

Learning-augmented count-min sketches via Bayesian nonparametrics

Emanuele Dolera

University of Pavia, Italy

EMANUELE.DOLERA@UNIPV.IT

Stefano Favaro

University of Torino and Collegio Carlo Alberto, Italy

STEFANO.FAVARO@UNITO.IT

Stefano Peluchetti

Cogent Labs, Tokyo, Japan

SPELUCHETTI@COGENT.CO.JP

Editor: Erik Sudderth

Abstract

The count-min sketch (CMS) is a time and memory efficient randomized data structure that provides estimates of tokens' frequencies in a data stream of tokens, i.e. point queries, based on random hashed data. A learning-augmented version of the CMS, referred to as CMS-DP, has been proposed by Cai, Mitzenmacher and Adams (*NeurIPS* 2018), and it relies on Bayesian nonparametric (BNP) modeling of the data stream of tokens via a Dirichlet process (DP) prior, with estimates of a point query being that are obtained as mean functionals of the posterior distribution of the point query, given the hashed data. While the CMS-DP has proved to improve on some aspects of CMS, it has the major drawback of arising from a “constructive” proof that builds upon arguments that are tailored to the DP prior, namely arguments that are not usable for other nonparametric priors. In this paper, we present a “Bayesian” proof of the CMS-DP that has the main advantage of building upon arguments that are usable under the popular Pitman-Yor process (PYP) prior, which generalizes the DP prior by allowing for a more flexible tail behaviour, ranging from geometric tails to heavy power-law tails. This result leads to develop a novel learning-augmented CMS under power-law data streams, referred to as CMS-PYP, which relies on BNP modeling of the data stream of tokens via a PYP prior. Under this more general framework, we apply the arguments of the “Bayesian” proof of the CMS-DP, suitably adapted to the PYP prior, in order to compute the posterior distribution of a point query, given the hashed data. Applications to synthetic data and real textual data show that the CMS-PYP outperforms the CMS and the CMS-DP in estimating low-frequency tokens, which are known to be of critical interest in textual data, and it is competitive with respect to a variation of the CMS designed to deal with the estimation of low-frequency tokens. An extension of our BNP approach to more general queries, such as range queries, is also discussed.

Keywords: Bayesian nonparametrics; count-min sketch; Dirichlet process prior; likelihood-free estimation; Pitman-Yor process prior; point query; power-law data stream; random hashing.

1. Introduction

When processing large data streams, it is critical to represent data in compact structures that allow to efficiently extract information. Sketches form a broad class of compact random-

ized data structures that can be easily updated and queried to perform time and memory efficient estimation in large data streams of tokens. They have found many applications in machine learning (Aggarwal and Yu, 2010), security analysis (Dwork et al., 2010), natural language processing (Goya et al., 2009), computational biology (Zhang et al., 2014; Leo Elworth et al., 2020), social networks (Song et al., 2009) and games (Harrison, 2010). We refer to the monographs of Cormode et al. (2012) and Cormode and Yi (2020), and references therein, for a comprehensive and up-to-date review on sketches. A notable problem involving sketches is the estimation or recovery of the frequency of a token in the stream, typically referred to as a “point query”. The count-min sketch (CMS) of Cormode and Muthukrishnan (2005) is arguably the most popular approach to estimate point queries, and it relies on random hashing to obtain a sketched representation of the data. The CMS achieves the goal of using a compact data structure to save time and memory, while having provable guarantees on the estimated frequency through hashed data. Nevertheless, there are several aspects of the CMS that may be improved. First, the CMS provides only point estimates, although the random hashing procedure may induce substantial uncertainty in the estimation, especially for low-frequency tokens. Second, the CMS relies on the assumption of a finite universe of tokens, although it is common for large data streams to assume an unbounded number of distinct tokens. Third, often there exists an a priori knowledge on the data, and therefore it may be desirable to incorporate such a knowledge into the CMS estimates.

Learning-augmented CMSs aim at improving the CMS through the use of statistical models for better exploiting the data (Aamand et al., 2019; Hsu et al., 2019). In such a context, Cai et al. (2018) first considered a Bayesian nonparametric (BNP) approach that assumes tokens in the stream to be modeled as random samples from an unknown distribution, which is endowed with a Dirichlet process (DP) prior (Ferguson, 1973). The resulting learning-augmented CMS, which is referred to as the CMS-DP, estimates a point query through mean functionals of the posterior distribution of the point query, given the hashed data. The posterior distribution is at the core of the BNP approach, and it is derived through an intriguing “constructive” proof that exploits a restriction property and a finite-dimensional projective property of the DP. The approach of Cai et al. (2018) allows for an unknown (unbounded) number of distinct tokens in the universe and, most importantly it allows to: i) make use of the DP prior for incorporating an a priori knowledge on the data into the CMS estimates; ii) make use of the posterior distribution for assessing the uncertainty of CMS estimates. Recently, Dolera et al. (2021) showed that the “constructive” proof of Cai et al. (2018) admits a non-trivial extension to the normalized inverse-Gaussian process (NIGP) prior, which features a distinguishing power-law tail behaviour (Lijoi et al., 2005), in contrast with the geometric tail behaviour of the DP prior. This has led to the introduction of the CMS-NIGP, which is a learning-augmented CMS under power-law data streams, though with the critical limitation that it can not be tuned to the power-law degree of the data.

1.1 Our contributions

In this paper, we further investigate the BNP approach to develop learning-augmented CMSs. The peculiar interplay between the predictive distribution and the restriction prop-

erty of the DP is the cornerstone of the “constructive” proof of Cai et al. (2018). While providing an intuitive derivation of the CMS-DP, such a proof builds upon some heuristic arguments that are tailored to the DP prior. This is a critical limitation of the approach of Cai et al. (2018), especially with respect to the flexibility of incorporating an a priori knowledge on the data into the CMS estimates. Here, we present a “Bayesian” proof of the CMS-DP, that is we present a rigorous computation of the (regular) conditional distribution of a point query, given the hashed data, and we show that such a distribution coincides with the posterior distribution derived in Cai et al. (2018). Besides strengthening the BNP approach of Cai et al. (2018), our proof improves its flexibility by avoiding the use of peculiar properties of the DP, paving the way to go beyond the use of the DP prior. In this respect, nonparametric priors with power-law tail behaviour are of special interest, as power-law distributions occur in many situations of scientific interest, and they have significant consequences for the understanding of natural and social phenomena (Clauset et al., 2009). As well as being well known in natural language or textual data (Zipf, 1949; Cancho and Solé, 2020; Harald, 2001), power-law phenomena have emerged for data arising from humans’ electronic activities, e.g. website visits, emails, interactions in social networks, password innovation, tags in annotation systems and editing of webpages (Huberman and Adamic, 1999; Barabási, 2005; Muchnik et al., 2013; Tria et al., 2014; Rybski, 2016; Monechi et al., 2017).

Our “Bayesian” proof of the CMS-DP leads to extend the BNP approach of Cai et al. (2018) to the Pitman-Yor process (PYP) prior, which is arguably the most popular nonparametric prior with power-law tail behaviour (Pitman and Yor, 1997). The PYP is indexed by a discount parameter that controls the tail behaviour of the prior, ranging from geometric tails to heavy power-law tails, and including the tail behaviour of the NIGP prior. The PYP has neither a restriction property nor a finite-dimensional projective property analogous to that of the DP, and hence: i) we apply the arguments developed for the “Bayesian” proof of the CMS-DP, suitably adapted to the PYP, to compute the posterior distribution of a point query, given the hashed data; ii) we introduce a likelihood-free approach, which relies on the minimum Wasserstein distance method (Bernton et al., 2019), to estimate the prior’s parameters. This leads to a learning-augmented CMS under power-law data, which is referred to as CMS-PYP. Besides generalizing the CMS-DP to power-law data streams, the CMS-PYP improves remarkably the CMS-NIGP, as it can be tuned to the power-law degree of the data through the discount parameter. Applications to synthetic and real data show that the CMS-PYP outperforms both the CMS and the CMS-DP in the estimation of low-frequency tokens, which are of interest in textual data (Pitel and Fouquier, 2015); the CMS-PYP also outperforms the CMS-NIGP for data with heavier power-law tails, whereas it is competitive with the CMS-NIGP for data with lighter power-law tails. In general, we show that CMS-PYP is competitive with respect to the count-mean-min (CMM) of Goyal et al. (2012), which is a variation of the CMS designed for the estimation of low-frequency tokens, and also with respect to the bootstrap-debiased-count-min (BDCM) of Ting (2018).

1.2 Organization of the paper

The paper is structured as follows. In Section 2 we briefly review the CMS-DP and its “constructive” proof, and then we present our “Bayesian” proof of the CMS-DP. In Section 3 we develop the CMS-PYP through the computation of posterior distribution of a point query, given the hashed data, and the estimation of the prior’s parameters. Section 4 contains a numerical illustration of the CMS-PYP, both on synthetic data and real data. In Section 5 we discuss our work, as well as its extension to the problem of estimating more general queries, and present some directions for future research. Proofs of our results, except for the “Bayesian” proof of the CMS-DP, and additional technical results are deferred to appendices.

