j}^{(V_j-1)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jH^{-v}}^{-v})}, \quad (25)$$

where τ_{jht}^{-v} is the t -th entry of the vector $\boldsymbol{\tau}_{jh}^{-v}$ obtained ordering the set \mathbb{T}_{jh}^{-v} in (11).

(2) If $\tau = \ell_{j\cdot}^{-v} + 1$ and $h \in [H^{-v} + 1]$:

$$\begin{aligned} \mathbb{P}(z_v = h, w_v = \ell_{j\cdot}^{-v} + 1 \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) \\ = \frac{\Phi_{H_h^{-v}, 0}^{(|\ell^{-v}|+1)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot h}^{-v} + 1, \dots, \ell_{\cdot H^{-v}}^{-v}) \Phi_{\ell_{j\cdot}^{-v}+1, j}^{(V_j)}(\mathbf{q}_{j1}^{-v}, \dots, (\mathbf{q}_{jh}^{-v}, 1), \dots, \mathbf{q}_{jH_h^{-v}}^{-v})}{\Phi_{H_h^{-v}, 0}^{(|\ell^{-v}|)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot H^{-v}}^{-v}) \Phi_{\ell_{j\cdot}^{-v}, j}^{(V_j-1)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jH^{-v}}^{-v})}, \end{aligned} \quad (26)$$

where $H_h^{-v} = H^{-v} \vee h$.

(3) For any other choice of $(h, \tau) \in [H^{-v} + 1] \times [\ell_{j\cdot}^{-v} + 1]$ the probability is 0.

Corollary 3.2 covers the H-DP case, whereas expressions for the H-NSP case are provided in the Supplementary Material.

COROLLARY 3.2. If $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \theta, \theta_0)$, then, for any generic node v in layer j , we have

$$\begin{aligned} \mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) \\ = \mathbb{1}_{\{\tau = \tau_{jht}^{-v}\}} \frac{q_{jht}^{-v}}{\theta + V_j - 1} + \mathbb{1}_{\{\tau = \ell_{j\cdot}^{-v} + 1\}} \frac{\theta}{\theta + V_j - 1} \left[\frac{\ell_{\cdot h}^{-v}}{\theta_0 + |\ell^{-v}|} + \mathbb{1}_{\{h = H^{-v} + 1\}} \frac{\theta_0}{\theta_0 + |\ell^{-v}|} \right], \end{aligned} \quad (27)$$

for any $h \in [H^{-v} + 1]$, $\tau \in [\ell_{j\cdot}^{-v} + 1]$, where $\boldsymbol{\tau}_{jh}^{-v}$ is defined as in Proposition 3.1.

Another key result, which can be deduced from equations (25)–(26), is the following

COROLLARY 3.3. Let \mathbf{z} be a random allocation vector distributed as in (7), and indicate by \mathbf{w} the corresponding subgroup allocation vector. Then $(\boldsymbol{\ell}, \mathbf{q})$ is a predictive sufficient statistics for (\mathbf{z}, \mathbf{w}) , i.e. $\mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \mathbb{P}(z_v = h, w_v = \tau \mid \boldsymbol{\ell}^{-v}, \mathbf{q}^{-v})$, for any $h \in [H^{-v} + 1]$, $\tau \in [\ell_{j\cdot}^{-v} + 1]$ and node $v \in [V]$ in layer $j \in [d]$.

REMARK 3.4. The previous result can be further refined in the H-DP case by stating that the vector of the column-wise sums $(\ell_{\cdot 1}, \dots, \ell_{\cdot H^{-v}})$ of the matrix $\boldsymbol{\ell}$, i.e., the frequencies of subgroups assigned to each already-observed sociability profile in the whole network, and the array \mathbf{q}_j , i.e., the frequencies of nodes assigned to each already-observed subgroup for each already-observed sociability profile in layer j , are predictive sufficient statistics.

Corollary 3.3, combined with (23)–(26), facilitates the implementation of a tractable collapsed Gibbs sampler, which iteratively simulates values from the full-conditional allocation probabilities $p(z_v, w_v \mid \mathbf{Y}, \mathbf{z}^{-v}, \mathbf{w}^{-v})$ of each node $v \in [V]$ via simple updating of the quantities (ℓ, \mathbf{q}) for the prior component in (25)–(26), and $(\mathbf{m}^{-v}, \bar{\mathbf{m}}^{-v}, \mathbf{r}_v, \bar{\mathbf{r}}_v)$ for the collapsed likelihood term in (24). The latter quantities denote, respectively, the two matrices counting the edges and non-edges between groups, after excluding node v , and the vectors with the number of edges and non-edges between v and the nodes in the different groups. The pseudo-code for the proposed computational strategy is outlined in Algorithm 1.

Although Step 3 of Algorithm 1 requires the computation of an $(H^{-v} + 1) \times (\ell_j^{-v} + 1)$ -dimensional probability table, this matrix is highly sparse. In fact, recalling Proposition 3.1, the probability of being assigned h at subgroup τ is non-zero only if τ is a previously-occupied subgroup with profile h or if τ is a new subgroup. Thus, each column of the probability table, except the last, has a single non-zero entry.

Leveraging on the samples of \mathbf{z} from Algorithm 1, we conduct posterior inference on the nodes' grouping structures via the variation of information (VI) approach set forth in Wade and Ghahramani (2018) for Bayesian clustering. The VI defines distances among partitions through a comparison of individual and joint entropies (Meilă, 2007), thus providing a metric which allows for point estimation and uncertainty quantification directly in the space of partitions. In this framework, a point estimate for \mathbf{z} is defined as $\hat{\mathbf{z}} = \operatorname{argmin}_{\mathbf{z}'} \mathbb{E}[\text{VI}(\mathbf{z}, \mathbf{z}') \mid \mathbf{Y}]$, whereas a $(1 - \alpha)$ -credible ball around $\hat{\mathbf{z}}$ is obtained by collecting all partitions with VI distance from $\hat{\mathbf{z}}$ less than a pre-specified threshold defined to ensure the

Algorithm 1 Gibbs sampler for pEx-SBM

Initialize coherently \mathbf{z} and \mathbf{w} ▷ Default: V different sociability profiles and subgroups (one for each node)
for $s = 1, \dots, n_{\text{iter}}$ (n_{iter} : number of Gibbs iterations) **do**
 for $v = 1 \dots V$ (V : total number of nodes in the network) **do**
 1. $j \leftarrow$ layer of node v
 2. remove node v
 2.1 reorder the labels in \mathbf{z}^{-v} and \mathbf{w}^{-v} so that only sociability profiles $h \in [H^{-v} + 1]$
 and subgroups $\tau \in [\ell_j^{-v} + 1]$ in layer j are non-empty
 2.2 compute the quantities $\ell^{-v}, \mathbf{q}^{-v}, \mathbf{m}^{-v}, \bar{\mathbf{m}}^{-v}, \mathbf{r}_v, \bar{\mathbf{r}}_v$
 for $h \in [H^{-v} + 1]$ and $\tau \in [\ell_j^{-v} + 1]$ **do**
 3. compute the full-conditional probabilities $\mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{Y}, \mathbf{z}^{-v}, \mathbf{w}^{-v})$
 (up to a proportionality constant) by direct application of equations (23)–(26)
 end for
 4. sample (z_v, w_v) from the bivariate discrete random variable with joint probabilities
 obtained by normalizing those computed in 3.
 5. update $\mathbf{m}, \ell, \mathbf{q}$ based on the new allocation vectors (z_v, \mathbf{z}^{-v}) and (w_v, \mathbf{w}^{-v})
 end for
 6. save the sampled vector $\mathbf{z}^{(s)}$ and $\mathbf{w}^{(s)}$
end for

ball to contain at least $1-\alpha$ posterior mass, while being of minimum size. These tasks can be performed via the R library `mcclust.ext` (Wade and Ghahramani, 2018), which requires the calculation of the $V \times V$ *posterior similarity (or co-clustering)* matrix \mathbf{C} , whose generic entry c_{vu} yields an estimate of $\mathbb{P}(z_v = z_u \mid \mathbf{Y})$ by computing the relative frequency of MCMC samples in which $z_v^{(s)} = z_u^{(s)}$.

Model assessment and comparison is performed through the WAIC information criterion, that has also a direct connection with Bayesian leave-one-out cross-validation (Watanabe, 2010; Gelman et al., 2014). This criterion is further useful to select, in a data-driven manner, specific priors in the H-NRMI class and tune the corresponding parameters. Alternatively, prior elicitation can proceed by studying the expected number of non-empty clusters and the growth of H as a function of the network size. For node-colored networks with $V_1 = \dots V_d = V/d$, such a growth is $\mathcal{O}(\log \log V)$ under H-DP(θ, θ_0) and $\mathcal{O}(V^{\sigma\sigma_0})$ for H-NSP(σ, σ_0) (e.g., Camerlenghi et al., 2019). Instead, within each layer j , this growth is $\mathcal{O}(\log V_j)$ and $\mathcal{O}(V_j^\sigma)$, respectively. A final option is to specify hyperpriors for the parameters of the selected H-NRMI. This requires simply adding a step in Algorithm 1 to sample from the full-conditionals of these parameters. As clarified in the Supplementary Material, in the H-DP case with gamma hyperpriors for θ and θ_0 , closed-form and tractable full conditionals can be straightforwardly derived by extending results of Escobar and West (1995) from the DP to the H-DP.

3.2 k -step-ahead prediction

The main focus of predictive strategies for SBMs has been on prediction of edges among the observed nodes, or forecasting the group allocations for new nodes given \mathbf{Y} and the observed connections between these new nodes and the previously observed ones (e.g., Lee and Wilkinson, 2019). These strategies fail to address the more challenging and realistic problem of predicting the allocations of new nodes, conditioned only on \mathbf{Y} , and, most importantly, the edges between these new nodes and the previously-observed ones. In fact, in practice one expects that for k new incoming nodes only the corresponding layers are observed, and no information on the associated edges and group allocations is available. By inheriting the Kolmogorov consistency from the pExp prior, the pExp-SBM opens the avenues for addressing these predictive tasks through innovative methods not yet available in current literature.

In particular, the sequential construction of the nonparametric priors employed in pEx-SBM allows for principled prediction of an arbitrary number k of new nodes, from any layer.

Let us first focus on predicting the allocations $\mathbf{z}_{\text{new}} = (z_{V+1}, \dots, z_{V+k})$ and suballocations $\mathbf{w}_{\text{new}} = (w_{V+1}, \dots, w_{V+k})$ for the k incoming nodes conditioned only on the network \mathbf{Y} among the previously-observed nodes. Since the new allocations are conditionally independent from \mathbf{Y} , given \mathbf{z} and \mathbf{w} , for any $v_{\text{new}} = V + 1, \dots, V + k$ and $u = 1, \dots, v_{\text{new}} - 1$ we have

$$\mathbb{P}(z_{v_{\text{new}}} = z_u \mid \mathbf{Y}) = \int \mathbb{P}(z_{v_{\text{new}}} = z_u \mid \mathbf{Y}, \mathbf{z}, \mathbf{w}) \mathcal{L}_{\mathbf{Y}}(d\mathbf{z}, d\mathbf{w}) = \int \mathbb{P}(z_{v_{\text{new}}} = z_u \mid \mathbf{z}, \mathbf{w}) \mathcal{L}_{\mathbf{Y}}(d\mathbf{z}, d\mathbf{w}), \quad (28)$$

where $\mathcal{L}_{\mathbf{Y}}$ denotes the joint posterior law of (\mathbf{z}, \mathbf{w}) .

By combining (28) with the results in Proposition 2.3 and Theorem 2.6, predictive co-clustering probabilities can be estimated via Monte Carlo as $\sum_{s=n_{\text{burn}}+1}^{n_{\text{iter}}} \mathbb{P}(z_{v_{\text{new}}} = z_u \mid \mathbf{z}^{(s)}, \mathbf{w}^{(s)}) / n_{\text{tot}}$ with n_{burn} denoting the number of discarded burn-in samples and $n_{\text{tot}} = n_{\text{iter}} - n_{\text{burn}}$ the number of subsequent MCMC draws from Algorithm 1 used for averaging. When u is a new node, $\mathbb{P}(z_{v_{\text{new}}} = z_u \mid \mathbf{z}^{(s)}, \mathbf{w}^{(s)})$ can be directly evaluated by Theorem 2.6. Conversely, when u denotes a previously-observed node, then $\mathbb{P}(z_{v_{\text{new}}} = z_u \mid \mathbf{z}^{(s)}, \mathbf{w}^{(s)}) = \mathbb{P}(z_{v_{\text{new}}} = h \mid \mathbf{z}^{-v(s)}, z_u^{(s)} = h, \mathbf{w}^{(s)})$, which is computed via Proposition 2.3.

The above procedure yields an estimate of predictive co-clustering probabilities among the k new nodes as well as between such new nodes and the previously-observed ones. Combining these with the quantities computed for the in-sample nodes from the MCMC yields an augmented $(V + k) \times (V + k)$ posterior similarity matrix \mathbf{C}_{aug} which quantifies uncertainty for both in-sample and predictive co-clustering probabilities. A VI-point estimate for \mathbf{z}_{new} is then obtained via the R library `mcclust.ext`. In contrast to other approaches (Nowicki and Snijders, 2001; Legramanti et al., 2022), knowledge on the edges of the new incoming nodes is not required. Only information on layer allocation is employed.