2. A “Bayesian” derivation of the CMS-DP

For any $m \geq 1$ let $x_{1:m} = (x_1, \dots, x_m)$ be a stream of \mathcal{V} -valued tokens, with \mathcal{V} being a space of types (symbols). Assuming $x_{1:m}$ to be available through summaries obtained by its random hashing, the goal is to estimate, or recover, the frequency of a new token x_{m+1} in $x_{1:m}$, i.e.

$$f_{x_{m+1}} = \sum_{i=1}^m \mathbb{1}_{\{x_i\}}(x_{m+1}).$$

The CMS (Cormode and Muthukrishnan, 2005) is the most popular approach to estimate the point query $f_{x_{m+1}}$. For positive integers J and N such that $[J] = \{1, \dots, J\}$ and $[N] = \{1, \dots, N\}$, let h_1, \dots, h_N , with $h_n : \mathcal{V} \rightarrow [J]$, be random hash functions that are i.i.d. according to a pairwise independent hash family \mathcal{H} . That is, $h \in \mathcal{H}$ is such that for all $v_1, v_2 \in \mathcal{V}$, with $v_1 \neq v_2$, the probability that v_1 and v_2 hash to any j_1 and j_2 , respectively, is $\Pr[h(v_1) = j_1, h(v_2) = j_2] = J^{-2}$. Pairwise independence is typically known as strong universality, and it implies uniformity, i.e. $\Pr[h(v) = j] = J^{-1}$ for any $j \in [J]$ (Cormode and Yi, 2020, Chapter 3). Hashing $x_{1:m}$ through h_1, \dots, h_N creates N vectors of J buckets, say $\{(C_{n,1}, \dots, C_{n,J})\}_{n \in [N]}$, as follows: $C_{n,j}$ is initialized at zero, and whenever a new token x_i with $h_n(x_i) = j$ is observed we set $C_{n,j} \leftarrow 1 + C_{n,j}$ for every $n \in [N]$. The CMS estimates $f_{x_{m+1}}$ by

$$\hat{f}^{(\text{CMS})} = \min\{C_{1,h_1(x_{m+1})}, \dots, C_{N,h_N(x_{m+1})}\}. \quad (1)$$

We refer to Cormode and Yi (2020, Chapter 3) for a detailed account on the CMS and generalizations thereof dealing with general small summaries for big data. In this section, we consider the CMS-DP (Cai et al., 2018), which is a learning-augmented version of the CMS that relies on BNP modeling of the stream $x_{1:m}$ through a DP prior. We briefly review the CMS-DP and its “constructive” proof, and then we present our “Bayesian” proof of the CMS-DP.

2.1 The CMS-DP and its “constructive” proof

2.1.1 THE DP PRIOR

A simple and intuitive definition of the DP follows from its stick-breaking construction (Ferguson, 1973; Sethuraman, 1994). For $\theta > 0$ let: i) $(B_i)_{i \geq 1}$ be random variables i.i.d. as a Beta distribution with parameter $(1, \theta)$; ii) $(V_i)_{i \geq 1}$ be random variables independent

of $(B_i)_{i \geq 1}$, and i.i.d. from a non-atomic distribution ν on \mathcal{V} . Then, define $P_1 = B_1$ and $P_j = B_j \prod_{1 \leq i \leq j-1} (1 - B_i)$ for $j \geq 2$, in such a way that $\sum_{i \geq 1} P_i = 1$ almost surely. The (discrete) random probability measure $P = \sum_{j \geq 1} P_j \delta_{V_j}$ is a DP on \mathcal{V} with (base) distribution ν and mass parameter θ . The law of P thus provides a prior distribution on the space of discrete distributions on \mathcal{V} . For short, $P \sim \text{DP}(\theta; \nu)$. See Ghosal and van der Vaart (2017) and references therein for a comprehensive account of the DP, including its definition in terms of the normalization of a Gamma completely random measure. For our work, it is useful to recall the restriction property and the finite-dimensional projective property of the DP (Ferguson, 1973; Regazzini, 2001). The restriction property is stated as follows: if $A \subset \mathcal{V}$ and P_A is the random probability measure on A induced by $P \sim \text{DP}(\theta; \nu)$ on \mathcal{V} , i.e. the renormalized restriction of P to A , then $P_A \sim \text{DP}(\theta \nu(A); \nu_A / \nu(A))$, where ν_A is the restriction of the measure ν to A . The finite-dimensional projective property is stated as follows: if $\{B_1, \dots, B_k\}$ is a measurable k -partition of \mathcal{V} , for $k \geq 1$, then $P \sim \text{DP}(\theta; \nu)$ is such that $(P(B_1), \dots, P(B_k))$ is distributed as a Dirichlet distribution with parameter $(\theta \nu(B_1), \dots, \theta \nu(B_k))$.

Because of the discreteness of $P \sim \text{DP}(\theta; \nu)$, a random sample $X_{1:m} = (X_1, \dots, X_m)$ from P induces a random partition of the set $\{1, \dots, m\}$ into $1 \leq K_m \leq m$ partition subsets, labelled by distinct types $\mathbf{v} = \{v_1, \dots, v_{K_m}\}$, with corresponding frequencies $(N_{1,m}, \dots, N_{K_m,m})$ such that $1 \leq N_{i,m} \leq m$ and $\sum_{1 \leq i \leq K_m} N_{i,m} = m$. For $1 \leq l \leq m$ let $M_{l,m}$ be the number of distinct types with frequency l , i.e. $M_{l,m} = \sum_{1 \leq i \leq K_m} \mathbb{1}_{\{N_{i,m}\}}(l)$ such that $\sum_{1 \leq l \leq m} M_{l,m} = K_m$ and $\sum_{1 \leq l \leq m} l M_{l,m} = m$. The distribution of $\mathbf{M}_m = (M_{1,m}, \dots, M_{m,m})$ is defined on $\mathcal{M}_{m,k} = \{(m_1, \dots, m_n) : m_l \geq 0, \sum_{1 \leq l \leq m} m_l = k, \sum_{1 \leq l \leq m} l m_l = m\}$, such that

$$\Pr[\mathbf{M}_m = \mathbf{m}] = \frac{\theta^k}{(\theta)_{(m)}} m! \prod_{i=1}^m \frac{1}{i^{m_i} m_i!} \mathbb{1}_{\mathcal{M}_{m,k}}(\mathbf{m}), \quad (2)$$

where $(a)_{(n)}$ denotes the rising factorial of a of order n , i.e. $(a)_{(n)} = \prod_{0 \leq i \leq n-1} (a + i)$. See (Pitman, 2006, Chapter 3), and references therein, for details on the sampling formula (2). Let $\mathbf{v}_l = \{v_i \in \mathbf{v} : N_{i,m} = l\}$, i.e. the labels of types with frequency l , and let $\mathbf{v}_0 = \mathcal{V} - \mathbf{v}$, i.e. the labels in of types not belonging to \mathbf{v} . The predictive distribution induced by $P \sim \text{DP}(\theta; \nu)$ is

$$\Pr[X_{m+1} \in \mathbf{v}_l \mid X_{1:m}] = \Pr[X_{m+1} \in \mathbf{v}_l \mid \mathbf{M}_m = \mathbf{m}] = \begin{cases} \frac{\theta}{\theta+m} & \text{if } l = 0 \\ \frac{l m_l}{\theta+m} & \text{if } l \geq 1, \end{cases} \quad (3)$$

for $m \geq 1$. The DP prior is characterized as the sole (discrete) nonparametric prior for which: i) the conditional probability that X_{m+1} belongs to \mathbf{v}_0 , given $X_{1:m}$, depends on $X_{1:m}$ only through m ; ii) the conditional probability that X_{m+1} belongs to \mathbf{v}_l , given $X_{1:m}$, depends on $X_{1:m}$ only through m and $M_{l,m}$. Such a characterization is typically referred to as the ‘‘sufficientness’’ postulate of the DP (Regazzini, 1978; Zabell, 1997; Bacallado et al., 2017).

2.1.2 THE CMS-DP

The CMS-DP of Cai et al. (2018) assumes that the stream $x_{1:m}$ is modeled as a random sample $X_{1:m}$ from an unknown discrete distribution P , which is endowed with a DP prior.

That is,

$$\begin{aligned} X_{1:m} | P &\stackrel{\text{iid}}{\sim} P \\ P &\sim \text{DP}(\theta; \nu) \end{aligned} \tag{4}$$

for $m \geq 1$. Let h_1, \dots, h_N be a collection of random hash functions that are i.i.d. from the strong universal family \mathcal{H} , and assume that h_1, \dots, h_N are independent of $X_{1:m}$ for any $m \geq 1$; in particular, by de Finetti's representation theorem, it holds that h_1, \dots, h_N are independent of $P \sim \text{DP}(\theta; \nu)$. Under the CMS-DP the X_i 's are hashed through h_1, \dots, h_N , thus creating $\{(C_{n,1}, \dots, C_{n,J})\}_{n \in [N]}$, and estimates of the point query $f_{X_{m+1}}$, with X_{m+1} being of an arbitrary type $v \in \mathcal{V}$, are obtained as suitable mean functionals of the posterior distribution of $f_{X_{m+1}}$ given $\{C_{n,h_n(X_{m+1})}\}_{n \in [N]}$. Cai et al. (2018) provided an intriguing “constructive” derivation of such a posterior distribution, which relies on two main arguments:

- A1) because of the strong universality of \mathcal{H} and the independence between h_n and $X_{1:m}$, the restriction property of the DP, in combination with the “sufficientness” postulate of the DP, implies that the tokens X_i 's hashed in the j -th bucket $C_{n,j}$ constitute random samples from a DP with scaled mass parameter θ/J , for any $j \in [J]$ and $n \in [N]$;
- A2) because of the strong universality of \mathcal{H} , the finite-dimensional projective property of the DP implies that the vector of hashed frequencies $\mathbf{C}_n = (C_{n,1}, \dots, C_{n,J})$ is distributed as a Dirichlet-Multinomial distribution with parameter $(\theta/J, \dots, \theta/J)$, for any $n \in [N]$.

For a single hash function h_n , the main result of Cai et al. (2018) may be summarized as follows. A random sample $X_{1:m}$ from $P \sim \text{DP}(\theta; \nu)$ induces a random partition of $\{1, \dots, m\}$ into subsets labelled by $\mathbf{v} \in \mathcal{V}$, and (3) is the posterior distribution, given $X_{1:m}$, over which subset X_{m+1} joins. The frequency of that subset is the point query $f_{X_{m+1}}$ we seek to estimate, i.e.

$$\Pr[f_{X_{m+1}} = l | X_{1:m}] = \Pr[X_{m+1} \in \mathbf{v}_l | X_{1:m}] \tag{5}$$

for $l = 0, 1, \dots, m$. However, we are assuming that the sampling information $X_{1:m}$ is available only through $\{C_{n,h_n(X_{m+1})}\}_{n \in [N]}$, and hence the posterior distribution (5) is not of interest in itself. Instead, it is of interest the distribution of $f_{X_{m+1}}$, which is obtained from (5) by marginalizing out $X_{1:m}$. By combining (5) with (2) (Cai et al., 2018, Section 3), it holds that

$$p_{f_{X_{m+1}}}(l; m, \theta) := \Pr[f_{X_{m+1}} = l] = \theta \frac{(m-l+1)_{(l)}}{(\theta+m-l)_{(l+1)}}. \tag{6}$$

For any $n \in [N]$, strong universality of \mathcal{H} and independence between h_n and $X_{1:m}$ imply that h_n induces a (fixed) J -partition of \mathcal{V} , say $\{B_{h_n,1}, \dots, B_{h_n,J}\}$, and the measure with respect to $P \sim \text{DP}(\theta; \nu)$ of each $B_{h_n,j}$ is $1/J$. Therefore, according to A1), h_n turns $P \sim \text{DP}(\theta; \nu)$ into J bucket-specific DPs, say $P_j \sim \text{DP}(\theta/J; J\nu_{B_{h_n,j}})$ for $j = 1, \dots, J$, such that P_j governs

the distribution of the sole X_i 's hashed in $B_{h_n, j}$. For any $l = 0, 1, \dots, c_n$, Cai et al. (2018) thus set

$$\Pr[f_{X_{m+1}} = l \mid C_{n, h_n(X_{m+1})} = c_n] = p_{f_{X_{m+1}}} \left(l; c_n, \frac{\theta}{J} \right). \quad (7)$$

This is an heuristic assignment, namely the left-hand side of (7) is not obtained through a rigorous computation of the (regular) conditional distribution of $f_{X_{m+1}}$ given $C_{n, h_n(X_{m+1})}$. We refer to such a derivation as the ‘‘constructive’’ proof of the posterior distribution of $f_{X_{m+1}}$ given $C_{n, h_n(X_{m+1})}$.

For the collection of hash functions h_1, \dots, h_N , the posterior distribution of $f_{X_{m+1}}$, given $\{C_{n, h_n(X_{m+1})}\}_{n \in [N]}$, follows from Equation (7) by means of the assumption that the h_n 's are i.i.d. according to the strong universal family \mathcal{H} . In particular, by a direct application of Bayes theorem, Cai et al. (2018, Section 3), showed that for $l = 0, 1, \dots, \min\{c_1, \dots, c_N\}$ it holds that

$$\Pr[f_{X_{m+1}} = l \mid \{C_{n, h_n(X_{m+1})}\}_{n \in [N]} = \{c_n\}_{n \in [N]}] = \frac{\prod_{n \in [N]} p_{f_{X_{m+1}}} \left(l; c_n, \frac{\theta}{J} \right)}{(p_{f_{X_{m+1}}} \left(l; m, \theta \right))^{N-1}}. \quad (8)$$

CMS-DP estimates of the point query $f_{X_{m+1}}$, with respect to a suitable choice of a loss function, are obtained as functionals of the posterior distribution (8), e.g. posterior mode, mean and median. We refer to Cai et al. (2018) for a detailed discussion on BNP estimators of $f_{X_{m+1}}$ and their interplay with the CMS. For a concrete application of (8), it remains to estimate the unknown prior's parameter $\theta > 0$, and this follows from A2). In particular, \mathbf{C}_n is distributed as a Dirichlet-Multinomial distribution with parameter $(\theta/J, \dots, \theta/J)$, and the distribution of $\{\mathbf{C}_n\}_{n \in [N]}$ follows by the assumption that the h_n 's are i.i.d. from \mathcal{H} , that is

$$\Pr[\{\mathbf{C}_n\}_{n \in [N]} = \{\mathbf{c}_n\}_{n \in [N]}] = \prod_{n \in [N]} \frac{m!}{(\theta)_{(m)}} \prod_{j=1}^J \frac{\binom{\theta}{J} (c_{n,j})}{c_{n,j}!}. \quad (9)$$

Equation (9) provides the (marginal) likelihood function of $\{\mathbf{c}_n\}_{n \in [N]}$. The explicit form of such a function allows for an easy implementation of a Bayesian estimation of the prior's parameter θ . Cai et al. (2018) adopted an empirical Bayes approach, which consists in estimating θ by maximizing, with respect to θ , the likelihood function (9). A fully Bayes, or hierarchical Bayes, approach can be also applied by setting a suitable prior distribution on θ .

2.2 A ‘‘Bayesian’’ proof of the CMS-DP

In Cai et al. (2018), the interplay between the predictive distribution and the restriction property of the DP is the cornerstone for the derivation of the posterior distribution (7). In particular, the ‘‘constructive’’ proof of the CMS-DP imposes two strong constraints with respect to the choice of the prior distribution: C1) the predictive distribution induced by the prior must have a simple analytical expression, i.e. the marginalization with respect to the sampling information $X_{1:m}$ must be doable explicitly, and it must satisfy a ‘‘sufficiency’’ postulate analogous to that of the DP prior; C2) the prior distribution must have a restriction property analogous to that of the DP prior, which allows us to make use of

the distribution of $f_{X_{m+1}}$ to assign the posterior distribution of $f_{X_{m+1}}$ given $C_{n,h_n(X_{m+1})}$. Nonparametric priors obtained by normalizing (homogeneous) completely random measures (James, 2002; Prünster, 2002; Pitman, 2003; Regazzini et al., 2003; James et al., 2009) form a broad class of priors that generalize the DP prior and satisfy C2); this follows from the Poisson process representation of completely random measures, for which the Poisson coloring theorem holds true (Kingman, 1993, Chapter 5). However, the DP is the sole normalized (homogeneous) completely random measure that satisfies C1) (Regazzini, 1978). Beyond normalized completely random measures, the PYP prior is a popular generalization of the DP prior that satisfies C1). However, the PYP does not satisfy C2); this is because the PYP is not a normalized completely random measure. To the best of our knowledge, the DP prior is the sole (discrete) nonparametric prior that satisfies both C1) and C2), and hence it is the sole prior for which the “constructive” proof of Cai et al. (2018) works. The “constructive” proof thus provides a limitation for the BNP approach of Cai et al. (2018), as it implies a lack of flexibility in the choice of the prior distribution for BNP modeling of $x_{1:m}$.

Here, we present a rigorous derivation of the posterior distribution of $f_{X_{m+1}}$ given $C_{n,h_n(X_{m+1})}$, which is referred to as the “Bayesian” proof of the CMS-DP. For any $n \in [N]$, we consider the problem of computing the (regular) conditional distribution of $f_{X_{m+1}}$ given $C_{n,h_n(X_{m+1})}$, i.e.

$$\Pr[f_{X_{m+1}} = l \mid C_{h_n(X_{m+1})} = c_n] = \frac{\Pr[f_{X_{m+1}} = l, \sum_{i=1}^m \mathbb{1}_{\{h_n(X_i)\}}(h_n(X_{m+1})) = c_n]}{\Pr[\sum_{i=1}^m \mathbb{1}_{\{h_n(X_i)\}}(h_n(X_{m+1})) = c_n]}, \quad (10)$$

for $l = 0, 1, \dots, c_n$. In the next theorem, we show that the (regular) conditional distribution (10) coincides with the posterior distribution (7) obtained by means of the “constructive” proof. That is, the “Bayesian” proof and the “constructive” proof lead to the same posterior distribution. As a critical feature, our “Bayesian” proof stands out for not relying on the peculiar restriction property of the DP; instead, by exploiting the strong universality of \mathcal{H} , the “Bayesian” proof relies on evaluating the numerator and the denominator of (10) through standard combinatorial arguments and well-known distributional properties of a random sample $X_{1:m}$ from the DP (Pitman, 2003, 2006; Sangalli, 2006), i.e. marginal properties. It emerges that the “Bayesian” proof has two main advantages with respect to the “constructive” proof: i) it provides a rigorous proof of the CMS-DP, which avoids any heuristic assignment of the posterior distribution, thus strengthening the BNP approach of Cai et al. (2018); ii) it avoids the use of the peculiar restriction property of the DP, thus paving the way to the use of more general classes of prior distributions than the sole DP prior.