We now focus on predicting the edges for these new incoming nodes. More specifically, the goal is to derive the predictive $p(\mathbf{Y}_{\text{new}} \mid \mathbf{Y})$, where \mathbf{Y}_{new} comprises the yet-unobserved edges $y_{v_{\text{new}}, u}$ with $v_{\text{new}} = V + 1, \dots, V + k$ and $u = 1, \dots, v_{\text{new}} - 1$. Let $\bar{\mathbf{z}} := (\mathbf{z}, \mathbf{z}_{\text{new}})$ be the complete allocation vector. Then

$$p(\mathbf{Y}_{\text{new}} \mid \mathbf{Y}) = \int p(\mathbf{Y}_{\text{new}} \mid \mathbf{Y}, \bar{\mathbf{z}}) \mathcal{L}_{\mathbf{Y}}(d\bar{\mathbf{z}}), \quad (29)$$

with $\mathcal{L}_{\mathbf{Y}}$ the posterior law of $\bar{\mathbf{z}}$. Notice that $p(\mathbf{Y} \mid \bar{\mathbf{z}}) = p(\mathbf{Y} \mid \mathbf{z})$, since $(\mathbf{Y} \mid \mathbf{z}) \perp\!\!\!\perp \mathbf{z}_{\text{new}}$. Hence, by Bayes'

rule and (1), $p(\mathbf{Y}_{\text{new}} \mid \mathbf{Y}, \bar{\mathbf{z}})$ can be expressed as

$$\begin{aligned}
p(\mathbf{Y}_{\text{new}} \mid \mathbf{Y}, \bar{\mathbf{z}}) &= \frac{p(\mathbf{Y}_{\text{new}}, \mathbf{Y} \mid \bar{\mathbf{z}})}{p(\mathbf{Y} \mid \mathbf{z})} \\
&= \prod_{h=1}^H \prod_{h'=1}^h \frac{B(a, b) B(a + m_{hh'} + m_{hh'}^{\text{new}}, b + \bar{m}_{hh'} + \bar{m}_{hh'}^{\text{new}})}{B(a + m_{hh'}, b + \bar{m}_{hh'}) B(a, b)} \prod_{h=H+1}^{H+H_k} \prod_{h'=1}^h \frac{B(a + m_{hh'}^{\text{new}}, b + \bar{m}_{hh'}^{\text{new}})}{B(a, b)}, \quad (30) \\
&= \prod_{h=1}^H \prod_{h'=1}^h \frac{B(a_{hh', \mathbf{Y}} + m_{hh'}^{\text{new}}, b_{hh', \mathbf{Y}} + \bar{m}_{hh'}^{\text{new}})}{B(a_{hh', \mathbf{Y}}, b_{hh', \mathbf{Y}})} \prod_{h=H+1}^{H+H_k} \prod_{h'=1}^h \frac{B(a + m_{hh'}^{\text{new}}, b + \bar{m}_{hh'}^{\text{new}})}{B(a, b)},
\end{aligned}$$

where $a_{hh', \mathbf{Y}} = a + m_{hh'}$ and $b_{hh', \mathbf{Y}} = b + \bar{m}_{hh'}$, while $m_{hh'}^{\text{new}}$ and $\bar{m}_{hh'}^{\text{new}}$ are the number of edges and non-edges among groups h and h' that involve at least a new node (if any) allocated to these clusters. In (30), the first factor corresponds to groups already occupied by the in-sample nodes, whereas the second refers to the H_k potential new clusters created by the k new incoming nodes.

The combination of (29)–(30) allows to evaluate the predictive probability of any configuration \mathbf{Y}_{new} via the Monte Carlo estimate $\sum_{s=n_{\text{burn}}+1}^{n_{\text{iter}}} p(\mathbf{Y}_{\text{new}} \mid \mathbf{Y}, \bar{\mathbf{z}}^{(s)}) / n_{\text{tot}}$, where $p(\mathbf{Y}_{\text{new}} \mid \mathbf{Y}, \bar{\mathbf{z}}^{(s)})$ coincides with (30) evaluated at the MCMC samples $\mathbf{z}^{(s)}$ and $\mathbf{z}_{\text{new}}^{(s)}$ from $p(\bar{\mathbf{z}} \mid \mathbf{Y})$. To obtain $\bar{\mathbf{z}}^{(s)}$, note that $p(\bar{\mathbf{z}} \mid \mathbf{Y}) = p(\mathbf{z}_{\text{new}} \mid \mathbf{z}, \mathbf{Y}) p(\mathbf{z} \mid \mathbf{Y})$, where, by the pEx-SBM formulation, $p(\mathbf{z}_{\text{new}} \mid \mathbf{z}, \mathbf{Y}) = p(\mathbf{z}_{\text{new}} \mid \mathbf{z})$ and the samples $\mathbf{z}^{(s)}$ from $p(\mathbf{z} \mid \mathbf{Y})$ are already available from the output of Algorithm 1. Therefore, for each $\mathbf{z}^{(s)}$ it suffices to simulate $\mathbf{z}_{\text{new}}^{(s)}$ from $p(\mathbf{z}_{\text{new}} \mid \mathbf{z}^{(s)})$ or, alternatively, from the augmented posterior $p(\mathbf{z}_{\text{new}}, \mathbf{w}_{\text{new}} \mid \mathbf{z}^{(s)}, \mathbf{w}^{(s)})$, similarly to Algorithm 1. In view of our construction, the entries in $\mathbf{z}_{\text{new}}^{(s)}$ can be coherently generated in a sequential manner. This is achieved via the joint urn scheme in (25)–(26) from $p(z_{v_{\text{new}}}, w_{v_{\text{new}}} \mid \mathbf{z}^{(s)}, \mathbf{w}^{(s)}, \mathbf{z}_{v_{\text{new}}-1}^{(s)}, \mathbf{w}_{v_{\text{new}}-1}^{(s)})$ for $v_{\text{new}} = V + 1, \dots, V + k$, where $\mathbf{z}_{v_{\text{new}}-1}^{(s)}$ and $\mathbf{w}_{v_{\text{new}}-1}^{(s)}$ are the simulated group and subgroup allocations for the new incoming nodes $V + 1, \dots, v_{\text{new}} - 1$.

The above strategy allows to evaluate predictive probabilities for any possible configuration of new edges and non-edges in \mathbf{Y}_{new} . However, in practice, the set of all configurations is excessively large and difficult to summarize in a meaningful manner. Hence, it is often convenient to complement a joint analysis with the predictive summaries for each new edge $y_{v_{\text{new}}, u}$ with $v_{\text{new}} = V + 1, \dots, V + k$ and $u = 1, \dots, v_{\text{new}} - 1$, namely $\mathbb{P}(y_{v_{\text{new}}, u} = 1 \mid \mathbf{Y}) = 1 - \mathbb{P}(y_{v_{\text{new}}, u} = 0 \mid \mathbf{Y})$.

Let $\bar{z}_v := \bar{z}_{v_{\text{new}}}$ be the allocation of v_{new} . Moreover, denote by $\mathcal{L}_{\mathbf{Y}}$ the posterior law of $\bar{\mathbf{z}}$, and by $\mathcal{L}_{\mathbf{Y}, \bar{\mathbf{z}}}$ the conditional law of the block-probability $\psi_{\bar{z}_v, \bar{z}_u}$ between clusters \bar{z}_v and \bar{z}_u given the observed

network \mathbf{Y} and all allocations $\bar{\mathbf{z}}$. Then, $\mathbb{P}(y_{v_{\text{new}},u} = 1 \mid \mathbf{Y})$ can be equivalently expressed as

$$\begin{aligned} & \mathbb{P}(y_{v_{\text{new}},u} = 1 \mid \mathbf{Y}) \\ &= \int \int \mathbb{E}[y_{v_{\text{new}},u} \mid \mathbf{Y}, \bar{\mathbf{z}}, \psi_{\bar{z}_v, \bar{z}_u}] \mathcal{L}_{\mathbf{Y}, \bar{\mathbf{z}}}(\mathrm{d}\psi_{\bar{z}_v, \bar{z}_u}) \mathcal{L}_{\mathbf{Y}}(\mathrm{d}\bar{\mathbf{z}}) = \int \frac{a + m_{\bar{z}_v, \bar{z}_u}}{a + b + m_{\bar{z}_v, \bar{z}_u} + \bar{m}_{\bar{z}_v, \bar{z}_u}} \mathcal{L}_{\mathbf{Y}}(\mathrm{d}\bar{\mathbf{z}}), \end{aligned} \quad (31)$$

after noticing that, due to beta-binomial conjugacy, $(\psi_{\bar{z}_v, \bar{z}_u} \mid \mathbf{Y}, \mathbf{z}, \mathbf{z}_{\text{new}}) \sim \text{beta}(a + m_{\bar{z}_v, \bar{z}_u}, b + \bar{m}_{\bar{z}_v, \bar{z}_u})$ where $m_{\bar{z}_v, \bar{z}_u}$ and $\bar{m}_{\bar{z}_v, \bar{z}_u}$ are the total number of edges and non-edges between pairs of in-sample nodes allocated to \bar{z}_v and \bar{z}_u . Therefore the last integral in (31) can be evaluated again via Monte Carlo using the samples $(\mathbf{z}^{(s)}, \mathbf{z}_{\text{new}}^{(s)})$ generated to compute (29). This yields the estimate of $\mathbb{P}(y_{v_{\text{new}},u} = 1 \mid \mathbf{Y})$, for each $v_{\text{new}} = V + 1, \dots, V + k$ and $u = 1, \dots, v_{\text{new}} - 1$, defined as $\sum_{s=n_{\text{burn}}+1}^{n_{\text{iter}}} [(a + m_{\bar{z}_v, \bar{z}_u}^{(s)}) / (a + m_{\bar{z}_v, \bar{z}_u}^{(s)} + b + \bar{m}_{\bar{z}_v, \bar{z}_u}^{(s)})] / n_{\text{tot}}$ where $m_{v_{\text{new}},u}^{(s)}$ and $\bar{m}_{v_{\text{new}},u}^{(s)}$ coincide with $m_{v_{\text{new}},u}$ and $\bar{m}_{v_{\text{new}},u}$ evaluated at $(\mathbf{z}^{(s)}, \mathbf{z}_{\text{new}}^{(s)})$.

4 Simulation Studies

To assess performance of pEx-SBM and illustrate its merits compared to both state-of-the-art (Zhang et al., 2016; Binkiewicz et al., 2017; Legramanti et al., 2022) and routinely-implemented (Blondel et al., 2008; Côme et al., 2021) competitors, we simulate complex binary undirected networks among $V = 80$ nodes divided in $d = 4$ layers, of size $V_1 = V_2 = 30$, $V_3 = 15$ and $V_4 = 5$, whose block interactions mimic those expected for the criminal network in Section 5. As shown in Figure 2, the first three layers comprise two within-layer groups, and one across-layer cluster, identified by the white color, which contains nodes from layers 1, 2 and 3. Recalling the criminal network in Figure 1, the larger groups in each layer can be thought of as *locale*-specific affiliates who mainly interact with peers in the same cluster and with an additional low-sized group of higher-level supervisors that administer activities in each *locale* and report to an across-layer group of bosses from the different *locali*. These individuals are the only ones entitled to establish dense connections with nodes in the fourth layer, which consists of a single group of bosses at the top of the criminal organization that connect with few, yet central, *locale*-specific bosses. As illustrated in Figure 2 the edge probabilities among the nodes in these $H_0 = 8$ groups are defined to be consistent with this pyramidal block-connectivity architecture, while accounting for two different scenarios in which separation among blocks is either more (Scenario 1) or less (Scenario 2) pronounced.

Based on the above probability matrices, we simulate the edges in the *supra*-adjacency matrix \mathbf{Y}

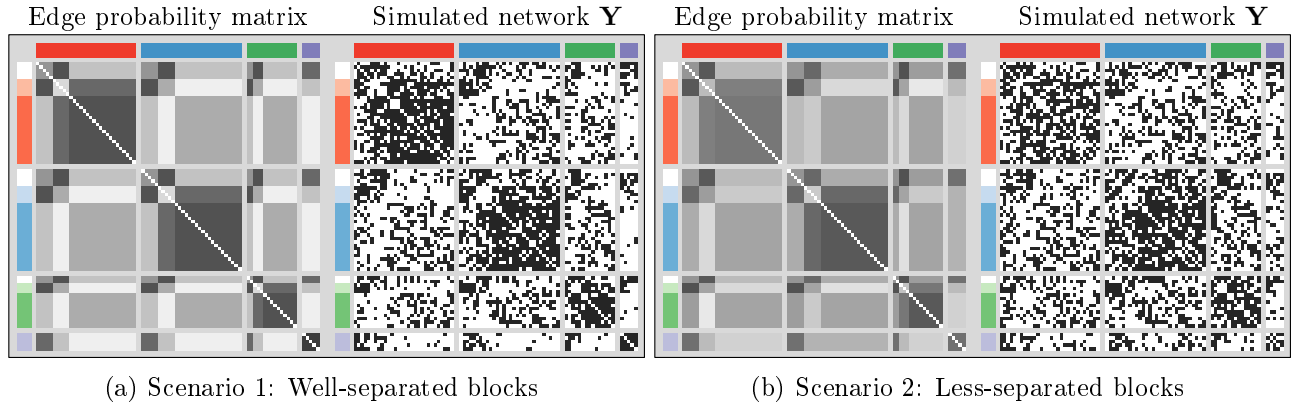


Figure 2: For scenario 1 (panel (a)) and 2 (panel (b)), graphical representation of the true underlying edge probability matrix (left), and of one example of *supra*-adjacency matrix simulated from this edge probability matrix (right). The color annotation of columns refers to layers division, while the row colors display the true underlying grouping structure. In the edge probability matrices, the color of each entry ranges from white to black as the corresponding edge probability goes from 0 to 1. In the *supra*-adjacency matrices, black and white entries indicate edges and non-edges, respectively.

from independent Bernoulli variables for each scenario (see Figure 2), and then perform the analysis under three key examples of pEx-SBMs in ten replicated experiments relying on different \mathbf{Y} simulated from the edge probability matrices in Figure 2. Computation and inference for both scenarios resort to the algorithms and methods in Section 3.1, for three pEx-SBMs with the customary $\text{beta}(1, 1)$ (*i.e.*, uniform) prior for the block-probabilities and pExP priors on \mathbf{z} induced by an H-DP($\theta = 0.5, \theta_0 = 4$), an H-NSP($\sigma = 0.2, \sigma_0 = 0.8$), and H-DP with $\text{gamma}(5, 10)$ and $\text{gamma}(12, 3)$ hyperpriors for θ and θ_0 , respectively. These settings serve to assess sensitivity and robustness with respect to misscalibrated priors. In fact, these choices enforces a prior expected number of groups of ≈ 5 for all the three pEx-SBM examples, lower than the true $H_0 = 8$. Posterior inference relies on 8,000 samples of \mathbf{z} from Algorithm 1, after a conservative burn-in of 2,000. As confirmed by the traceplots for the logarithm of the likelihood in (1) (see the Supplementary Material), in practice, a few MCMC iterations suffice for convergence of the collapsed Gibbs samplers. A basic R implementation of Algorithm 1 requires ≈ 1.5 minutes to simulate 10,000 values of \mathbf{z} on a MacBook Air (M1, 2020), CPU 8-core and 8GB RAM.