Theorem 1 *For $m \geq 1$, let $x_{1:m}$ be a stream of tokens that are modeled as a random sample $X_{1:m}$ from $P \sim DP(\theta; \nu)$, and let X_{m+1} be an additional random sample from P . Moreover, let h_n be a random hash function distributed as the strong universal family \mathcal{H} , and let h_n be independent of $X_{1:m}$ for any $m \geq 1$, that is h_n is independent of P . Then, for $l = 0, 1, \dots, c_n$*

$$\Pr[f_{X_{m+1}} = l \mid C_{n,h_n(X_{m+1})} = c_n] = \frac{\theta}{J} \frac{(c_n - l + 1)_{(l)}}{(J + c_n - l)_{(l+1)}}. \quad (11)$$

Proof The proof consists of three steps: i) evaluate the numerator of (10); ii) evaluate the denominator of (10); iii) evaluate (10) with respect to what obtained in step i) and step ii). First, we observe that the independence between h_n and $X_{1:m}$ allows us to invoke the “freezing lemma” (Baldi, 2017, Lemma 4.1), according to which we can treat h_n as it was fixed, i.e. non-random. To simplify the notation, we remove the subscript n from h_n and c_n . We start with the denominator of (10). Uniformity of h implies that h induces a (fixed) J -partition $\{B_1, \dots, B_J\}$ of \mathcal{V} such that $B_j = \{v \in \mathcal{V} : h(v) = j\}$ and $\nu(B_j) = J^{-1}$ for $j = 1, \dots, J$. Then, the finite-dimensional projective property of the DP implies that $P(B_j)$ is distributed as a Beta distribution with parameter $(\theta/J, \theta(1-1/J))$ for $j = 1, \dots, J$. Hence, we write

$$\begin{aligned}
 & \Pr \left[\sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] \\
 &= J \binom{m}{c} \mathbb{E}[(P(B_j))^{c+1} (1 - P(B_j))^{m-c}] \\
 &= J \binom{m}{c} \int_0^1 p^{c+1} (1-p)^{m-c} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J})\Gamma(\theta - \frac{\theta}{J})} p^{\frac{\theta}{J}-1} (1-p)^{\theta - \frac{\theta}{J}-1} dp \\
 &= J \binom{m}{c} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J})\Gamma(\theta - \frac{\theta}{J})} \frac{\Gamma(\frac{\theta}{J} + c + 1)\Gamma(\theta - \frac{\theta}{J} + m - c)}{\Gamma(\theta + m + 1)}.
 \end{aligned} \tag{12}$$

This completes the study of the denominator of (10). Now, we consider the numerator of (10). Let us define the event $B(m, l) = \{X_1 = \dots = X_l = X_{m+1}, \{X_{l+1}, \dots, X_m\} \cap \{X_{m+1}\} = \emptyset\}$. Then,

$$\begin{aligned}
 & \Pr \left[f_{X_{m+1}} = l, \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] \\
 &= \binom{m}{l} \Pr \left[B(m, l), \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] \\
 &= \binom{m}{l} \Pr \left[B(m, l), \sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right].
 \end{aligned} \tag{13}$$

That is, the distribution of $(f_{X_{m+1}}, C_j)$ is completely determined by the distribution of the random variables (X_1, \dots, X_{m+1}) . Let $\Pi(s, k)$ denote the set of all possible partitions of the set $\{1, \dots, s\}$ into k disjoint subsets π_1, \dots, π_k such that n_i is the cardinality of π_i . In particular, from Sangalli (2006, Equation 3.5), for any measurable A_1, \dots, A_{m+1} we have that

$$\Pr[X_1 \in A_1, \dots, X_{m+1} \in A_{m+1}] = \sum_{k=1}^{m+1} \frac{\theta^k}{(\theta)_{(m+1)}} \sum_{(\pi_1, \dots, \pi_k) \in \Pi(m+1, k)} \prod_{i=1}^k (n_i - 1)! \nu(\cap_{r \in \pi_i} A_r)$$

for $m \geq 1$. Let \mathcal{V} be the Borel σ -algebra of \mathcal{V} . Let $\nu_{\pi_1, \dots, \pi_k}$ be a probability measure on $(\mathcal{V}^{m+1}, \mathcal{V}^{m+1})$ defined as

$$\nu_{\pi_1, \dots, \pi_k}(A_1 \times \dots \times A_{m+1}) = \prod_{1 \leq i \leq k} \nu(\cap_{r \in \pi_i} A_r),$$

and attaching to $B(m, l)$ a value that is either 0 or 1. In particular, $\nu_{\pi_1, \dots, \pi_k}(B(m, l)) = 1$ if and only if one of the π_i 's is equal to the set $\{1, \dots, l, m+1\}$. Hence, based on the measure $\nu_{\pi_1, \dots, \pi_k}$, we write

$$\begin{aligned}
 & \Pr \left[B(m, l), \sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right] \\
 &= \sum_{k=2}^{m-l+1} \frac{\theta^k}{(\theta)_{(m+1)}} \\
 & \quad \times \sum_{(\pi_1, \dots, \pi_{k-1}) \in \Pi(m-l, k-1)} l! \prod_{i=1}^{k-1} (n_i - 1)! \nu_{\pi_1, \dots, \pi_k} \left(\sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right) \\
 &= \theta \frac{(\theta)_{(m-l)} l!}{(\theta)_{(m+1)}} \\
 & \quad \times \sum_{r=1}^{m-l} \frac{\theta^r}{(\theta)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (n_i - 1)! \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=1}^{m-l} \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right).
 \end{aligned}$$

Now,

$$\sum_{r=1}^{m-l} \frac{\theta^r}{(\theta)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (n_i - 1)! \nu_{\pi_1, \dots, \pi_r}(\cdot)$$

is the distribution of a random sample (X_1, \dots, X_{m-l}) under $P \sim \text{DP}(\theta; \nu)$. Again, the distribution of (X_1, \dots, X_{m-l}) is given in Sangalli (2006, Equation 3.5). Using the fact that $P(B_j)$ is distributed as a Beta distribution with parameter $(\theta/J, \theta(1-1/J))$, for $j = 1, \dots, J$, we write

$$\begin{aligned}
 & \Pr \left[B(m, l), \sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right] \\
 &= \theta \frac{(\theta)_{(m-l)} l!}{(\theta)_{(m+1)}} \\
 & \quad \times \sum_{r=1}^{m-l} \frac{\theta^r}{(\theta)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (n_i - 1)! \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=1}^{m-l} \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right) \\
 &= \theta \frac{(\theta)_{(m-l)} l!}{(\theta)_{(m+1)}} \binom{m-l}{c-l} \mathbb{E}[(P(B_j))^{c-l} (1 - P(B_j))^{m-c}] \\
 &= \theta \frac{(\theta)_{(m-l)} l!}{(\theta)_{(m+1)}} \binom{m-l}{c-l} \int_0^1 p^{c-l} (1-p)^{m-c} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J}) \Gamma(\theta - \frac{\theta}{J})} p^{\frac{\theta}{J}-1} (1-p)^{\theta - \frac{\theta}{J} - 1} dp \\
 &= \theta \frac{(\theta)_{(m-l)} l!}{(\theta)_{(m+1)}} \binom{m-l}{c-l} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J}) \Gamma(\theta - \frac{\theta}{J})} \frac{\Gamma(\frac{\theta}{J} + c - l) \Gamma(\theta - \frac{\theta}{J} + m - c)}{\Gamma(\theta + m - l)},
 \end{aligned}$$

where the second identity follows from an application of Sangalli (2006, Proposition 3.1) under the DP prior; see also the formule displayed at page 469 of Sangalli (2006). From

(13) we write

$$\begin{aligned}
 \Pr \left[f_{X_{m+1}} = l, \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] & \quad (14) \\
 &= \binom{m}{l} \theta \frac{(\theta)_{(m-l)}}{(\theta)_{(m+1)}} l! \binom{m-l}{c-l} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J})\Gamma(\theta-\frac{\theta}{J})} \frac{\Gamma(\frac{\theta}{J}+c-l)\Gamma(\theta-\frac{\theta}{J}+m-c)}{\Gamma(\theta+m-l)} \\
 &= \theta \frac{m!}{(c-l)!(m-c)!} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J})\Gamma(\theta-\frac{\theta}{J})} \frac{\Gamma(\frac{\theta}{J}+c-l)\Gamma(\theta-\frac{\theta}{J}+m-c)}{\Gamma(\theta+m+1)}.
 \end{aligned}$$

This completes the study of the numerator of (10). Then, by combining (14) and (12), for $l = 0, 1, \dots, m$

$$\begin{aligned}
 \Pr[f_{X_{m+1}} = l | C_{h(X_{m+1})} = c] & \quad (15) \\
 &= \frac{\theta \frac{m!}{(c-l)!(m-c)!} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J})\Gamma(\theta-\frac{\theta}{J})} \frac{\Gamma(\frac{\theta}{J}+c-l)\Gamma(\theta-\frac{\theta}{J}+m-c)}{\Gamma(\theta+m+1)}}{J \binom{m}{c} \frac{\Gamma(\theta)}{\Gamma(\frac{\theta}{J})\Gamma(\theta-\frac{\theta}{J})} \frac{\Gamma(\frac{\theta}{J}+c+1)\Gamma(\theta-\frac{\theta}{J}+m-c)}{\Gamma(\theta+m+1)}} \\
 &= \frac{\theta}{J} \frac{(c-l+1)_{(l)}}{(\frac{\theta}{J}+c-l)_{(l+1)}},
 \end{aligned}$$

which follows directly from the (regular) conditional distribution (10), whose denominator and numerator are replaced by Equation (12) and Equation (14), respectively. The proof is completed. \blacksquare

The next proposition is an interesting complement to Theorem 1: i) it provides a novel (constructive) representation of the posterior distribution (11) in terms of a mixture of Binomial distributions; ii) it characterizes the large c_m asymptotic behaviour of the posterior distribution (11). Hereafter, we denote by \xrightarrow{w} the convergence in distribution or weak convergence.

Proposition 2 *Let $B_{a,b}$ be a Beta random variable, and denote by $f_{B_{a,b}}$ its density function. Under the setting of Theorem 1, if $F_{X_{m+1}}$ is a random variable distributed as (11), then as $c_n \rightarrow +\infty$*

$$\frac{F_{X_{m+1}}}{c_n} \xrightarrow{w} B_{1, \frac{\theta}{J}}. \quad (16)$$

Moreover, if $\text{Binomial}(n, p)$ denotes the Binomial distribution with parameter (n, p) , then for $l = 0, 1, \dots, c_n$

$$\Pr[f_{X_{m+1}} = l | C_{n, h_n(X_{m+1})} = c_n] = \int_0^1 \text{Binomial}(l; c_n, p) f_{B_{1, \frac{\theta}{J}}}(p) dp. \quad (17)$$

See Appendix A for the proof of Proposition 2. Let $F_{X_{m+1}}$ be a random variable whose distribution coincides with the posterior distribution (11). Then, Equation (17) shows that the distribution of $F_{X_{m+1}}$ is a mixture of Binomial distributions, with the mixing

distribution on the success probability being a Beta distribution with parameter $(1, \theta/J)$. That is, the posterior distribution (11) admits a straightforward representation in terms of a Beta-Binomial distribution with parameter $(c_n, 1, \theta/J)$ (Johnson et al., 2005, Chapter 6). Moreover, Equation (16) shows that the mixing distribution is precisely the limiting distribution of the proportion $c_n^{-1}F_{X_{m+1}}$ as $c_n \rightarrow +\infty$. Then, according to Proposition 2, we write $F_{X_{m+1}} = \sum_{1 \leq i \leq c_n} Z_i$, where, by means of de Finetti’s representation theorem, $(Z_i)_{i \geq 1}$ is an exchangeable sequence of Bernoulli random variables with de Finetti’s measure being the Beta distribution with parameter $(1, \theta/J)$. Concerning large m behavior, Berry-Esseen estimates for the vicinity of the (rescaled) law of $F_{X_{m+1}}$ to the aforesaid de Finetti’s measure are contained in Dolera and Favaro (2020b), while other similar estimates can be found in Dolera (2013). Then, for a collection of hash functions h_1, \dots, h_N , one may combine Proposition 2 with Equation (8) to obtain an alternative representation, in terms of product of Beta-Binomial distributions, of the posterior distribution of $f_{X_{m+1}}$, given $\{C_{n, h_n(X_{m+1})}\}_{n \in [N]}$. Despite Proposition 2 has not a direct impact with respect to the implementation of the CMS-DP, in the sense of improving its computation, we believe it is of interest in shedding light on distributional properties of the posterior distribution of $f_{X_{m+1}}$, given $C_{h_n(X_{m+1})}$.

3. A learning-augmented CMS under power-law streams

Our “Bayesian” proof relies on less prior’s assumptions than the “constructive” proof of Cai et al. (2018), and it paves the way to extend the BNP approach of Cai et al. (2018) to the PYP prior. The PYP prior is arguably the most popular generalization of the DP prior, with a flexible tail behaviour that ranges from geometric tails to heavy power-law tails (Pitman, 2006, Chapter 3). In this section, we apply the PYP prior to develop a novel learning-augmented CMS under power-law streams of tokens. In particular, we assume the stream $x_{1:m}$ to be modeled as a random sample from an unknown discrete distribution P , which is endowed with a PYP prior \mathcal{Q} . Within the class of (discrete) nonparametric priors with power-law tail behaviour, the PYP prior stands out for both its mathematical tractability, flexibility and interpretability, and therefore it is the natural candidate for applications (De Blasi et al., 2015; Bacallado et al., 2017). We recall that the PYP has neither a restriction property nor a “sufficientness” postulate analogous to those featured by the DP, and therefore the “constructive” proof of Cai et al. (2018) cannot be applied to obtain the posterior distribution of a point query. Moreover, we recall that the PYP does not feature a finite-dimensional projective property analogous to that of the DP, and therefore prior’s parameters cannot be estimated through an empirical Bayes procedure, as discussed in Cai et al. (2018), or through a hierarchical (fully) Bayes procedure. In this section, we adapt the “Bayesian” proof of Section 2 in order to compute the posterior distribution of the point query $f_{X_{m+1}}$, given $\{C_{n, h_n(X_{m+1})}\}_{n \in [N]}$, under a PYP prior. Then, we exploit the predictive distribution of the PYP prior to implement a likelihood-free approach to estimate the PYP prior’s parameters. Our work leads to a generalization of the CMS-DP, which is referred to as the CMS-PYP, providing a novel learning-augmented CMS under power-law streams.

3.1 The CMS-PYP

3.1.1 THE PYP PRIOR

A simple and intuitive definition of the PYP follows from its stick-breaking construction (Perman et al., 1992; Pitman, 1995; Pitman and Yor, 1997). For $\alpha \in [0, 1)$ and $\theta > -\alpha$ let: i) $(B_i)_{i \geq 1}$ be independent random variables such that B_i is distributed as a Beta distribution with parameter $(1 - \alpha, \theta + i\alpha)$; ii) $(V_i)_{i \geq 1}$ random variables, independent of $(B_i)_{i \geq 1}$, and i.i.d. from a non-atomic distribution ν on \mathcal{V} . Then, define $P_1 = B_1$ and $P_j = B_j \prod_{1 \leq i \leq j-1} (1 - B_i)$ for $j \geq 2$, in such a way that $\sum_{i \geq 1} P_i = 1$ almost surely. The (discrete) random probability measure $P = \sum_{j \geq 1} P_j \delta_{V_j}$ is a PYP on \mathcal{V} with (base) distribution ν , discount parameter α and mass parameter θ . For short, we write $P \sim \text{PYP}(\alpha, \theta; \nu)$. We refer to Perman et al. (1992) and Pitman and Yor (1997) for an alternative definition of the PYP through a suitable transformation of the α -stable completely random measure Kingman (1993). See also Pitman (2006, Chapter 4) and references therein. The DP arises as a special case of the PYP by setting $\alpha = 0$. For the purposes of the present paper, it is useful to recall the power-law tail behaviour featured by the PYP prior. In particular, let $P \sim \text{PYP}(\alpha, \theta; \nu)$ with $\alpha \in (0, 1)$, and let $(P_{(j)})_{j \geq 1}$ be the decreasing ordered random probabilities P_j 's of P (Pitman, 2006, Chapter 3). Then, as $j \rightarrow +\infty$ the $P_{(j)}$'s follow a power-law distribution of exponent $c = \alpha^{-1}$ (Pitman and Yor, 1997). That is, $\alpha \in (0, 1)$ controls the power-law tail behaviour of the PYP through small probabilities $P_{(j)}$'s: the larger α the heavier the tail of P . See also Gnedin et al. (2007, Section 10) for a detailed account on the tail behaviour of the PYP prior.

As for the DP, the discreteness of $P \sim \text{PYP}(\alpha, \theta; \nu)$ implies that a random sample $X_{1:m} = (X_1, \dots, X_m)$ from P induces a random partition of the set $\{1, \dots, m\}$ into $1 \leq K_m \leq m$ partition subsets, labelled by distinct types $\mathbf{v} = \{v_1, \dots, v_{K_m}\}$, with corresponding frequencies $(N_{1,m}, \dots, N_{K_m,m})$ such that $1 \leq N_{i,m} \leq m$ and $\sum_{1 \leq i \leq K_m} N_{i,m} = m$. For $1 \leq l \leq m$ let $M_{l,m}$ be the number of distinct types with frequency l , i.e. $M_{l,m} = \sum_{1 \leq i \leq K_m} \mathbb{1}_{\{N_{i,m} = l\}}$ such that $\sum_{1 \leq l \leq m} M_{l,m} = K_m$ and $\sum_{1 \leq l \leq m} l M_{l,m} = m$. The distribution of \mathbf{M}_m is

$$\Pr[\mathbf{M}_m = \mathbf{m}] = \frac{\left(\frac{\theta}{\alpha}\right)_{(k)}}{(\theta)_{(m)}} m! \prod_{i=1}^m \left(\frac{\alpha(1-\alpha)_{(i-1)}}{i!} \right)^{m_i} \frac{1}{m_i!} \mathbb{1}_{\mathcal{M}_{m,k}}(\mathbf{m}), \quad (18)$$

such that

$$\Pr[K_m = k] = \frac{\left(\frac{\theta}{\alpha}\right)_{(k)}}{(\theta)_{(m)}} \mathcal{E}(m, k; \alpha) \quad (19)$$

for $k = 1, \dots, m$, where $\mathcal{E}(m, k; \alpha) = (k!)^{-1} \sum_{0 \leq i \leq k} \binom{k}{i} (-1)^i (-i\alpha)_{(m)}$ denotes the generalized factorial coefficient (Charalambides, 2005), with the proviso that $\mathcal{E}(0, 0; \alpha) = 1$ and $\mathcal{E}(m, 0; \alpha) = 0$. See Pitman (2006, Chapter 3) for details on (18) and (19). Now, let $\mathbf{v}_l = \{v_i \in \mathbf{v} : N_{i,m} = l\}$, i.e. the labels of types with frequency l and let $\mathbf{v}_0 = \mathcal{V} - \mathbf{v}$, i.e. the labels of types not belonging to \mathbf{v} . The predictive distribution induced by $P \sim \text{PYP}(\alpha, \theta; \nu)$ is

$$\Pr[X_{m+1} \in \mathbf{v}_l | X_{1:m}] = \Pr[X_{m+1} \in \mathbf{v}_l | \mathbf{M}_m = \mathbf{m}] = \begin{cases} \frac{\theta + k\alpha}{\theta + m} & \text{if } l = 0 \\ \frac{m_l(l-\alpha)}{\theta + m} & \text{if } l \geq 1, \end{cases} \quad (20)$$

for $m \geq 1$. In particular, the PYP prior is characterized as the sole (discrete) nonparametric prior for which: i) the conditional probability that X_{m+1} belongs to \mathbf{v}_0 , given $X_{1:m}$ depends on $X_{1:m}$ only through m and K_m ; ii) the conditional probability that X_{m+1} belongs to \mathbf{v}_l , given $X_{1:m}$, depends on $X_{1:m}$ only through m and $M_{l,m}$ (Bacallado et al., 2017, Proposition 1).