As shown in the first three lines of Table 1, all priors within the proposed pEx-SBM class learn the true underlying block structures in both Scenario 1 and 2 with high accuracy, across replicated experiments. In Scenario 2 the reduced group separation leads, as expected, to a slight performance deterioration. However, even in this challenging context, pEx-SBMs are still accurate and systematically outperform relevant Bayesian and non-Bayesian competitors (see Table 1). These results clarify that, regardless of the specific pExP prior considered, as long as such a prior incorporates the most suitable

Table 1: Performance comparison between three key examples of pEx-SBMs and relevant competitors for Scenarios 1–2 under several metrics: VI distance $\text{VI}(\hat{\mathbf{z}}, \mathbf{z}_0)$ of the estimated $\hat{\mathbf{z}}$ from the truth \mathbf{z}_0 (number of unique clusters in $\hat{\mathbf{z}}$ inside brackets), posterior median of the number of groups H (interquartile range inside brackets), posterior mean $\mathbb{E}[\text{VI}(\mathbf{z}, \mathbf{z}_0) \mid \mathbf{Y}]$ of the VI distance from the true \mathbf{z}_0 , distance $\text{VI}(\hat{\mathbf{z}}, \mathbf{z}_b)$ between the estimated partition $\hat{\mathbf{z}}$ and the 95% credible bound \mathbf{z}_b , WAIC. Only pEx-SBMs and Supervised ESBM rely on a Bayesian approach and, hence, posterior quantities and WAIC are only available for them. For JCDC and SBM (**greed**) a favorable implementation is considered with routines’ initializations at the true number of groups $H_0 = 8$. For CASC, the number of clusters is set exactly equal to $H_0 = 8$. Results are averaged over 10 replicated experiments. Lower VI distances and WAIC indicate more accurate performance.

SCENARIO	$\text{VI}(\hat{\mathbf{z}}, \mathbf{z}_0)$ (\hat{H})		$\text{med}(H \mid \mathbf{Y})$		$\mathbb{E}[\text{VI}(\mathbf{z}, \mathbf{z}_0) \mid \mathbf{Y}]$		$\text{VI}(\hat{\mathbf{z}}, \mathbf{z}_b)$		WAIC	
	1	2	1	2	1	2	1	2	1	2
pEx-SBM (H-DP)	0.00 (8.0)	0.45 (7.1)	8.0 [0.1]	7.5 [0.7]	0.02	0.62	0.12	0.82	3508	4046
pEx-SBM (H-NSP)	0.00 (8.0)	0.39 (7.3)	8.2 [0.9]	8.2 [1.5]	0.04	0.60	0.17	0.84	3510	4049
pEx-SBM (H-DP hyp)	0.00 (8.0)	0.43 (7.0)	8.0 [0.2]	7.6 [1.0]	0.03	0.63	0.14	0.85	3509	4045
Supervised ESBM (DP)	0.02 (8.1)	0.74 (6.3)	8.1 [0.7]	7.2 [1.6]	0.06	0.98	0.18	1.14	3512	4082
Louvain	1.21 (3.9)	2.21 (3.8)	—	—	—	—	—	—	—	—
SBM (greed)	0.06 (7.5)	2.18 (2.6)	—	—	—	—	—	—	—	—
JCDC ($w_n = 5$)	1.54 (7.7)	2.41 (8.0)	—	—	—	—	—	—	—	—
JCDC ($w_n = 1.5$)	1.45 (7.1)	1.58 (7.0)	—	—	—	—	—	—	—	—
CASC (rCASC)	0.83 (8.0)	1.49 (8.0)	—	—	—	—	—	—	—	—

notion of exchangeability for the multilayer network analyzed, remarkable gains can be achieved over competitors, further motivating our broad focus on the general H-NRMI class.

The improvements w.r.t. the Louvain algorithm for community detection (Blondel et al., 2008) and the **greed** implementation by Côme et al. (2021) of classical single-layer SBM highlights the importance of including layer information for the analysis of node-colored networks. This can be included via state-of-the-art attribute-assisted methods in Zhang et al. (2016) (JCDC with either $w_n = 5$ or $w_n = 1.5$), Binkiewicz et al. (2017) (CASC leveraging the optimized parameter tuning in the **rCASC** library and a favorable implementation setting the number of clusters exactly equal to $H_0 = 8$) and Legramanti et al. (2022) (supervised ESBM with DP prior inducing the same expected number of groups as for pEx-SBM). As illustrated in Table 1, both JCDC and CASC are not competitive with pEx-SBM. The former mainly searches for assortative block structures and hence lacks sufficient flexibility, whereas the latter inherits the difficulties of the underlying spectral clustering along with potential challenges in finding a proper balance between the exogenous layer division and the endogenous grouping structures in the network. The ESBM improves over JCDC and CASC, but it is still not competitive with pEx-SBM. Recall that ESBM employs exchangeable priors for \mathbf{z} supervised by layers in a way that does not yield a projective formulation and, hence, cannot induce a prediction rule as the one we derived for pEx-SBMs.

The empirical evidence in Figure 3 suggests that pEx-SBMs are also able to recover the true partition \mathbf{z}_0 as the network size V grows, in replicated studies, thereby hinting at the validity of frequentist




Figure 3: Empirical assessment of frequentist posterior consistency for H-DP, H-DP hyperprior, and H-NSP monitored via the boxplots for $VI(\hat{\mathbf{z}}, \mathbf{z}_0)$ from 10 replicated experiments on each network size from $V = 40$ to $V = 160$. As benchmark we consider CASC (Binkiewicz et al., 2017) for which consistency in estimating \mathbf{z}_0 has been proved theoretically.

posterior consistency for our proposed methods, even in the challenging Scenario 2. This finding is further strengthened by the comparison with CASC for which consistency has been proved in Binkiewicz et al. (2017), though under a number of assumptions that rule out several data-generating processes of direct interest in practice. Relaxing such assumptions while extending the results to the node-colored network setting motivates future theoretical studies on pEx-SBM posterior consistency; see Section 6.

As clarified in Section 3.2, pEx-SBM stands out also for its novel and principled k -step-ahead predictive strategies for both the edges and the allocations of future incoming nodes. By evaluating these strategies on 10 randomly-selected held-out nodes varying across the ten replicated studies in Scenario 1, yields an average mean squared error of 0.03 between the true and predicted edge probabilities, and an average of 2.8 misallocated nodes out of 10 for all the three pEx-SBM examples. These are remarkable results when considering that only layer information is employed in these predictive tasks.

Assessments on robustness to initialization, hyperparameter specification and inclusion of uninformative layers in the three pEx-SBM examples analyzed can be found in the Supplementary Material. Results suggest that all these three pEx-SBMs are robust. Among them, the H-DP with hyperpriors on θ and θ_0 stands out due to its ability of learning the parameters that control clustering properties together with the strength of layer information. As such, we will consider this prior in Section 5.

5 Application to Criminal Networks

We now showcase the potential of our pEx-SBM approach on real-world data and consider the challenging criminal network application hinted at in Figure 1. The raw data were retrieved from the judicial

acts of a law enforcement operation, named *Operazione Infinito* (e.g., [Calderoni et al., 2017](#)), that was conducted in Italy to disrupt a branch of the 'Ndrangheta Mafia operating in the Milan area.

Original data and judicial acts are available at <https://sites.google.com/site/ucinetsoftware/datasets/covert-networks>. Here we consider the pre-processed version studied by [Legramanti et al. \(2022\)](#), which contains information on the presence or absence of at least one monitored co-attendance to a 'Ndrangheta summit for each pair of the $V = 84$ registered criminals. Besides connectivity information, there is also knowledge on the role of each member (simple affiliate or boss) and on membership to the so-called *locali*, namely structural coordinated sub-units administering crime in specific territories. In analyzing this network, [Calderoni et al. \(2017\)](#) focus on detecting simple community structures, while [Legramanti et al. \(2022\)](#) identify modular architectures via an ESBM supervised with both role and locale affiliation. Nonetheless, obtaining reliable information on the roles of all criminals is often challenging, or even impossible, and, in general, is not available as prior information. Therefore, an important goal is that of deducing the roles of criminals from the inferred group structures.

Motivated by this remark, we only use *locali* to define the layers in the pEx-SBM, disregarding information on roles. Posterior inference is performed under the same settings of Section 4 leveraging a H-DP with $\text{gamma}(10, 2.5)$ and $\text{gamma}(5, 0.45)$ hyperpriors for the parameters θ and θ_0 , so to induce a prior expected number of groups of ≈ 15 . Since there are five layers with nodes possibly covering different roles, it is reasonable to expect, a priori, the total number of groups to be at least double or triple the number of layers. In fact, as illustrated in Figure 4, the minimum-VI point estimate $\hat{\mathbf{z}}$ of \mathbf{z} points toward $\hat{H} = 14$ groups, which is in line with a number of network dynamics reported in the judicial acts. Such an accurate characterization of the block patterns in the observed network for pEx-SBM is also confirmed by a WAIC of 1282, which improves the WAIC of 1299 achieved by the most competitive alternative in the empirical studies in Section 4, i.e., ESBM with DP prior and *locali* supervision.

As is apparent from Figure 4, pEx-SBM is able to disentangle core-periphery structures associated with boss-affiliate dynamics in each *locale*. In fact, even if the role attribute is not incorporated in the model, for 4 of the 5 *locali*, affiliates and bosses are grouped separately, hinted at a clear role separation in these *locali*. Note that each small group of bosses also contains few affiliates. This suggests more nu-




Figure 4: Graphical representation of the *Infinito* network along with the grouping structure estimated by the pEx-SBM. Node positions are obtained via force directed placement (Fruchterman and Reingold, 1991), whereas colors indicate the division in layers corresponding to *locali* affiliation. The node size is proportional to the corresponding betweenness, while the shape indicates affiliates (circles) and bosses (squares). Gray areas highlight the groups inferred by pEx-SBM.

anced roles beyond boss-affiliate separation, with some of the affiliates displaying connectivity behavior closer to those of bosses. This is further confirmed by the single-node group detected by pEx-SBM at the core of the network. Surprisingly, such a node is an affiliate that, however, displays a unique and central role in the connectivity architecture. According to the judicial acts, this node corresponds to a high-rank member who is in charge of coordinating the different *locali* and reporting to the leading Calabria-based families. Hence, its status appears to be even higher than that of *locali* bosses.

Across-layer groups are also inferred, as desired. These clusters reflect collaborative schemes between criminals in different *locali* or more nuanced dynamics within the network, such as attempts of some members to change affiliation. For instance, the inferred group of red affiliates also includes a member of the blue *locale* trying to create a new one by looking for affiliates in a different influence area.

A peculiar feature of this network is represented by the murder, during the investigations, of a high-rank member that destabilized the green *locale*. As shown in Figure 4, pEx-SBM estimate of \mathbf{z} unveils the consequences of this event, with the green *locale* being the most fragmented in both within- and across-layer groups. Some of its affiliates are allocated to the nodes' group of the blue *locale*. Moreover, even if a cluster of green bosses is detected, there is also a small group of three green affiliates seemingly forming a core-periphery structure with the more peripheral members of this green *locale*. Interestingly, these three nodes are among those more closely involved in the murder, according to the judicial acts.

The VI distance between $\hat{\mathbf{z}}$ and the partition \mathbf{z}_b at the edge of the 95% credible ball is 0.233, much lower than the maximum achievable $\log_2 84 = 6.392$ distance among two generic partition of $V = 84$

nodes. Such a concentration of the posterior for \mathbf{z} around $\hat{\mathbf{z}}$ further strengthens the above claims.

We conclude with an assessment of the predictive schemes in Section 3.2, focusing in particular on the out-of-sample accuracy in predicting the edges associated with 10 randomly-selected held-out criminals. This test set comprises both affiliates and bosses from different *locali*, thus allowing for a comprehensive evaluation. For these criminals we compute the predictive edge probabilities via (31), conditioned only on the corresponding *locale* affiliation and the network among the in-sample criminals. Despite the use of such a limited information, we obtain an area under the ROC curve of 0.93. This remarkable result highlights the practical effectiveness of the newly proposed predictive methods. Such gains are confirmed by the posterior similarity matrix comprising in-sample and predictive co-clustering probabilities for estimating the allocations of out-of-sample nodes (see the Supplementary Material).

6 Conclusions and Future Research

We proposed and developed a new class of SBMs for node-colored multilayer networks that infer complex block-connectivity structures both within and across layers. The layers’ division is accounted for in the formation of blocks via a rigorous and interpretable probabilistic construction. The proposed pEx-SBM allows for uncertainty quantification, derivation of clustering and co-clustering properties, and, crucially, prediction of connectivity patterns for future nodes. Moreover, it is computationally tractable and, as shown in Sections 4–5, yields remarkable practical improvements. To the best of our knowledge, current models for multilayer networks do not display all these properties in a single formulation.

While we focused on binary undirected networks, weighted edges are readily tackled by replacing the beta-binomial likelihood (1) with a Poisson-gamma for count edges, or a Gaussian-Gaussian for continuous ones. Directed and bipartite networks can be instead addressed by modeling the row and column partitions of the non-symmetric or rectangular *supra*-adjacency matrix via two distinct pExP priors.