At the sampling level, the power-law tail behaviour of $P \sim \text{PYP}(\alpha, \theta; \nu)$ emerges from the analysis of the large m asymptotic behaviour of K_m and $M_{r,m}/K_m$ (Pitman, 2006, Chapter 3). Let $X_{1:m}$ be a random sample from P . Pitman (2006, Theorem 3.8) shows that, as $m \rightarrow +\infty$,

$$\frac{K_m}{m^\alpha} \xrightarrow{\text{a.s.}} S_{\alpha, \theta}^{-\alpha}, \quad (21)$$

where $S_{\alpha, \theta}$ is a polynomially tilted α -stable random variable, that is the distribution of $S_{\alpha, \theta}$ has density function $f_{S_{\alpha, \theta}}(x) \propto x^{-\theta} g_\alpha(x) \mathbb{1}_{\mathbb{R}^+}(x)$ for g_α being the positive α -stable density function. According to (21), it holds $K_m \approx m^\alpha S_{\alpha, \theta}^{-\alpha}$ for large m , or equivalently $K_m \approx [(\theta + m)^\alpha - \theta^\alpha] S_{\alpha, \theta}^{-\alpha}$ for large m (Favaro et al., 2009). It follows from (21) that, as $m \rightarrow +\infty$,

$$\frac{M_{l,m}}{K_m} \xrightarrow{\text{a.s.}} \frac{\alpha(1-\alpha)_{(l-1)}}{l!}. \quad (22)$$

Equation (21) shows that the number K_m of distinct types in $X_{1:m}$, for large m , grows as m^α . This is precisely the growth of the number of distinct types in $m \geq 1$ random samples from a power-law distribution of exponent $c = \alpha^{-1}$. Moreover, Equation (22) shows that $p_{\alpha, l} = \alpha(1-\alpha)_{(l-1)}/l!$ is the large m asymptotic proportion of the number of distinct types with frequency l in $X_{1:m}$. Then, it holds $p_{\alpha, l} \approx c_\alpha l^{-\alpha-1}$ for large l , for a constant c_α . This is precisely the distribution of the number of distinct types with frequency l in $m \geq 1$ random samples from a power-law distribution of exponent $c = \alpha^{-1}$. See Figure 1 for an illustration of the large m behaviour of K_m and $M_{l,n}$ under the PYP prior, for some choices of the parameter (α, θ) .

3.1.2 THE CMS-PYP

To introduce the CMS-PYP, we assume that the stream $x_{1:m}$ is modeled as random samples $X_{1:m}$ from an unknown discrete distribution P , which is endowed with a PYP prior. That is,

$$X_{1:m} | P \stackrel{\text{iid}}{\sim} P \quad (23)$$

$$P \sim \text{PYP}(\alpha, \theta; \nu)$$

for $m \geq 1$. Let h_1, \dots, h_N be a collection of random hash functions that are i.i.d. from the strong universal family \mathcal{H} , and assume that h_1, \dots, h_N are independent of $X_{1:m}$ for any $m \geq 1$; in particular, by de Finetti's representation theorem, h_1, \dots, h_N are independent of $P \sim \text{PYP}(\alpha, \theta; \nu)$. Under the CMS-PYP the X_i 's are hashed through h_1, \dots, h_N , thus creating $\{(C_{n,1}, \dots, C_{n,J})\}_{n \in [N]}$, and estimates of the point query $f_{X_{m+1}}$, with X_{m+1} being of an arbitrary type $v \in \mathcal{V}$, are obtained as functionals of the posterior distribution of

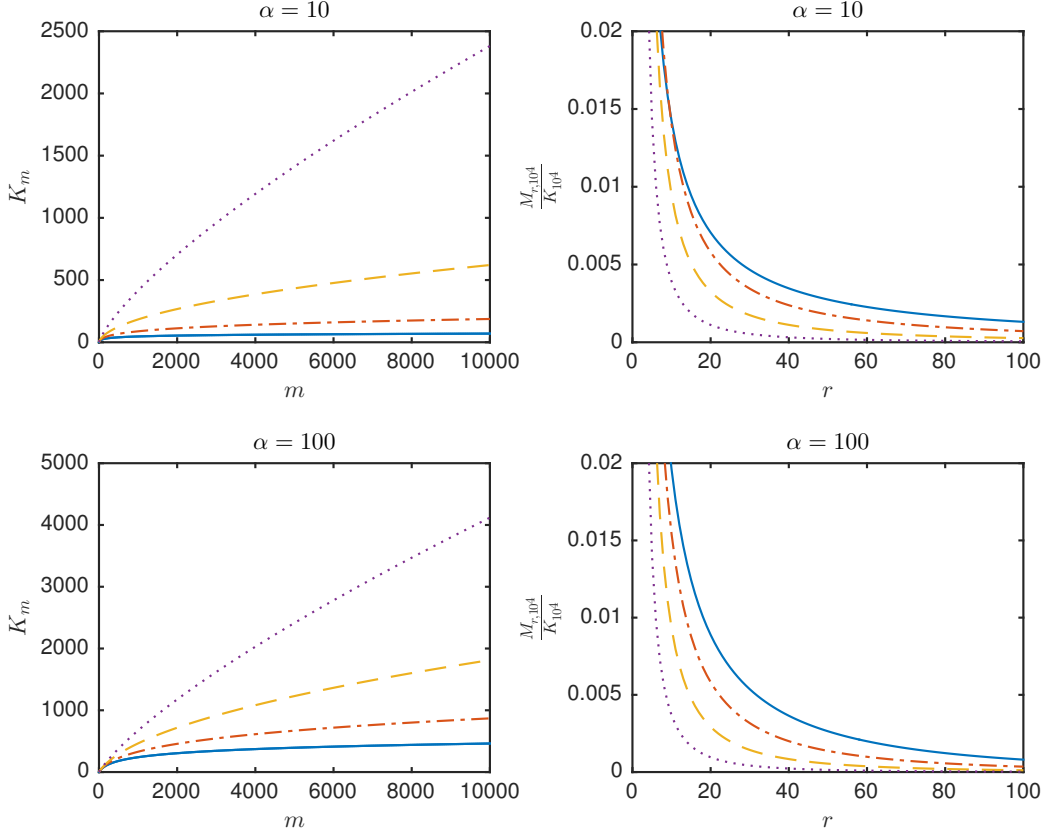


Figure 1: Behaviours in $m \geq 1$ of the statistics K_m and $M_{r,m}/K_m$, for $1 \leq m \leq 10^4$ under $P \sim \text{PYP}(\alpha, \theta; \nu)$: $\alpha = 0$ (blue -), $\alpha = .25$ (red -), $\alpha = .5$ (yellow -) and $\alpha = .75$ (purple :).

$f_{X_{m+1}}$ given $\{C_{n,h_n(X_{m+1})}\}_{n \in [N]}$. As for the derivation of the CMS-DP in Section 2, the assumption of independence between the h_n 's and $X_{1:m}$ plays a critical role to obtain the posterior distribution of $f_{X_{m+1}}$ given $\{C_{n,h_n(X_{m+1})}\}_{n \in [N]}$; that is, it allows us to treat the h_n 's as they were fixed, i.e. non-random hash functions. For a single hash function h_n , in the next theorem we provide a rigorous derivation of the posterior distribution of $f_{X_{m+1}}$, given $C_{n,h_n(X_{m+1})}$.

Theorem 3 For $m \geq 1$, let $x_{1:m}$ be a stream of tokens that are modeled as a random sample $X_{1:m}$ from $P \sim \text{PYP}(\alpha, \theta; \nu)$, and let X_{m+1} be an additional random sample from P . Moreover, let h_n be a random hash function distributed as the strong universal family \mathcal{H} , and let h_n be independent of $X_{1:m}$ for any $m \geq 1$, that is h_n is independent of P . Then, for $l = 0, 1, \dots, c_n$

$$pf_{X_{m+1}}(l; m, c_n, \alpha, \theta) \tag{24}$$

$$\begin{aligned}
 &:= \Pr[f_{X_{m+1}} = l \mid C_{n, h_n}(X_{m+1}) = c_n] \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n - l, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n + 1, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)}.
 \end{aligned}$$

See Appendix B for the proof of Theorem 3. Theorem 3 provides an extension of Theorem 1 to the more general BNP model (23); in particular, Theorem 1 can be recovered from Theorem 3 by setting $\alpha = 0$. See Appendix C for details. The proof of Theorem 3 relies on two properties of the distribution (18) for the random partition \mathbf{M}_n induced by $P \sim \text{PYP}(\alpha, \theta; \nu)$: P1) the distribution of \mathbf{M}_n is of Gibbs-type (Pitman, 2003; Gnedin and Pitman, 2006), i.e.

$$\Pr[\mathbf{M}_m = \mathbf{m}] = V_{m,k} m! \prod_{i=1}^m \left(\frac{(1-\alpha)_{(i-1)}}{i!} \right)^{m_i} \frac{1}{m_i!} \mathbb{1}_{\mathcal{M}_{m,k}}(\mathbf{m}), \quad (25)$$

where $(V_{m,k})_{m \geq 1, k \leq m}$ forms a collection of nonnegative (deterministic) weights that satisfy the triangular recursion $V_{m,k} = V_{m+1,k}(m - \alpha k) + V_{m+1,k+1}$ with the proviso $V_{1,1} := 1$; P2) for any constant $c > -\alpha$, the $V_{m,k}$'s are such that for any $m \geq 1$ and any $k \leq m$ it holds that

$$\frac{V_{m+1,k+1}}{V_{m,k}} = \frac{c + k\alpha}{c + m}. \quad (26)$$

In particular, within the proof of Theorem 3, P2) is critical to exploit the quasi-conjugacy property of the PYP prior (Lijoi et al., 2008). While P1) holds true for a broad class of nonparametric priors, which are referred to as Gibbs-type priors (Pitman, 2003; Gnedin and Pitman, 2006; Pitman, 2006; De Blasi et al., 2015; Bacallado et al., 2017), P2) holds true only for the PYP prior. That is, the PYP prior is characterized as the sole Gibbs-type prior for which P2) holds true (Zabell, 1997; Bacallado et al., 2017). Therefore, we expect that an extension of Theorem 3 to Gibbs-type priors may require a non trivial effort to overcome the lack of P2), though the general strategy underlying the ‘‘Bayesian’’ proof remains valid.

For $\alpha \in [0, 1)$, an alternative expression for (24) may be given in terms of the distribution (19) of the number K_m of distinct types in a random sample from the PYP. If $c_n > 0$, then for $l = 0, 1, \dots, c_n$

$$\begin{aligned}
 &p_{f_{X_{m+1}}}(l; m, c_n, \alpha, \theta) \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{(\theta)_{(c_n-l)} \mathbb{E} \left[\frac{\left(\frac{\theta+\alpha}{\alpha}\right)_{(K_{c_n-l} + K_{m-c_n})}}{\left(\frac{\theta}{\alpha}\right)_{(K_{c_n-l})} \left(\frac{\theta}{\alpha}\right)_{(K_{m-c_n})}} \left(\frac{1}{J}\right)^{K_{c_n-l}} \left(1 - \frac{1}{J}\right)^{K_{m-c_n}} \right]}{(\theta)_{(c_n+1)} \mathbb{E} \left[\frac{\left(\frac{\theta}{\alpha}\right)_{(K_{c_n+1} + K_{m-c_n})}}{\left(\frac{\theta}{\alpha}\right)_{(K_{c_n+1})} \left(\frac{\theta}{\alpha}\right)_{(K_{m-c_n})}} \left(\frac{1}{J}\right)^{K_{c_n+1}} \left(1 - \frac{1}{J}\right)^{K_{m-c_n}} \right]}
 \end{aligned} \quad (27)$$

with the proviso that $K_0 = 0$, where K_{c_n-l} and K_{m-c_n} in the numerator of (27) are independent random variables for any $l = 0, 1, \dots, c_n - 1$, and K_{c_n+1} and K_{m-c_n} in the denominator of (27) are independent random variables. See Appendix D for the proof of Equation (27). Equation (27) gives a probabilistic representation of the posterior distribution (24), whose critical terms are the expected value of a suitable functional of (K_{c_n-l}, K_{m-c_n}) , i.e. the

numerator of (27), and the expected value of a suitable functional of (K_{c_n+1}, K_{m-c_n}) , i.e. the denominator of (27). See Appendix E for a further alternative expression of (24), which is in terms of exponentially tilted α -stable random variables (Zolotarev, 1986). Figure 2 shows the shape behaviour of the posterior distribution (27) for different values of (α, θ) , keeping m, J and c_n fixed. For $\alpha = 0$, i.e. under the DP prior, Cai et al. (2018) showed that the posterior distribution of $f_{X_{m+1}}$, given $C_{n, h_n(X_{m+1})}$ is monotonically decreasing or increasing. Under the PYP, the additional parameter $\alpha \in (0, 1)$ allows for a more flexible shape behaviour.

Remark 4 Equation (27) is useful for the numerical evaluation of the posterior distribution (24), since it avoids numerical issues that arise in evaluating the generalized factorial coefficients. In particular, (27) allows for a Monte Carlo (MC) evaluation of (24), which requires to sample from the random variable K_m , for suitable choices of m . Sampling K_m is straightforward, and it exploits the predictive probabilities of the PYP. In particular, from (20), $\text{Bernoulli}(p)$ is the Bernoulli distribution with parameter p , for $p \in (0, 1)$, then sampling K_m reduces to sample $(m-1)$ Bernoulli random variables. See Algorithm 1 in Section 4.

Under the PYP prior, Theorem 3 shows that the posterior distribution of $f_{X_{m+1}}$, given $C_{n, h_n(X_{m+1})}$, depends on the sampling information through c_n and m . This is a critical difference with respect to the DP prior, where Theorem 3 shows that the posterior distribution of $f_{X_{m+1}}$, given $C_{n, h_n(X_{m+1})}$, depends on the sampling information only through m . Therefore, under the PYP prior, one may consider different large m asymptotic behaviours for the posterior distribution (24). Here, we start by considering a local limit theorem of (24) for $m \rightarrow +\infty$, while c_n is fixed. Under the setting of Theorem 3, for any $l = 0, 1, \dots, c_n$ it holds

$$p(l; c_n, \alpha, \theta) := \lim_{m \rightarrow +\infty} p_{f_{X_{m+1}}}(l; m, c_n, \alpha, \theta) = \binom{c_n}{l} (1-\alpha)^l \frac{(\theta + 2\alpha)_{(c_n-l)}}{(\theta + \alpha + 1)_{(c_n)}} \quad (28)$$

and

$$p(l; c_n, \alpha, \theta) = \int_0^1 \text{Binomial}(l; c_n, p) f_{B_{1-\alpha, \theta+2\alpha}}(p) dp, \quad (29)$$

where $f_{B_{a,b}}$ is the density function of the distribution of a Beta random variable $B_{a,b}$. See Appendix F for the proof of Equation (28) and Equation (29). The next proposition is an interesting complement to Theorem 3, providing the large m asymptotic behaviour of the posterior distribution (24). In particular, we consider $m \rightarrow +\infty$ and $c_n \rightarrow +\infty$ with the assumption that $c_n = \lambda m$ for some choice of $\lambda \in (0, 1)$. Such a large m asymptotic behaviour is in line with the large c_n asymptotic behaviour presented in Proposition 2 under the DP prior.

Proposition 5 For $\alpha \in (0, 1)$ and $c > 0$ let $S_{\alpha,c}$ be a polynomially tilted α -stable random variable, i.e. the distribution of $S_{\alpha,c}$ has density function $f_{S_{\alpha,c}}(x) \propto x^{-c} g_\alpha(x) \mathbb{1}_{\mathbb{R}^+}(x)$ for g_α being the positive α -stable density function; moreover, set $Z_{\alpha, \theta+\alpha} = (J-1)^{1/\alpha} S_{\alpha,0}/S_{\alpha, \theta+\alpha}$ and $W_{\alpha, \theta} = (J-1)^{1/\alpha} S_{\alpha,0}/S_{\alpha, \theta}$, with $S_{\alpha,0}$ being independent of $S_{\alpha, \theta+\alpha}$ and of $S_{\alpha, \theta}$, and