From a general perspective, our contribution showcases in the node-colored setting the potential of a broadly-impactful and novel idea. Namely that of designing models for multilayer networks that consist in matching the distinctive structures of a given multilayer network sub-class with the most suitable notion of probabilistic invariance encoded in a specific node-exchangeability assumption. For example, for edge-colored (multiplex) networks the natural pairing would be with *separate exchangeability*, which,

unlike for partial exchangeability, preserves the identity of nodes replicated across layers; see [Rebaudo et al. \(2024\)](#) for a review focusing on mixture models. A similar reasoning can be applied to dynamic networks, where the time dimension demands for further restrictions to the invariance structure.

Finally, proving frequentist posterior consistency of pEx-SBMs, as suggested by the empirical findings in [Figure 3](#), is an important, yet challenging, direction for future research. To this end, there are two main issues that need to be addressed. First, the available Bayesian asymptotic theory for single-layered SBMs (*e.g.*, [Geng et al., 2019](#)) has to be extended to node-colored multilayer networks and beyond the simplifying assumptions that are currently imposed in the literature (even for proving consistency of frequentist methods). Second, one should deal with the additional difficulty posed by the condition of partial exchangeability for which the investigation of frequentist properties is still limited (see, *e.g.*, [Catalano et al., 2022](#)).

Acknowledgements

Daniele Durante is funded by the European Union (ERC, NEMESIS, project number: 101116718). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. Antonio Lijoi and Igor Prünster were partially supported by the European Union – NextGenerationEU PRIN-PNRR (project P2022H5WZ9). This research was largely conducted while Francesco Gaffi was a Ph.D. student at Bocconi University.

Supplementary Material

S1 Background Material

Section [S1](#) reviews the core concepts underlying the proposed pEx-SBM model. More specifically, Section [S1.1](#) summarizes the general multilayer network framework presented in [Kivelä et al. \(2014\)](#) and clarifies how the node-colored networks we consider in our contribution represent a remarkable special

case of such a framework. Sections [S1.2–S1.3](#), discuss the connection between exchangeability and non-parametric priors defined via completely random measures, and then clarify the link between partial exchangeability and hierarchical compositions of such priors.

S1.1 Multilayer networks

Recalling the comprehensive review by [Kivelä et al. \(2014\)](#), the term *multilayer networks* refers to a general construction that encompasses most of the layered network structures treated in the literature. In the following, we summarize such a construction, and clarify how it includes, among others, the node-colored networks we consider in our contribution. Refer to Section 2.1 in [Kivelä et al. \(2014\)](#) for a more extensive and in-depth treatment.

Let $\mathbb{V} := \{1, \dots, V\}$ denote a set of V nodes in a network, for $V \in \mathbb{N}$, and define with $\delta \in \mathbb{N}$ the number of *aspects* of such a network, that is the number of directions of layerization. Depending on the type of multilayer network analyzed, each *aspect* could denote, *e.g.*, a division of the nodes into subpopulations (*e.g.*, node-colored networks), the presence of different types of relationships monitored (*e.g.*, edge-colored networks), or a time dimension (*e.g.*, dynamic networks), among others. Within this framework, any aspect $l \in \{1, \dots, \delta\}$ generates a set of *elementary layers* $L_l = \{1, \dots, d_l\}$ for every $d_l \in \mathbb{N}$. For example, in node-colored networks the elementary layers correspond to the different subpopulations to which the nodes belong. Under these settings, a *layer* is an element of the product space $L_1 \times \dots \times L_\delta$, (namely, a vector of coordinates identifying a location in the product space of elementary layers), whereas the set of *existing nodes* is defined as $\mathbb{V}_M \subset \mathbb{V} \times L_1 \times \dots \times L_\delta$. In symbols, node i exists in layer (j_1, \dots, j_δ) whenever $(i, j_1, \dots, j_\delta) \in \mathbb{V}_M$, for $j_l \in L_l$ and $l \in \{1, \dots, \delta\}$. Consistent with these definitions, in the aforementioned node-colored network example each layer coincides with an elementary layer, whereas an existing node is a pair that defines the node index in \mathbb{V} and the subpopulation to which it belongs. If such a node-colored network were also observed at different time points, this would generate an additional aspect (*i.e.*, time) with layers defined as pairs (subpopulation, time point) and existing nodes as triplets (node index, subpopulation, time point). Finally, denote with $\mathbb{E}_M \subset \mathbb{V}_M \times \mathbb{V}_M$ the set of edges. Such a set comprises all those couples $((i, j_1, \dots, j_\delta), (i', j'_1, \dots, j'_\delta))$ of existing nodes

$(i, j_1, \dots, j_\delta)$ and $(i', j'_1, \dots, j'_\delta)$ among which an edge has been observed. With these settings one can give the following definition of multilayer network.

DEFINITION S1.1. *A δ -multilayer network is a quadruplet $M := (\mathbb{V}_M, \mathbb{E}_M, \mathbb{V}, L)$ where $L = L_1 \times \dots \times L_\delta$ is the product space of layers, \mathbb{V} is the set of nodes, \mathbb{V}_M is the set of coordinates of existing nodes, and \mathbb{E}_M is the set of edges.*

Definition S1.1 allows for the existence of edges between nodes in the same layer, as well as across-layers, and even between copies of the same node existing in different layers. Moreover, notice that any multilayer network M , which is, in this formulation, a highly general object, can be flattened to obtain the node-labelled graph given by the couple $(\mathbb{V}_M, \mathbb{E}_M)$, which is called the *supra*-graph of M . See, *e.g.*, Figure 2 in Kivelä et al. (2014). Hence, there is a close relationship between multilayer structures, of any kind that such general construction can encompass, and *node-level covariate* structures. A δ -multilayer network can indeed be represented by a *supra*-adjacency matrix \mathbf{Y} , that is the adjacency matrix of the *supra*-graph. Once chosen an order in the product space of the layers L , *e.g.*, the lexicographic order, based on a natural order of each entry, then it is easy to describe the matrix \mathbf{Y} as a block matrix whose diagonal blocks are the adjacency matrices of each layer, and the off-diagonal blocks are the matrices of across-layers connections.

Under the above formulation, network structures comprising different types of relationships — with each layer encoding connections with respect to one of these types (*e.g.*, friendship, kinship, advice, ...) — can be represented through a block-diagonal *supra*-adjacency matrix excluding connections across layers. When $\delta = 1$, these are called *edge-colored* networks. If one considers the case in which each node exists in every layer, the network can be seen as a superposition of different sets of edges on the same collection of nodes, hence the layerization is simply given by considering a copy of the nodes for each set of edges. In this case, the term *multiplex network* is increasingly used in the literature on network data (see, *e.g.*, Mucha et al., 2010; Barbillon et al., 2017; MacDonald et al., 2022; Pensky and Wang, 2024; Noroozi and Pensky, 2024; Amini et al., 2024).

On the other hand, by forcing each node in the network to belong only to a single layer via the dis-

jointness condition

$$\forall i \in \mathbb{V} \quad \exists! (j_1, \dots, j_\delta) \in L \quad \text{s.t.} \quad (i, j_1, \dots, j_\delta) \in \mathbb{V}_M \quad (\text{S.1})$$

one can recover networks in which layers represent a division of nodes into distinct subpopulations, so that a node cannot have copies in different layers. If $\delta = 1$, these are called *node-colored* networks. In the main article, we propose a stochastic block model for such networks, based on the novel and general idea of matching the specific structure of each sub-class of multilayer network analyzed with the most suitable notion of probabilistic invariance encoded in a specific node-exchangeability assumption. As we clarify in the main article, in the context of node-colored networks the most natural invariance structure is the one encoded in the assumption of partial exchangeability.

S1.2 Exchangeable sequences and completely random measures

Let \mathbb{X} be a complete and separable metric space equipped with the Borel σ -algebra \mathcal{X} . \mathcal{P} is the space of probability distributions defined on $(\mathbb{X}, \mathcal{X})$ and endowed with the topology of weak convergence. $\sigma(\mathcal{P})$ is then the Borel σ -algebra of subsets of \mathcal{P} . Moreover, $[n]$ indicates the set of the first n integers $\{1, \dots, n\}$, for any $n \in \mathbb{N}$. For a sequence of \mathbb{X} -valued random elements $\mathbf{x} = (x_n)_{n \geq 1}$, defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ the following definition can be given.

DEFINITION S1.2. *\mathbf{x} is exchangeable if, for every $n \geq 1$ and any permutation π of the indices in $[n]$, it holds $(x_1, \dots, x_n) \stackrel{d}{=} (x_{\pi(1)}, \dots, x_{\pi(n)})$.*

By virtue of the representation result in [de Finetti \(1937\)](#), it is possible to state the following.

THEOREM S1.1 (de Finetti). *A sequence \mathbf{x} is exchangeable if and only if there exists a probability measure Q on the space of probability distributions $(\mathcal{P}, \sigma(\mathcal{P}))$ such that*

$$\mathbb{P}(x_1 \in A_1, \dots, x_n \in A_n) = \int_{\mathcal{P}} \prod_{i=1}^n P(A_i) Q(dP), \quad (\text{S.2})$$

for any A_1, \dots, A_n in \mathcal{X} and $n \geq 1$.

The probability measure Q directing the exchangeable sequence $(x_n)_{n \geq 1}$ is also termed *de Finetti*

measure and can be naturally interpreted as a prior distribution in the Bayesian framework. The representation theorem in (S.2) can be equivalently rephrased by stating that, given an exchangeable sequence $(x_n)_{n \geq 1}$, there exists a random probability measure \tilde{P} , defined on $(\mathbb{X}, \mathcal{X})$ and taking values in $(\mathcal{P}, \sigma(\mathcal{P}))$, such that, for any $n \in \mathbb{N}$, it holds

$$x_1, \dots, x_n \mid \tilde{P} \stackrel{\text{iid}}{\sim} \tilde{P}, \quad \text{with} \quad \tilde{P} \sim Q. \quad (\text{S.3})$$

The most popular instance of nonparametric prior Q is the Dirichlet process (DP) prior, introduced in Ferguson (1973), which selects discrete distributions with probability 1. As shown in Lijoi and Prünster (2010) most classes of discrete nonparametric priors, including the DP, can be seen as suitable transformations of *completely random measures*.

More specifically, let \mathcal{M} be the set of boundedly finite measures on \mathbb{X} equipped with the corresponding Borel σ -algebra $\sigma(\mathcal{M})$. For details on the definition of this σ -algebra, see Daley and Vere-Jones (2007). A completely random measure (CRM) $\tilde{\mu}$ on $(\mathbb{X}, \mathcal{X})$ is a measurable function on $(\Omega, \mathcal{F}, \mathbb{P})$ taking values in \mathcal{M} such that for any $k \geq 2$ and pairwise disjoint sets A_1, \dots, A_k in \mathcal{X} the random variables $\tilde{\mu}(A_1), \dots, \tilde{\mu}(A_k)$ are independent. CRMs have been introduced in Kingman (1967). Detailed treatments can be found in Kingman (1993) and Daley and Vere-Jones (2007). Any CRM $\tilde{\mu}$ fulfills the following representation. For $A \in \mathcal{X}$, we have

$$\tilde{\mu}(A) = \sum_{k \geq 1} u_k \delta_{x_k}(A) + \beta(A) + \int_0^\infty s \tilde{N}(\text{d}s, A), \quad (\text{S.4})$$

where $(x_k)_{k \geq 1}$ is a sequence in \mathbb{X} , $(u_k)_{k \geq 1}$ is a sequence of independent non-negative random variables, β is a fixed non-atomic boundedly finite measure on $(\mathbb{X}, \mathcal{X})$, and \tilde{N} is a Poisson process on $\mathbb{R}^+ \times \mathbb{X}$ independent of $(u_k)_{k \geq 1}$ and whose parameter measure ν satisfies $\int_{\mathbb{R}^+} \int_B \min\{s, 1\} \nu(\text{d}s, \text{d}x) < \infty$ for any bounded B in \mathcal{X} . Intuitively, a CRM can be seen as a superposition of a random measure with fixed atoms, a deterministic non-atomic drift and a part characterized by random jumps and random locations, whose intensities, distribution and mutual dependence are governed by a Poisson process. Here, as customary in the Bayesian nonparametric literature, we focus on CRMs $\tilde{\mu}$ with no fixed atoms and no drift. These are almost surely discrete and the corresponding Laplace functional admits the

following *Lévy–Khintchine* representation

$$\mathbb{E} \left[e^{-\int_{\mathbb{X}} f(x) \tilde{\mu}(dx)} \right] = \exp \left\{ - \int_{\mathbb{R}^+ \times \mathbb{X}} [1 - e^{-sf(x)}] \nu(ds, dx) \right\}, \quad (\text{S.5})$$

where $f : \mathbb{X} \rightarrow \mathbb{R}$ is a measurable function such that $\int |f| d\tilde{\mu} < \infty$ almost surely. The measure ν is known as the *Lévy intensity* of $\tilde{\mu}$ and regulates the intensity of the jumps of a CRM and their locations. By virtue of (S.5), it characterizes the CRM $\tilde{\mu}$. Two noteworthy examples are given by the σ -stable CRM $\tilde{\mu}_\sigma$ with Lévy intensity defined as

$$\nu(ds, dx) = \frac{\sigma}{\Gamma(1-\sigma)} s^{-1-\sigma} ds \alpha(dx), \quad (\text{S.6})$$

where α is a measure on $(\mathbb{X}, \mathcal{X})$ and $\sigma \in (0, 1)$, and by the *gamma CRM* $\tilde{\mu}$, which corresponds to

$$\nu(ds, dx) = e^{-s} s^{-1} ds \alpha(dx). \quad (\text{S.7})$$

If we impose the condition $0 < \tilde{\mu}(\mathbb{X}) < \infty$ almost surely, which is implied by $\nu(\mathbb{R}^+ \times \mathbb{X}) = \infty$ and $\int_{\mathbb{R}^+ \times \mathbb{X}} [1 - e^{-\lambda s}] \nu(ds, dx) < \infty$ for any $\lambda > 0$, then, as proved in Regazzini et al. (2003), it is possible to define

$$\tilde{P}(A) := \frac{\tilde{\mu}(A)}{\tilde{\mu}(\mathbb{X})}, \quad (\text{S.8})$$

for any $A \in \mathcal{X}$. Random probability measures defined as in (S.8) form the class of *normalized random measures with independent increments* (NRMIs), introduced in Regazzini et al. (2003). As clarified in (S.5), such measures are characterized by ν . Moreover, $\tilde{\mu}$ is termed *homogeneous CRM* if its Lévy intensity factorizes as $\nu(ds, dx) = \rho(s) ds \alpha(dx)$ for some measurable positive function ρ on \mathbb{R}^+ and some measure α on \mathbb{X} ; in this case $\tilde{\mu}$ is identified by ρ and α . Intensities in (S.6) and (S.7) are in this class. Note that for homogeneous CRMs $\int_{\mathbb{R}^+ \times \mathbb{X}} [1 - e^{-\lambda s}] \nu(ds, dx) < \infty$ is equivalent to the finiteness of α . Setting $c := \alpha(\mathbb{X})$, we have

$$\nu(ds, dx) = \rho(s) ds c G(dx) \quad (\text{S.9})$$

where $G(\cdot) := \alpha(\cdot)/\alpha(\mathbb{X})$ is now a probability measure on \mathbb{X} . Therefore, we can write

$$\tilde{P} \sim \text{NRMI}(\rho, c, G) \quad (\text{S.10})$$

for some measurable positive $\rho, c > 0$ and probability measure G . Note that $\mathbb{E}[\tilde{P}(A)] = G(A)$ for any

$A \in \mathcal{X}$, which means that, if $x \mid \tilde{P} \sim \tilde{P}$ then $\mathbb{P}(x \in A) = G(A)$ for any $A \in \mathcal{X}$. For particular choices of ρ it is possible to recover noteworthy nonparametric priors. For example, taking $\rho(s) = e^{-s}s^{-1}$ as in (S.7), we are normalizing a gamma CRM and, hence, the correspondent \tilde{P} in (S.8) is a DP. In this case the total mass c , also called *concentration parameter*, is often denoted as $\theta > 0$. Hence, we shall also write $\tilde{P} \sim \text{DP}(\theta, P_0)$, where $\theta = \alpha(\mathbb{X})$ and $P_0 = \mathbb{E}[\tilde{P}]$. If instead we choose ρ as in (S.6), then \tilde{P} in (S.8) is called *normalized stable process* (NSP) (Kingman, 1975).

S1.3 Partially exchangeable arrays and hierarchical processes

As discussed in the main article, we consider a generalization of exchangeability in order to induce a random partition prior on the nodes which accounts for information from the layer division. Such a generalization is known as *partial exchangeability* (de Finetti, 1938) and is a more natural hypothesis of dependence for random elements divided in a finite number d of layers. In fact, as clarified in Definition 2.1 of the main article, it assumes that the joint distribution of the entries in a given infinite random array $\mathbf{X}^\infty = \{(x_{ji})_{i \geq 1} \mid j \in [d]\}$ is invariant with respect to within-layer permutations, but not necessarily with respect to across-layer ones.

Under the partial exchangeability assumption, the following extension of representation (S.2) holds.

THEOREM S1.2 (de Finetti). *The infinite random array \mathbf{X}^∞ is partially exchangeable if and only if there exists a measure Q on the space of d -dimensional vectors of probability measures \mathcal{P}^d endowed with the product σ -algebra $\bigotimes_{j=1}^d \sigma(\mathcal{P})$ such that*

$$\mathbb{P}(x_{11} \in A_{11}, \dots, x_{1V_1} \in A_{1V_1}, \dots, x_{d1} \in A_{d1}, \dots, x_{dV_d} \in A_{dV_d}) = \int_{\mathcal{P}^d} \prod_{j=1}^d \prod_{i=1}^{V_j} P_j(A_{ji}) Q(dP_1, \dots, dP_d), \quad (\text{S.11})$$

for any $A_{11}, \dots, A_{dV_d} \in \mathcal{X}$ and any $(V_1, \dots, V_d) \in \mathbb{N}^d$.

Again the quantity Q is called *de Finetti measure*. By extending the representation in (S.3) to the partial exchangeability setting, one obtains (4) of the main article, which is equivalent to Theorem S1.2. Namely, given a partially exchangeable infinite array \mathbf{X}^∞ , there exists a vector of random probability measures $(\tilde{P}_1, \dots, \tilde{P}_d)$, each defined on $(\mathbb{X}, \mathcal{X})$, and taking values in $(\mathcal{P}^d, \bigotimes_{j=1}^d \sigma(\mathcal{P}))$, such that, for

any $j_1, \dots, j_k \in [d]$ and $i_1, \dots, i_k \geq 1$, one has

$$\begin{aligned} x_{j_1 i_1}, \dots, x_{j_k i_k} \mid (\tilde{P}_1, \dots, \tilde{P}_d) &\stackrel{\text{iid}}{\sim} \tilde{P}_{j_1} \times \dots \times \tilde{P}_{j_k}, \\ (\tilde{P}_1, \dots, \tilde{P}_d) &\sim Q. \end{aligned} \tag{S.12}$$

Partial exchangeability reduces to exchangeability if $\tilde{P}_j = \tilde{P}$ almost surely for any $j \in [d]$ and some random probability measure \tilde{P} . In this case, (S.12) reduces to (S.3). Conversely, if $(\tilde{P}_j)_{j \in [d]}$ are independent, that is Q is a product law on \mathcal{P}^d , then the rows of the array \mathbf{X}^∞ are independent. These particular cases represent the extremes in the range of borrowing of information among nodes coming from different layers. Defining a partially exchangeable model for an array, then, amounts to choosing a de Finetti measure Q for a vector of random probability measures which can allow flexible borrowing of information both within and across layers. To this end, and recalling the considerations for the classical exchangeability setting, a natural and routinely-employed solution is to rely on hierarchical compositions of the nonparametric priors presented in Section S1.2. This yields the so-called *hierarchical normalized random measures with independent increments* (H-NRMIs) priors; see Definition 2.2 in the main article — which extends (S.10) to the partial exchangeability setting — and [Camerlenghi et al. \(2019\)](#) for an in-depth treatment.

Example S1.3 provides two popular instances of H-NRMIs, which further clarify the related construction via hierarchical compositions of the nonparametric priors.

EXAMPLE S1.3. *If for both levels of the hierarchy we consider a DP, namely*

$$\begin{aligned} \tilde{P}_1, \dots, \tilde{P}_d \mid \tilde{P}_0 &\stackrel{\text{iid}}{\sim} \text{DP}(\theta, \tilde{P}_0), \\ \tilde{P}_0 &\sim \text{DP}(\theta_0, P_0), \end{aligned} \tag{S.13}$$

then $(\tilde{P}_1, \dots, \tilde{P}_d) \sim \text{H-DP}(\theta, \theta_0, P_0)$ ([Teh et al., 2006](#)).

If instead we consider two levels of NSP, i.e.,

$$\begin{aligned} \tilde{P}_1, \dots, \tilde{P}_d \mid \tilde{P}_0 &\stackrel{\text{iid}}{\sim} \text{NSP}(\sigma, \tilde{P}_0), \\ \tilde{P}_0 &\sim \text{NSP}(\sigma_0, P_0), \end{aligned} \tag{S.14}$$

then, $(\tilde{P}_1, \dots, \tilde{P}_d) \sim \text{H-NSP}(\sigma, \sigma_0, P_0)$ ([Camerlenghi et al., 2019](#)).

Since H-NRMIs represent a particular specification of the de Finetti measure Q in (S.12), if the j -th row, $(x_{j1}, \dots, x_{jV_j})$, of a finite array $\mathbf{X} \subset \mathbf{X}^\infty$ is a conditionally iid sample from \tilde{P}_j in Definition 2.2 of the main article, for any $V_j \geq 1$ and $j \in [d]$, then \mathbf{X}^∞ is partially exchangeable according to its rows. Moreover, because of the almost sure discreteness of each \tilde{P}_j and \tilde{P}_0 , the probability of a tie among the entries in the array \mathbf{X} is positive both within and across rows (*i.e.*, layers), that is $\mathbb{P}(x_{ij} = x_{i'j'}) > 0$, for any $i \in [V_j]$, $i' \in [V_{j'}]$ and $j, j' \in [d]$. As clarified in the main article, the proposed partially exchangeable partition prior for the node allocations to groups in the pEx-SBM model is directly obtained from the within- and across-layer clustering structures induced by these ties.

S2 Results for the Hierarchical Normalized Stable Case

In the main article, all the results obtained for the general class of H-NRMI are also specialized to the H-DP case. Here we specialize the general results also to the H-NSP case.

From Proposition 2.3, we can deduce the following.

COROLLARY S2.1. *Let $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \sigma, \sigma_0)$ where $\text{pExp}(\mathbf{V}; \sigma, \sigma_0)$ denotes the partition structure induced by a H-NSP with parameters $\sigma, \sigma_0 \in (0, 1)$. Then*

$$\mathbb{P}(z_v = h \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \mathbb{1}_{\{\ell_{\cdot h}^{-v} = 0\}} \frac{H^{-v} \sigma_0}{|\ell^{-v}|} \frac{\ell_{j \cdot}^{-v} \sigma}{V_j - 1} + \mathbb{1}_{\{\ell_{\cdot h}^{-v} \neq 0\}} \left[\frac{\ell_{\cdot h}^{-v} - \sigma_0}{|\ell^{-v}|} \frac{\ell_{j \cdot}^{-v} \sigma}{V_j - 1} + \frac{n_{jh}^{-v} - \ell_{jh}^{-v} \sigma}{V_j - 1} \right], \quad (\text{S.15})$$

for each $h \in [H^{-v} + 1]$.

Theorem 2.6 yields, instead, the following expressions.

COROLLARY S2.2. *Let $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \sigma, \sigma_0)$. Then, if both v and u are in the same layer j , we have*

$$\begin{aligned} \mathbb{P}(z_v = z_u \mid \mathbf{z}^{-vu}, \mathbf{w}^{-vu}) &= \frac{1}{(V_j - 2)(V_j - 1)} \left\{ \sum_{h=1}^{H^{-vu}} (n_{jh}^{-vu} - \ell_{jh}^{-vu} \sigma) (n_{jh}^{-vu} - \ell_{jh}^{-vu} \sigma + 1) \right. \\ &\quad + \ell_{j \cdot}^{-vu} \sigma \left[(1 - \sigma) + \frac{2}{|\ell^{-vu}|} \sum_{h=1}^{H^{-vu}} (n_{jh}^{-vu} - \ell_{jh}^{-vu} \sigma) (\ell_{\cdot h}^{-vu} - \sigma_0) \right] \\ &\quad \left. + \frac{\ell_{j \cdot}^{-vu} (\ell_{j \cdot}^{-vu} + 1) \sigma^2}{|\ell^{-vu}| (|\ell^{-vu}| + 1)} \left[\sum_{h=1}^{H^{-vu}} (\ell_{\cdot h}^{-vu} - \sigma_0) (\ell_{\cdot h}^{-vu} - \sigma_0 + 1) + H^{-vu} \sigma_0 (1 - \sigma_0) \right] \right\}, \end{aligned} \quad (\text{S.16})$$

whereas, if node v is in layer j and node u is in layer j' , with $j \neq j'$, it follows that

$$\begin{aligned} \mathbb{P}(z_v = z_u \mid \mathbf{z}^{-vu}, \mathbf{w}^{-vu}) &= \frac{1}{(V_j - 1)(V_{j'} - 1)} \left\{ \sum_{h=1}^{H^{-vu}} (n_{jh}^{-vu} - \ell_{jh}^{-vu} \sigma) (n_{j'h}^{-vu} - \ell_{j'h}^{-vu} \sigma) \right. \\ &\quad + \frac{\sigma}{|\ell^{-vu}|} \left[\ell_{j\cdot}^{-vu} \sum_{h=1}^{H^{-vu}} (n_{j'h}^{-vu} - \ell_{j'h}^{-vu} \sigma) (\ell_{\cdot h}^{-vu} - \sigma_0) + \ell_{j'\cdot}^{-vu} \sum_{h=1}^{H^{-vu}} (n_{jh}^{-vu} - \ell_{jh}^{-vu} \sigma) (\ell_{\cdot h}^{-vu} - \sigma_0) \right] \\ &\quad \left. + \frac{\ell_{j\cdot}^{-vu} (\ell_{j'\cdot}^{-vu}) \sigma^2}{|\ell^{-vu}| (|\ell^{-vu}| + 1)} \left[\sum_{h=1}^{H^{-vu}} (\ell_{\cdot h}^{-vu} - \sigma_0) (\ell_{\cdot h}^{-vu} - \sigma_0 + 1) + H^{-vu} \sigma_0 (1 - \sigma_0) \right] \right\}. \end{aligned} \quad (\text{S.17})$$

The following statement provides a prior elicitation result for the H-NSP case, similar to the one in Proposition 2.9.

PROPOSITION S2.3. *Let $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \sigma, \sigma_0)$. Then, for any generic node v in layer j , we have*

$$\sigma \left(1 + \frac{|\mathbf{H}_j^{-v}|}{|\ell^{-v}|} \sigma_0 \right) \leq \frac{V_j - 1}{2\ell_{j\cdot}^{-v}} \implies \mathbb{P}(z_v \in \mathbf{H}_j^{-v} \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) \geq \mathbb{P}(z_v \notin \mathbf{H}_j^{-v} \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}), \quad (\text{S.18})$$

for every node $v \in [V]$ and corresponding layer $j \in [d]$. In particular, the left hand side of (S.18) is implied if $\sigma \leq 0.5(1 + (V_j - 1)\sigma_0/d)^{-1}$. As discussed in Remark 2.10 in the article, this condition also guarantees $\mathbb{P}(z_v \in \mathbf{H}_j^{-v} \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) \geq \mathbb{P}(z_v \in \mathbf{H}^{-v} \setminus \mathbf{H}_j^{-v} \mid \mathbf{z}^{-v}, \mathbf{w}^{-v})$.

We conclude by tailoring the results on the conditional probabilities in Proposition 3.1 to the H-NSP.

COROLLARY S2.4. *If $\mathbf{z} \sim \text{pExp}(\mathbf{V}; \sigma, \sigma_0)$, then, for any node v in layer j , we have*

$$\begin{aligned} \mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) \\ = \mathbb{1}_{\{\tau = \tau_{jht}^{-v}\}} \frac{q_{jht}^{-v} - \sigma}{V_j - 1} + \mathbb{1}_{\{\tau = \ell_{j\cdot}^{-v} + 1\}} \frac{\ell_{j\cdot}^{-v} \sigma}{V_j - 1} \left[\mathbb{1}_{\{h \leq H^{-v}\}} \frac{\ell_{\cdot h}^{-v} - \sigma_0}{|\ell^{-v}|} + \mathbb{1}_{\{h = H^{-v} + 1\}} \frac{H^{-v} \sigma_0}{|\ell^{-v}|} \right], \end{aligned} \quad (\text{S.19})$$

for any $h \in [H^{-v} + 1]$, $\tau \in [\ell_{j\cdot}^{-v} + 1]$, where τ_{jh}^{-v} is defined as in Proposition 3.1 of the main article.

S3 Proofs of Theorems, Propositions and Corollaries

Section S3 provides the proofs of the theorems, propositions and corollaries stated in the main article and in these supplementary materials.

We start proving the expressions for the joint conditional probabilities given in Proposition 3.1 and its Corollaries 3.2 and S2.4, which we will also leverage on in the proofs of other results.

Proof of Proposition 3.1. By Bayes rule

$$\mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \frac{p(z_v = h, w_v = \tau, \mathbf{z}^{-v}, \mathbf{w}^{-v})}{p(\mathbf{z}^{-v}, \mathbf{w}^{-v})}, \quad (\text{S.20})$$

where both the numerator and the denominator can be directly retrieved from the pEPPF in (8) of the main article. To this end, it suffices to notice that, having fixed a specific configuration for the group and subgroup allocations — instead of just an array of frequencies — the summation over all the configurations and the multiplication by the product of multinomial factors in (8) is not required since these operations yield the total mass assigned to equiprobable configurations with equal array of frequencies. Hence, the denominator in (S.20) can be analytically evaluated via

$$p(\mathbf{z}^{-v}, \mathbf{w}^{-v}) = \Phi_{H^{-v},0}^{(|\ell^{-v}|)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot H^{-v}}^{-v}) \Phi_{\ell_{j\cdot}^{-v},j}^{(V_j-1)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jH^{-v}}^{-v}) \prod_{j'=1, j' \neq j}^d \Phi_{\ell_{j'\cdot}^{-v},j'}^{(V_{j'})}(\mathbf{q}_{j'1}^{-v}, \dots, \mathbf{q}_{j'H^{-v}}^{-v}), \quad (\text{S.21})$$

where j is the layer of node v .

To derive a similar expression for the numerator, it is worth analyzing separately the two situations in which τ corresponds either to an already-existing subgroup in layer j or to a new one. In the former case, for any $\tau \in [\ell_{j\cdot}^{-v}]$ the probability in (S.20) is non-zero just when h corresponds to the sociability profile associated to subgroup τ . It is easy to see that for any h , the set of such subgroup indices is given by \mathbb{T}_{jh}^{-v} defined in (11) of the main article. Note that $|\mathbb{T}_{jh}^{-v}| = \ell_{jh}^{-v}$. Now, if τ_{jh}^{-v} denotes the vector obtained ordering \mathbb{T}_{jh}^{-v} , the subgroup corresponding to τ_{jht}^{-v} is the t -th subgroup with sociability profile h in layer j . Hence, to compute the numerator in (S.20) for $\tau = \tau_{jht}^{-v}$, it suffices to increase the frequency q_{jht}^{-v} by one, whereas ℓ^{-v} remains unchanged. As a result

$$p(z_v = h, w_v = \tau_{jht}^{-v}, \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \Phi_{H^{-v},0}^{(|\ell^{-v}|)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot H^{-v}}^{-v}) \cdot \Phi_{\ell_{j\cdot}^{-v},j}^{(V_j)}(\mathbf{q}_{j1}^{-v}, \dots, \mathbf{q}_{jh}^{-v} + \mathbf{e}_t, \dots, \mathbf{q}_{jH^{-v}}^{-v}) \prod_{j'=1, j' \neq j}^d \Phi_{\ell_{j'\cdot}^{-v},j'}^{(V_{j'})}(\mathbf{q}_{j'1}^{-v}, \dots, \mathbf{q}_{j'H^{-v}}^{-v}) \quad (\text{S.22})$$

where \mathbf{e}_t is a ℓ_{jh}^{-v} -dimensional vector of 0s, with a 1 in the t -th entry.

When, instead, τ is a new subgroup in layer j , *i.e.*, $\tau = \ell_{j\cdot}^{-v} + 1$, then node v might be assigned either a sociability profile already present in the network or, alternatively, a previously-unseen one,

i. e., $h = H^{-v} + 1$. This requires adding a new entry to \mathbf{q}_{jh}^{-v} and increasing ℓ^{-v} , for the creation of the new subgroup. Note that, when $h = H^{-v} + 1$ then $(\mathbf{q}_{jh}^{-v}, 1) = 1$. Hence, in this case we have

$$p(z_v = h, w_v = \ell_{j\cdot}^{-v} + 1, \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \Phi_{H_h^{-v}, 0}^{(|\ell^{-v}|+1)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot h}^{-v} + 1, \dots, \ell_{\cdot H^{-v}}^{-v}) \cdot \Phi_{\ell_{j\cdot}^{-v}+1, j}^{(V_j)}(\mathbf{q}_{j1}^{-v}, \dots, (\mathbf{q}_{jh}^{-v}, 1), \dots, \mathbf{q}_{jH_h^{-v}}^{-v}) \prod_{j'=1, j' \neq j}^d \Phi_{\ell_{j'\cdot}^{-v}, j'}^{(V_{j'})}(\mathbf{q}_{j'1}^{-v}, \dots, \mathbf{q}_{j'H^{-v}}^{-v}) \quad (\text{S.23})$$

Replacing (S.21)–(S.23) in (S.20) yields the expressions in (25)–(26) of the main article. \square

Proof of Corollary 3.2. Recall that the EPPF induced by a DP with concentration parameter $\bar{\theta} > 0$ implies that the probability of any partition of n objects in K clusters is given by

$$\Phi_K^{(n)}(n_1, \dots, n_K) = \frac{\bar{\theta}^K}{[\bar{\theta}]_n} \prod_{\kappa=1}^K (n_\kappa - 1)! \quad (\text{S.24})$$

for any frequency vector (n_1, \dots, n_K) such that $n = \sum_{\kappa=1}^K n_\kappa$, where $[\cdot]_n$ denotes the n -th ascending factorial.

When $\tau \in \mathbb{T}_{jh}^{-v}$, it suffices to substitute (S.24) in (25) to get the first summand in (27). As for the case $\tau = \ell_{j\cdot}^{-v} + 1$, if $h \in [H^{-v}]$ then $H_h^{-v} = H^{-v}$ and we have

$$\frac{\Phi_{H^{-v}, 0}^{(|\ell^{-v}|+1)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot h}^{-v} + 1, \dots, \ell_{\cdot H^{-v}}^{-v})}{\Phi_{H^{-v}, 0}^{(|\ell^{-v}|)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot H^{-v}}^{-v})} = \frac{\ell_{\cdot h}^{-v}}{\theta_0 + |\ell^{-v}|} \quad (\text{S.25})$$

while if $h = H^{-v} + 1$, since $H_h^{-v} = H^{-v} + 1$,

$$\frac{\Phi_{H^{-v}+1, 0}^{(|\ell^{-v}|+1)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot H^{-v}}^{-v}, 1)}{\Phi_{H^{-v}, 0}^{(|\ell^{-v}|)}(\ell_{\cdot 1}^{-v}, \dots, \ell_{\cdot H^{-v}}^{-v})} = \frac{\theta_0}{\theta_0 + |\ell^{-v}|}. \quad (\text{S.26})$$

In both cases the second ratio in (26) gives the common factor of the second summand in (27). \square

Proof of Corollary S2.4. The EPPF induced by a NSP with parameter $\bar{\sigma} \in (0, 1)$ gives

$$\Phi_K^{(n)}(n_1, \dots, n_K) = \frac{(K-1)! \bar{\sigma}^{K-1}}{(n-1)!} \prod_{\kappa=1}^K [1 - \bar{\sigma}]_{n_\kappa - 1} \quad (\text{S.27})$$

Substituting (S.27) in (25)–(26), with considerations as in proof of Corollary 3.2, yields (S.19). \square

We now provide proofs of the remaining results, in order of appearance.

Proof of Proposition 2.3. It suffices to note that

$$\mathbb{P}(z_v = h \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}) = \sum_{\tau=1}^{\ell_{j\cdot}^{-v}+1} \mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v}), \quad (\text{S.28})$$

where expressions of $\mathbb{P}(z_v = h, w_v = \tau \mid \mathbf{z}^{-v}, \mathbf{w}^{-v})$ for any choice of $(h, \tau) \in [H^{-v} + 1] \times [\ell_{j\cdot}^{-v} + 1]$ are given in Proposition 3.1. \square

Proof of Corollary 2.5. Equation (14) in Corollary 2.5 can be obtained from (12) by replacing the general $\Phi^{(\cdot)}$ in (13) with those specific to the DP case, given in (S.24), or alternatively, by summing the H-DP joint conditional probabilities in (27) over all the possible choices of $\tau \in [\ell_{j\cdot}^{-v} + 1]$. \square

Proof of Theorem 2.6. We can write

$$\begin{aligned} \mathbb{P}(\{z_v = z_u\}, w_v = \tau, w_u = \tau' \mid \mathbf{z}^{-vu}, \mathbf{w}^{-vu}) &= \frac{p(\{z_v = z_u\}, w_v = \tau, w_u = \tau', \mathbf{z}^{-vu}, \mathbf{w}^{-vu})}{p(\mathbf{z}^{-vu}, \mathbf{w}^{-vu})} \\ &= \frac{\sum_{h=1}^{H^{-vu}+1} p(z_v = h, z_u = h, w_v = \tau, w_u = \tau', \mathbf{z}^{-vu}, \mathbf{w}^{-vu})}{p(\mathbf{z}^{-vu}, \mathbf{w}^{-vu})}. \end{aligned} \quad (\text{S.29})$$

Now, if v and u are in the same layer j , the probability within the summation at the numerator will include, for any h, τ and τ' the common factor

$$\prod_{j'=1, j' \neq j}^d \Phi_{\ell_{j\cdot}^{-vu}, j'}^{(V_{j'})}(\mathbf{q}_{j'1}^{-vu}, \dots, \mathbf{q}_{j'H^{-vu}}^{-vu}),$$

given by the allocations in all the other layers.

The form of the remaining factor depends, instead, on τ and τ' . If both τ and τ' are already-existing subgroups, then $\tau = \tau_{jht}^{-vu}$ and $\tau' = \tau_{jht'}^{-vu}$ for some (possibly equal) $t, t' \in [\ell_{jh}^{-vu}]$, where the vector τ_{jh}^{-vu} is defined as in Proposition 3.1. In this case the multiplicative factor is equal to

$$\Phi_{H^{-vu}, 0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot H^{-vu}}^{-vu}) \Phi_{\ell_{j\cdot}^{-vu}, j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jh}^{-vu} + \mathbf{e}_t + \mathbf{e}_{t'}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu}).$$

If, instead, τ is already occupied and τ' is new, that is $\tau = \tau_{jht}^{-vu}$ for some $t \in [\ell_{jh}^{-vu}]$ and $\tau' = \ell_j^{-vu} + 1$, then the factor is

$$\Phi_{H^{-vu},0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot h}^{-vu} + 1, \dots, \ell_{\cdot H^{-vu}}^{-vu}) \Phi_{\ell_j^{-vu},j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, (\mathbf{q}_{jh}^{-vu} + \mathbf{e}_t, 1), \dots, \mathbf{q}_{jH^{-vu}}^{-vu}).$$

When both τ and τ' correspond to the same new subgroup, *i.e.*, $\tau = \tau' = \ell_j^{-vu} + 1$, we have

$$\Phi_{H^{-vu},0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot h}^{-vu} + 1, \dots, \ell_{\cdot H^{-vu}}^{-vu}) \Phi_{\ell_j^{-vu},j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, (\mathbf{q}_{jh}^{-vu}, 2), \dots, \mathbf{q}_{jH^{-vu}}^{-vu}),$$

while if both are new but different, *i.e.*, $\tau = \ell_j^{-vu} + 1$, $\tau' = \ell_j^{-vu} + 2$ (or vice versa), the multiplicative factor is

$$\Phi_{H^{-vu},0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot h}^{-vu} + 2, \dots, \ell_{\cdot H^{-vu}}^{-vu}) \Phi_{\ell_j^{-vu},j}^{(V_j)}(\mathbf{q}_{j1}^{-vu}, \dots, (\mathbf{q}_{jh}^{-vu}, 1, 1), \dots, \mathbf{q}_{jH^{-vu}}^{-vu}).$$

As in (S.21), it is also easy to check that the denominator in (S.29) is equal to

$$\Phi_{H^{-vu},0}^{(|\ell^{-vu}|)}(\ell_{\cdot 1}^{-vu}, \dots, \ell_{\cdot H^{-vu}}^{-vu}) \Phi_{\ell_j^{-vu},j}^{(V_j-2)}(\mathbf{q}_{j1}^{-vu}, \dots, \mathbf{q}_{jH^{-vu}}^{-vu}) \prod_{j'=1, j' \neq j}^d \Phi_{\ell_{j'}^{-vu},j'}^{(V_{j'})}(\mathbf{q}_{j'1}^{-vu}, \dots, \mathbf{q}_{j'H^{-vu}}^{-vu}).$$

Replacing the above factors within the ratio in (S.29) and summing over all the possible choices of τ and τ' , yields expressions (15)–(16). The proof of (17)–(18) in Theorem 2.6 follows along the same line of reasoning. It suffices to note that when v and u are not in the same layer it is not possible for v and u to be allocated to the same subgroup, neither new nor already-occupied. \square

Proof of Corollary 2.8. Similarly to the proof of Corollary 2.5, (19) and (20) in Corollary 2.8 can be derived from (15) and (17), respectively, by replacing the general functions $\Phi^{(\cdot)}$ in (16) and (18) with those specific to the DP case in (S.24). \square

Proof of Proposition 2.9. Summing the predictive probabilities in (14) over $h \in \mathbf{H}_j^{-v}$ and $h \in \mathbf{H}^{-v} \setminus \mathbf{H}_j^{-v}$, we have that the inequality in the right hand side of (21) is satisfied if and only if

$$\frac{|\ell^{-v}|p_v}{\theta_0 + |\ell^{-v}|} \frac{\theta}{\theta + V_j - 1} + \frac{V_j - 1}{\theta + V_j - 1} \geq \frac{|\ell^{-v}|(1 - p_v)}{\theta_0 + |\ell^{-v}|} \frac{\theta}{\theta + V_j - 1}, \quad (\text{S.30})$$

where $p_v = (1/|\ell^{-v}|) \sum_{h \in \mathbf{H}_j^{-v}} \ell_h^{-v}$. This is equivalent to

$$p_v \geq \frac{1}{2} \left(1 - \frac{|\ell^{-v}| + \theta_0}{|\ell^{-v}|} \frac{V_j - 1}{\theta} \right). \quad (\text{S.31})$$

Therefore, whenever the right hand side of (S.31) is negative, regardless of the value of the proportion p_v of subgroups with sociability profiles already observed in layer j , the inequality is always satisfied. This is implied by the left hand side of (21).

Summing to the right hand side of (S.30) the probability of node v being assigned a new sociability profile, *i.e.*,

$$\frac{\theta_0}{\theta_0 + |\ell^{-v}|} \frac{\theta}{\theta + V_j - 1} \quad (\text{S.32})$$

we obtain the inequality in the right hand side of (22). With the previous strategy we can again retrieve the sufficient condition. \square

Proof of Proposition S2.3. Summing the predictive probabilities in (S.15) over $h \in \mathbf{H}_j^{-v}$ and $h \in \mathbf{H}^{-v} \setminus \mathbf{H}_j^{-v}$, we have that the inequality in the right hand side of (S.18) is satisfied if and only if

$$p_v - \frac{|\mathbf{H}_j^{-v}| \sigma_0}{|\ell^{-v}|} + \frac{V_j - 1}{\sigma \ell_{j \cdot -v}} - 1 \geq (1 - p_v) + \frac{|\mathbf{H}_j^{-v}| \sigma_0}{|\ell^{-v}|}, \quad (\text{S.33})$$

where $p_v = (1/|\ell^{-v}|) \sum_{h \in \mathbf{H}_j^{-v}} \ell_h^{-v}$. This is equivalent to

$$1 - p_v \geq \frac{V_j - 1}{2\sigma \ell_{j \cdot -v}} - \frac{|\mathbf{H}_j^{-v}| \sigma_0}{|\ell^{-v}|}. \quad (\text{S.34})$$

Therefore, whenever the right hand side of (S.34) is greater than 1, regardless of the value of the proportion p_v of subgroups with sociability profiles already observed in layer j , the inequality is always satisfied. This is implied by the left hand side of (S.18). In particular, since $|\mathbf{H}_j^{-v}| \leq V_j - 1$, $|\ell| \geq d$, and $\ell_{j \cdot -v} \leq V_j - 1$, we have that (S.34) is implied by $\sigma \leq 0.5(1 + (V_j - 1)\sigma_0/d)^{-1}$. \square

Proof of Corollary 3.3. Corollary 3.3 can be directly deduced from (25)–(26). \square

Proof of Corollary S2.1. Equation (S.15) in Corollary S2.1 can be also obtained from (12) by replacing the general functions $\Phi^{(\cdot)}$ in (13) with those specific to the NSP case, given in (S.27), or

alternatively, by summing the H-NSP joint conditional probabilities in (S.19) over all the possible choices of $\tau \in [\ell_j^- + 1]$. \square

Proof of Corollary S2.2. Similarly to the proof of Corollary S2.1, (S.16) and (S.17) in Corollary S2.2 can be derived from (15) and (17), respectively, by replacing the general functions $\Phi^{(\cdot)}$ in (16) and (18) with those specific to the NSP case in (S.27). \square

S4 Additional Methodological and Empirical Results

We provide below details on $\text{pExp}(\mathbf{V}; \theta, \theta_0)$ with hypepriors on θ and θ_0 , along with additional empirical results which complement the analyses in Sections 4 and 5 of the main article.

S4.1 $\text{pExp}(\mathbf{V}; \theta, \theta_0)$ with hypepriors on θ and θ_0

As discussed in Section 3.1 of the main article it is possible to include gamma hyperpriors for the parameters θ and θ_0 of $\text{pExp}(\mathbf{V}; \theta, \theta_0)$ through a tractable construction that yields closed-form conjugate full conditionals. This allows for a direct modification of Algorithm 1 that includes an additional step to sample from these full-conditionals of the $\text{pExp}(\mathbf{V}; \theta, \theta_0)$ parameters.

More specifically, extending and adapting the results of Escobar and West (1995) from the DP to the H-DP setting within our formulation, we have that, when $\theta \sim \text{gamma}(\alpha, \beta)$ and $\theta_0 \sim \text{gamma}(\alpha_0, \beta_0)$, the two full-conditionals for θ and θ_0 are

$$p(\theta_0 \mid \mathbf{z}, \mathbf{w}, \mathbf{Y}, \theta) \propto \frac{\theta_0^H}{[\theta_0]_{|\ell|}} \times p(\theta_0) \quad \text{and} \quad p(\theta \mid \mathbf{z}, \mathbf{w}, \mathbf{Y}, \theta_0) \propto \prod_{j=1}^d \frac{\theta^{\ell_j}}{[\theta]_{V_j}} \times p(\theta),$$

where $p(\theta_0)$ and $p(\theta)$ are the densities of the gamma hypepriors for θ_0 and θ , respectively, whereas $[\theta_0]_{|\ell|}$ and $[\theta]_{V_j}$ correspond to the two ascending factorials that can be also expressed as $\Gamma(\theta_0 + |\ell|)/\Gamma(\theta_0)$ and $\Gamma(\theta + V_j)/\Gamma(\theta)$. Therefore, $1/[\theta_0]_{|\ell|} \propto \text{B}(\theta_0, |\ell|) = \int_0^1 \eta_0^{\theta_0-1} (1 - \eta_0)^{|\ell|-1} d\eta_0$ and $1/[\theta]_{V_j} \propto \text{B}(\theta, V_j) = \int_0^1 \eta_j^{\theta-1} (1 - \eta_j)^{V_j-1} d\eta_j$, for each $j = 1, \dots, d$. Hence, introducing augmented variables η_0 and η_1, \dots, η_d

such that

$$(\eta_0 \mid \theta_0, \mathbf{z}, \mathbf{w}) \sim \text{beta}(\theta_0, |\ell|) \quad \text{and} \quad (\eta_j \mid \theta) \stackrel{\text{ind}}{\sim} \text{beta}(\theta, V_j), \quad j = 1, \dots, d, \quad (\text{S.35})$$

and defining $\nu_0 := \log(\eta_0) \leq 0$ and $\nu := \sum_{j=1}^d \log(\eta_j) \leq 0$, we have

$$(\theta_0 \mid \eta_0, \mathbf{z}, \mathbf{w}) \sim \text{gamma}(\alpha_0 + H, \beta_0 - \nu_0) \quad \text{and} \quad (\theta \mid \eta_1, \dots, \eta_d, \mathbf{z}, \mathbf{w}) \sim \text{gamma}(\alpha + |\ell|, \beta - \nu). \quad (\text{S.36})$$

Hence, the inclusion of an additional step in Algorithm 1 sampling first from (S.35) and then from (S.36) allows for tractable posterior computation also when including hyperpriors in $\text{pExp}(\mathbf{V}; \theta, \theta_0)$.

S4.2 Section 4 (additional empirical results)

As discussed within Section 4, convergence of the Gibbs sampling schemes for H-DP($\theta = 0.5, \theta_0 = 4$), H-NSP($\sigma = 0.2, \sigma_0 = 0.8$), and H-DP with $\text{gamma}(5, 10)$ and $\text{gamma}(12, 3)$ hyperpriors for θ and θ_0 is assessed from the analysis of the traceplots for the logarithm of the likelihood in (1). Figure S.1 provides a graphical representation of these traceplots for one of the ten replicated experiments considered in Section 4, under both Scenario 1 and 2. The results in Figure S.1 suggest rapid convergence and ad-




Figure S.1: Traceplots of $\log p(\mathbf{Y} \mid \mathbf{z}^{(s)})$, $s = 1, \dots, n_{\text{iter}}$ for H-DP, H-NSP, and H-DP with hyperpriors on θ and θ_0 , in one of the ten replicated experiments considered in Section 4, under both Scenario 1 and 2.

Table S.1: Robustness assessments for three key examples of pEx-SBMs in Scenarios 1–2: Posterior mean $\mathbb{E}[\text{VI}(\mathbf{z}, \mathbf{z}_0) \mid \mathbf{Y}]$ of the VI distance from the true \mathbf{z}_0 , when changing the initialization of the collapsed Gibbs samplers and the hyperparameter values w.r.t. the default settings in Section 4. Values within brackets are the posterior means $\mathbb{E}[\text{VI}(\mathbf{z}, \mathbf{z}_0) \mid \mathbf{Y}]$ obtained under the default settings in the main article. Results are averaged over ten replicated experiments.

SCENARIO	initialization $\mathbb{E}[\text{VI}(\mathbf{z}, \mathbf{z}_0) \mid \mathbf{Y}]$		hyperparameters $\mathbb{E}[\text{VI}(\mathbf{z}, \mathbf{z}_0) \mid \mathbf{Y}]$	
	1	2	1	2
pEx-SBM (H-DP)	0.03 (0.02)	0.62 (0.62)	0.07 (0.02)	0.67 (0.62)
pEx-SBM (H-NSP)	0.04 (0.04)	0.61 (0.60)	0.07 (0.04)	0.61 (0.60)
pEx-SBM (H-DP hyp)	0.03 (0.03)	0.62 (0.63)	0.05 (0.03)	0.63 (0.63)

equate mixing in both Scenario 1 and 2, for the three pEx-SBM examples under analysis. The higher variability of the traceplots for Scenario 2 correctly unveils the increased posterior uncertainty on \mathbf{z} arising from such a more challenging scenario with reduced separation among the different groups.

Table S.1 quantifies the robustness of the three pEx-SBMs examples analyzed in Section 4, with a focus on the initialization of the Gibbs sampling routine in Section 3.1 and hyperparameters specification. Recalling Algorithm 1, as a default setting for the initialization of \mathbf{z} we consider V different groups, each comprising a single node. Table S.1 clarifies that the inferred posterior concentration around \mathbf{z}_0 in replicated studies does not change when initializing Algorithm 1 to the other extreme setting characterized by a single group containing all the V nodes in the network. Similar robustness is observed also when varying the hyperparameters of the three priors under analysis. In Table S.1 we assess, in particular, an H-DP($\theta = 2, \theta_0 = 7$), an H-NSP($\sigma = 0.35, \sigma_0 = 0.9$), and an H-DP with gamma(12, 6) and gamma(14, 2) hyperpriors for θ and θ_0 . Although these alternative hyperparameter settings induce a prior expected number of groups of ≈ 10 (i.e., twice the one obtained under the prior specifications in Section 4), as shown in Table S.1 the overall posterior concentration around \mathbf{z}_0 in replicated studies remains almost the same as the one achieved under the original settings, with H-DP based on hyperpriors showcasing the improved robustness.

The slightly improved robustness of H-DP with hyperprior can be explained by its ability to learn the parameters that control the clustering properties and the strength of the layer information in informing inference on \mathbf{z} . Such an advantage is also observed when studying the performance of the three pEx-SBMs examples in Scenario 1 under a non-informative layer division obtained by considering a random permutation of the original layer labels in Section 4 (see also Figure 2). By performing poste-

rior inference in this context under the default settings from Section 4 for the three pEx-SBMs examples yields a $VI(\hat{\mathbf{z}}, \mathbf{z}_0)$ distance, averaged over ten replicated studies, of 0.16 for H-DP and H-NSP, while H-DP with hyperprior achieves a slightly improved robustness with a value of 0.14. As expected, when layers are non-informative, performance slightly deteriorates relative to the results displayed in Table 1 within the main article. Nonetheless, a closer inspection of Table 1 highlights that, even in this highly challenging setting, pEx-SBMs still display higher accuracy than the one achieved by state-of-the-art and routinely-implemented methods (e.g., Louvain (Blondel et al., 2008), JCDC (Zhang et al., 2016) and CASC (Binkiewicz et al., 2017)) supervised by the original informative layer division. Considering non-informative layers for CASC yields a $VI(\hat{\mathbf{z}}, \mathbf{z}_0)$ distance, averaged over ten replicated studies, of 2.93, which is orders of magnitude higher than the one achieved under the three pEx-SBMs examples. This suggest that pEx-SBM is less sensitive to non-informative layers than state-of-the-art methods.

S4.3 Section 5 (additional empirical results)

Figure S.2 complements the predictive studies on the criminal network in Section 5 of the main article with a focus on the estimated posterior similarity matrix comprising both in-sample and predictive co-




Figure S.2: For the criminal network application in Section 5, graphical representation of the posterior similarity matrix comprising in-sample and predictive co-clustering probabilities. The two gray lines separate in-sample nodes (larger group) and out-of-sample ones (smaller group). Side colors correspond to *locale* affiliation (i.e., layers), whereas the color of each entry in the matrix ranges from white to black as the estimated co-clustering probability goes from 0 to 1.

clustering probabilities. The results nicely illustrate one of the important advantages of the proposed pEx-SBM. Namely its ability to properly quantify in-sample and predictive clustering uncertainty. This is evident in the higher co-clustering uncertainty for pairs of nodes involving at least an out-of-sample one, whereas those comprising both in-sample nodes display, as expected, lower variability. Interestingly, the higher heterogeneity is observed for the co-clustering probabilities involving out-of-sample nodes from the green *locale*. As discussed in Section 5 (see also Figure 4), such a *locale* is the most fragmented in different groups. This structure properly translates into a higher co-clustering uncertainty in Figure S.2 for out-of-sample nodes from the green *locale*.

S5 Glossary

The table below provides a glossary of the main quantities involved in our adaptation of the *Chinese restaurant franchise* (CRF) metaphor to the pEx-SBM construction.

QUANTITY	DESCRIPTION
V	total number of nodes in the network
d	total number of layers in the network
H	total number of sociability profiles (i.e., clusters) in the network
V_j	number of nodes in layer j
w_{ji}	label of the subgroup to which node i in layer j has been assigned
ℓ_{jh}	number of subgroups in layer j with sociability profile h
q_{jht}	number of nodes in layer j assigned to the t -th subgroup with sociability profile h
$\ell_{j\cdot}$	number of subgroups in layer j
$\ell_{\cdot h}$	total number of subgroups with sociability profile h
n_{jh}	number of nodes in layer j with sociability profile h
$\Phi_{\ell_{j\cdot}, j}^{(V_j)}$	EPPF regulating the distribution of the division in subgroups within layer j
$\Phi_{H,0}^{(\ell)}$	EPPF driving the sociability profile assignment to the subgroups in the different layers

References

- Abbe, E. (2017), “Community detection and stochastic block models: recent developments,” *Journal of Machine Learning Research*, 18, 6446–6531.
- Amini, A. A., Paez, M. S., and Lin, L. (2024), “Hierarchical stochastic block model for community detection in multiplex networks,” *Bayesian Analysis*, 19, 319–345.
- Arroyo, J., Athreya, A., Cape, J., Chen, G., Priebe, C. E., and Vogelstein, J. T. (2021), “Inference for multiple heterogeneous networks with common invariant subspace,” *Journal of Machine Learning Research*, 22, 1–49.

- Barbillon, P., Donnet, S., Lazega, E., and Bar-Hen, A. (2017), “Stochastic block models for multiplex networks: an application to a multilevel network of researchers,” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180, 295–314.
- Binkiewicz, N., Vogelstein, J. T., and Rohe, K. (2017), “Covariate-assisted spectral clustering,” *Biometrika*, 104, 361–377.
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008), “Fast unfolding of communities in large networks,” *Journal of Statistical Mechanics: Theory and Experiment*, 2008, P10008.
- Briatte, F. (2016), “Network patterns of legislative collaboration in twenty parliaments,” *Network Science*, 4, 266–271.
- Bullmore, E., and Sporns, O. (2009), “Complex brain networks: graph theoretical analysis of structural and functional systems,” *Nature Reviews Neuroscience*, 10, 186–198.
- Calderoni, F., Brunetto, D., and Piccardi, C. (2017), “Communities in criminal networks: A case study,” *Social Networks*, 48, 116–125.
- Camerlenghi, F., Lijoi, A., Orbanz, P., and Prünster, I. (2019), “Distribution theory for hierarchical processes,” *The Annals of Statistics*, 47, 67–92.
- Catalano, M., De Blasi, P., Lijoi, A.— (2022), “Posterior asymptotics for boosted hierarchical Dirichlet process mixtures,” *Journal of Machine Learning Research*, 23, 1–23.
- Côme, E., Jouvin, N., Latouche, P., and Bouveyron, C. (2021), “Hierarchical clustering with discrete latent variable models and the integrated classification likelihood,” *Advances in Data Analysis and Classification*, 15, 957–986.
- Daley, D., and Vere-Jones, D. (2007), *An Introduction to the Theory of Point Processes: Volume II: General Theory and Structure*, Probability and Its Applications, Springer New York.
- de Finetti, B. (1937), “La prévision: ses lois logiques, ses sources subjectives,” in *Annales de l’institut Henri Poincaré*, vol. 17, pp. 1–68.
- (1938), “Sur la condition d’équivalence partielle,” *Atualités Scientifiques et Industrielles*, 5–18.
- Dempsey, W., Oselio, B., and Hero, A. (2022), “Hierarchical network models for exchangeable structured interaction processes,” *Journal of the American Statistical Association*, 117, 2056–2073.
- Durante, D., Dunson, D. B., and Vogelstein, J. T. (2017a), “Nonparametric Bayes modeling of populations of networks,” *Journal of the American Statistical Association*, 112, 1516–1530.
- Durante, D., Paganin, S., Scarpa, B., and Dunson, D. B. (2017b), “Bayesian modelling of networks in complex business intelligence problems,” *Journal of the Royal Statistical Society. Series C*, 66, 555–580.
- Escobar, M. D., and West, M. (1995), “Bayesian density estimation and inference using mixtures,” *Journal of the American Statistical Association*, 90, 577–588.
- Ferguson, T. S. (1973), “A Bayesian analysis of some nonparametric problems,” *The Annals of Statistics*, 1, 209 – 230.
- Fortunato, S., and Hric, D. (2016), “Community detection in networks: A user guide,” *Physics Reports*, 659, 1–44.
- Franzolini, B. (2022), *On Dependent Processes in Bayesian Nonparametrics*, PhD Thesis, Bocconi University.
- Fruchterman, T. M., and Reingold, E. M. (1991), “Graph drawing by force-directed placement,” *Software: Practice and Experience*, 21, 1129–1164.
- Gao, L. L., Witten, D., and Bien, J. (2022), “Testing for association in multiview network data,” *Biometrics*, 78, 1018–1030.
- Gelman, A., Hwang, J., and Vehtari, A. (2014), “Understanding predictive information criteria for Bayesian

- models,” *Statistics and Computing*, 24, 997–1016.
- Geng, J., Bhattacharya, A., and Pati, D. (2019), “Probabilistic community detection with unknown number of communities,” *Journal of the American Statistical Association*, 114, 893–905.
- Girvan, M., and Newman, M. E. (2002), “Community structure in social and biological networks,” *Proceedings of the National Academy of Sciences*, 99, 7821–7826.
- Handcock, M. S., Raftery, A. E., and Tantrum, J. M. (2007), “Model-based clustering for social networks,” *Journal of the Royal Statistical Society: Series A*, 170, 301–354.
- Holland, P. W., Laskey, K. B., and Leinhardt, S. (1983), “Stochastic blockmodels: First steps,” *Social Networks*, 5, 109–137.
- James, L. F., Lijoi, A., and Prünster, I. (2009), “Posterior analysis for normalized random measures with independent increments,” *Scandinavian Journal of Statistics*, 36, 76–97.
- Jing, B.-Y., Li, T., Lyu, Z., and Xia, D. (2021), “Community detection on mixture multilayer networks via regularized tensor decomposition,” *The Annals of Statistics*, 49, 3181–3205.
- Josephs, N., Amini, A. A., Paez, M., and Lin, L. (2023), “Nested stochastic block model for simultaneously clustering networks and nodes,” *arXiv preprint arXiv:2307.09210*.
- Kemp, C., Tenenbaum, J. B., Griffiths, T. L., Yamada, T., and Ueda, N. (2006), “Learning systems of concepts with an infinite relational model,” in *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1*, pp. 381–388.
- Kingman, J. F. C. (1967), “Completely random measures,” *Pacific Journal of Mathematics*, 21, 59 – 78.
- (1975), “Random discrete distributions,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 37, 1–22.
- (1978), “The representation of partition structures,” *Journal of London Mathematical Society*, s2-18, 374–380.
- (1993), *Poisson processes*, Oxford Studies in Probability, Oxford University Press.
- Kivelä, M., Arenas, A., Barthélemy, M., Gleeson, J. P., Moreno, Y., and Porter, M. A. (2014), “Multilayer networks,” *Journal of Complex Networks*, 2, 203–271.
- Lee, C., and Wilkinson, D. J. (2019), “A review of stochastic block models and extensions for graph clustering,” *Applied Network Science*, 4, 1–50.
- Legramanti, S., Rigon, T., Durante, D., and Dunson, D. B. (2022), “Extended stochastic block models with application to criminal networks,” *The Annals of Applied Statistics*, 16, 2369–2395.
- Lei, J., Chen, K., and Lynch, B. (2020), “Consistent community detection in multi-layer network data,” *Biometrika*, 107, 61–73.
- Lijoi, A., and Prünster, I. (2010), in *Bayesian nonparametrics*, eds. Hjort, N., Holmes, C., Müller, P., and Walker, S., Cambridge: Cambridge University Press, pp. 80–136.
- MacDonald, P. W., Levina, E., and Zhu, J. (2022), “Latent space models for multiplex networks with shared structure,” *Biometrika*, 109, 683–706.
- Meilä, M. (2007), “Comparing clusterings — an information based distance,” *Journal of Multivariate Analysis*, 98, 873–895.
- Mele, A., Hao, L., Cape, J., and Priebe, C. E. (2022), “Spectral estimation of large stochastic blockmodels with discrete nodal covariates,” *Journal of Business & Economic Statistics*, 1–13.
- Mucha, P. J., Richardson, T., Macon, K., Porter, M. A., and Onnela, J.-P. (2010), “Community structure in time-dependent, multiscale, and multiplex networks,” *Science*, 328, 876–878.
- Newman, M. E., and Clauset, A. (2016), “Structure and inference in annotated networks,” *Nature communi-*

- cations, 7, 1–11.
- Noroozi, M., and Pensky, M. (2024), “Sparse subspace clustering in diverse multiplex network model,” *Journal of Multivariate Analysis*, 203, 105333.
- Nowicki, K., and Snijders, T. A. B. (2001), “Estimation and prediction for stochastic blockstructures,” *Journal of the American Statistical Association*, 96, 1077–1087.
- Paul, S., and Chen, Y. (2020), “Spectral and matrix factorization methods for consistent community detection in multi-layer networks,” *The Annals of Statistics*, 48, 230–250.
- Pensky, M., and Wang, Y. (2024), “Clustering of diverse multiplex networks,” *IEEE Transactions on Network Science and Engineering*, 11, 3441–3454.
- Pitman, J. (1995), “Exchangeable and partially exchangeable random partitions,” *Probability Theory and Related Fields*, 102, 145–158.
- Rebaudo, G., Lin, Q., and Mueller, P. (2024), “Separate exchangeability as modeling principle in Bayesian nonparametrics,” *arXiv preprint arXiv:2112.07755*.
- Regazzini, E., Lijoi, A., and Prünster, I. (2003), “Distributional results for means of normalized random measures with independent increments,” *The Annals of Statistics*, 31, 560–585.
- Schmidt, M. N., and Morup, M. (2013), “Nonparametric Bayesian modeling of complex networks: An introduction,” *IEEE Signal Processing Magazine*, 30, 110–128.
- Stanley, N., Bonacci, T., Kwitt, R., Niethammer, M., and Mucha, P. J. (2019), “Stochastic block models with multiple continuous attributes,” *Applied Network Science*, 4, 1–22.
- Stanley, N., Shai, S., Taylor, D.— (2016), “Clustering network layers with the strata multilayer stochastic block model,” *IEEE Transactions on Network Science and Engineering*, 3, 95–105.
- Sweet, T. M. (2015), “Incorporating covariates into stochastic blockmodels,” *Journal of Educational and Behavioral Statistics*, 40, 635–664.
- Tallberg, C. (2004), “A Bayesian approach to modeling stochastic blockstructures with covariates,” *Journal of Mathematical Sociology*, 29, 1–23.
- Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. (2006), “Hierarchical Dirichlet processes,” *Journal of the American Statistical Association*, 101, 1566–1581.
- Von Luxburg, U. (2007), “A tutorial on spectral clustering,” *Statistics and Computing*, 17, 395–416.
- Wade, S., and Ghahramani, Z. (2018), “Bayesian cluster analysis: Point estimation and credible balls,” *Bayesian Analysis*, 13, 559–626.
- Watanabe, S. (2010), “Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory,” *Journal of Machine Learning Research*, 11, 3571–3594.
- Wilson, J. D., Palowitch, J., Bhamidi, S., and Nobel, A. B. (2017), “Community extraction in multilayer networks with heterogeneous community structure,” *Journal of Machine Learning Research*, 18, 5458–5506.
- Xu, Z., Ke, Y., Wang, Y., Cheng, H., and Cheng, J. (2012), “A model-based approach to attributed graph clustering,” in *Proceedings of the 2012 ACM SIGMOD Conference on Management of Data*, pp. 505–516.
- Yan, B., and Sarkar, P. (2021), “Covariate regularized community detection in sparse graphs,” *Journal of the American Statistical Association*, 116, 734–745.
- Yang, J., McAuley, J., and Leskovec, J. (2013), “Community detection in networks with node attributes,” in *2013 IEEE 13th International Conference on Data Mining*, IEEE, pp. 1151–1156.
- Zhang, Y., Levina, E., and Zhu, J. (2016), “Community detection in networks with node features,” *Electronic Journal of Statistics*, 10, 3153–3178.