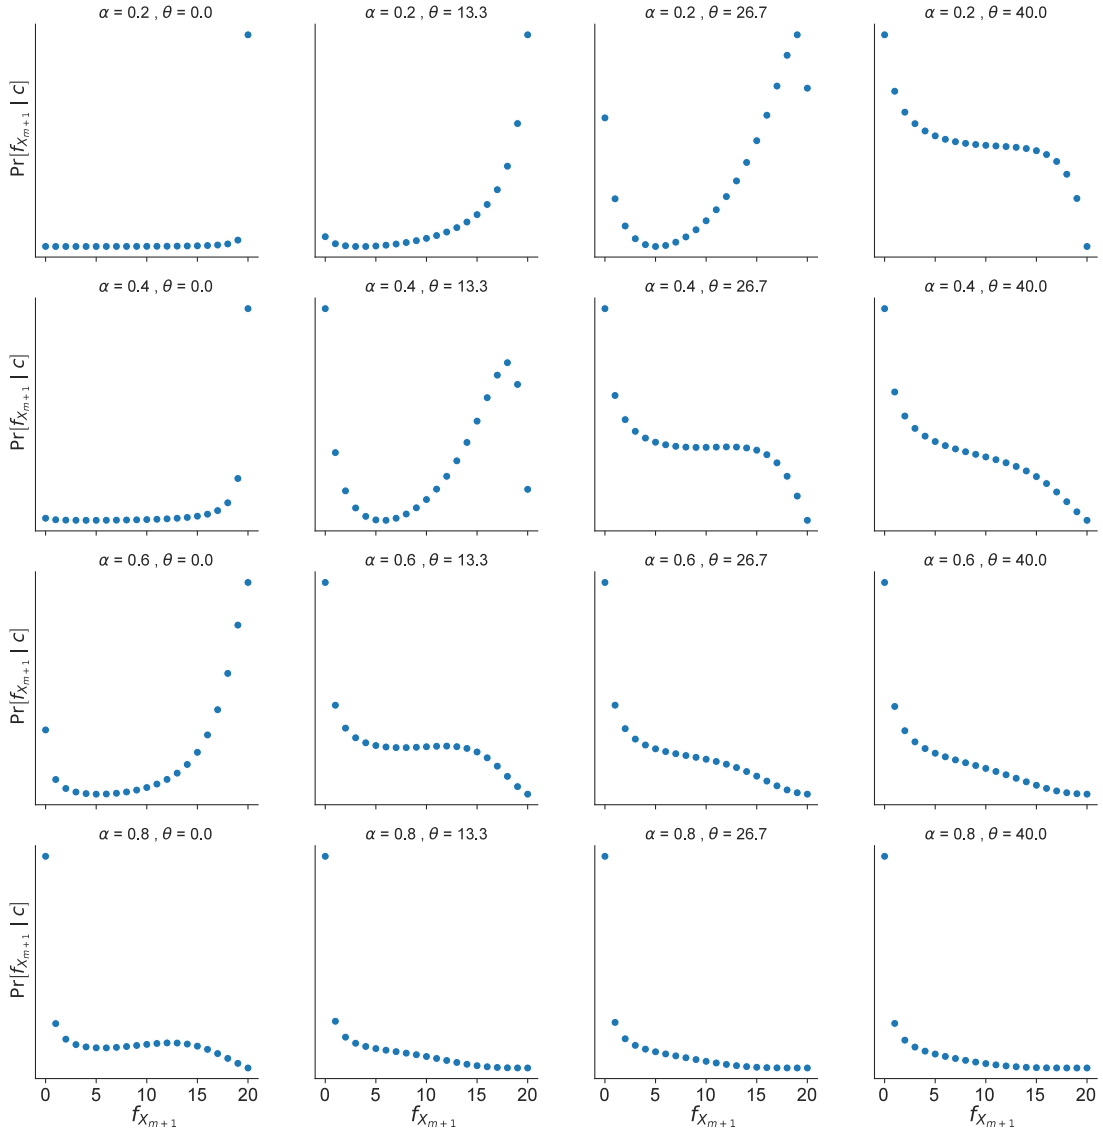


Figure 2: Posterior distribution of $f_{X_{m+1}}$ given $C_{n, h_n}(X_{m+1}) = c_n$ under $P \sim \text{PYP}(\alpha, \theta; \nu)$: $m = 1000$, $J = 50$ and $c_n = 20$.

denote by $f_{Z_{\alpha, \theta + \alpha}}$ and $f_{W_{\alpha, \theta}}$ the density functions of the distributions of $Z_{\alpha, \theta + \alpha}$ and $W_{\alpha, \theta}$, respectively. Under the setting of Theorem 3, let $F_{X_{m+1}}$ be a random variable with distribution (24). As $m \rightarrow +\infty$ and under the large m asymptotic regime $c_n = \lambda m$, for some choice of $\lambda \in (0, 1)$,

$$\frac{F_{X_{m+1}}}{c_n} \xrightarrow{w} B_{1-\alpha, \theta + \alpha}^{(\lambda)}, \quad (30)$$

where $B_{1-\alpha, \theta+\alpha}^{(\lambda)}$ is a random variable whose distribution has density function of the following form

$$f_{B_{1-\alpha, \theta+\alpha}^{(\lambda)}}(x) = \frac{\frac{\Gamma(\theta+1)}{\Gamma(\theta+\alpha)\Gamma(1-\alpha)}}{f_{W_{\alpha, \theta}}(\lambda^{-1}-1)} \frac{f_{Z_{\alpha, \theta+\alpha}}\left(\frac{\lambda^{-1}-1}{1-x}\right)}{1-x} x^{1-\alpha-1} (1-x)^{\theta+\alpha-1} \mathbb{1}_{(0,1)}(x).$$

See Appendix G for the proof of Proposition 5. As in the context of the DP prior discussed in Section 2, Proposition 5 shows that the posterior distribution of $f_{X_{m+1}}$ given $C_{h_n(X_{m+1})}$ admits a representation in terms of a mixture of Binomial distributions. In particular, Proposition 5 may be viewed as the natural counterpart of Proposition 2, though the resulting mixing distribution is not as simple as the Beta distribution of Proposition 2. For the collection of hash functions h_1, \dots, h_N , the posterior distribution of $f_{X_{m+1}}$, given $\{C_{n, h_n(X_{m+1})}\}_{n \in [N]}$, follows from Theorem 3 by means of the assumption that the h_n 's are i.i.d. according to the strong universal family \mathcal{H} . In particular, by a direct application of Bayes theorem, straightforward calculations show that for $l = 0, 1, \dots, \min\{c_1, \dots, c_N\}$ it holds that

$$\Pr[f_{X_{m+1}} = l \mid \{C_{n, h_n(X_{m+1})}\}_{n \in [N]} = \{c_n\}_{n \in [N]}] = \frac{\prod_{n \in [N]} p_{f_{X_{m+1}}}(l; m, c_n, \alpha, \theta)}{(p_{f_{X_{m+1}}}(l; m, \alpha, \theta))^{N-1}}, \quad (31)$$

where

$$p_{f_{X_{m+1}}}(l; m, \alpha, \theta) := \Pr[f_{X_{m+1}} = l] = \binom{m}{l} (1-\alpha)^{(l)} \frac{(\theta+\alpha)_{(m-l)}}{(\theta+1)_{(m)}}$$

for $l = 0, 1, \dots, m$, where $f_{B_{1-\alpha, \theta+\alpha}}$ denotes the density function of the distribution of a Beta random variable with parameter $(1-\alpha, \theta+\alpha)$. See Appendix H for the proof of Equation (31). CMS-PYP estimates of the point query $f_{X_{m+1}}$, with respect to a suitable choice of a loss function, are obtained as functionals of the posterior distribution (31), e.g. posterior mode, mean and median. In general, the numerical evaluation of the posterior distribution (31), as well as the evaluation of its alternative expression in terms of (27), requires care in order to achieve numerical stability and efficiency, as it is discussed in the last part of this section.

To apply (31), it remains to estimate the prior's parameter (α, θ) . For ease of exposition, we denote by \mathbf{C} the $N \times J$ matrix with entries $C_{n,j}$ for $n \in [N]$ and $j \in [J]$. Assuming that the matrix \mathbf{C} has been computed from m tokens, the sum of the entries of each row of \mathbf{C} is equal to the sample size m . Since the PYP does not have a restriction property analogous to that of the DP, under the BNP model (23) the distribution of \mathbf{C} is not available in closed-form. Hence, the prior's parameter (α, θ) cannot be estimated following the empirical Bayes approach adopted by Cai et al. (2018) in the context of the DP prior. Instead, here we estimate (α, θ) by relying on the minimum Wasserstein distance method (Bernton et al., 2019). This method estimates (α, θ) by selecting the value of (α, θ) that minimizes the expected Wasserstein distance between a summary statistic of the data and the corresponding summary statistic of synthetic data generated under the BNP model (23). In our context, a natural choice for the summary statistic is the matrix \mathbf{C} . By construction, the rows of \mathbf{C} are i.i.d.; moreover, since \mathcal{H} is assumed to be a perfectly random hash family, each column of \mathbf{C} is exchangeable. Then, we can define the reference summary statistic $\overline{\mathbf{C}}$

as a vector of length NJ containing the (unordered) entries of the matrix \mathbf{C} . For any fixed $m' \geq 1$ and a any fixed prior's parameter (α, θ) , let $\tilde{X}_{1:m'} = (\tilde{X}_1, \dots, \tilde{X}_{m'})$ be a random sample from $P \sim \text{PYP}(\alpha, \theta; \nu)$, i.e. $\tilde{X}_{1:m'}$ is modeled as (23). For a moderate sample size m' , generating random variates from $\tilde{X}_{1:m'}$ is straightforward by means of the predictive distribution (20) of the PYP. These random variates, by a direct transformation through the hash functions h_1, \dots, h_N drawn at random from \mathcal{H} , lead to random variates from the hashed frequencies and to random variates from reference summary statistic, denoted by $\tilde{\mathbf{C}}(\alpha, \theta, m')$.

In practice, m is such that $m \gg m'$ and the computational cost of sampling from (20) scales super-linearly in m' . To account for this mis-match we scale the entries of $\tilde{\mathbf{C}}(\alpha, \theta, m')$ by m/m' , so that each row of $\tilde{\mathbf{C}}(\alpha, \theta, m')m/s$ sum to m . Now, we are interested in finding $(\hat{\alpha}, \hat{\theta})$ such that

$$(\hat{\alpha}, \hat{\theta}) = \arg \min_{(\alpha, \theta)} \mathbb{E} \left[\mathcal{W}_1 \left(\bar{\mathbf{C}}, \tilde{\mathbf{C}}(\alpha, \theta, m') \frac{m}{m'} \right) \right], \quad (32)$$

where \mathcal{W}_1 is the Wasserstein distance of order 1, and the expectation is taken with respect to $\tilde{\mathbf{C}}$. To fully specify the optimization problem we choose $\rho(x, y) = |x - y|$ as distance underlying \mathcal{W}_1 (Bernton et al., 2019). We use an MC approximation of the expectation in (32), i.e.,

$$\frac{1}{R} \sum_{r=1}^R \mathcal{W}_p \left(\bar{\mathbf{C}}, \tilde{\mathbf{C}}_r(\alpha, \theta, m') \frac{m}{m'} \right) \quad (33)$$

for $R \geq 1$, where $(\tilde{\mathbf{C}}_1(\alpha, \theta, m'), \dots, \tilde{\mathbf{C}}_R(\alpha, \theta, m'))$ are i.i.d. according to $\tilde{\mathbf{C}}(\alpha, \theta, m')$. We refer to Bernton et al. (2019) for a theoretical and empirical analysis of the minimum distance Wasserstein method. To improve the MC approximation displayed in (33), which might be detrimental for the minimization problem in (32), we fix the same random numbers underlying the routines used for generating random variates from the predictive distribution (20) of the PYP over all values of (α, θ) . Moreover the optimization is carried out via noise-robust Gaussian optimization (Letham et al., 2019). We report experimental results in Section 4.

3.2 Computational aspects of the CMS-PYP

For the CMS-PYP estimator of $f_{x_{m+1}}$ we consider the posterior mean $\hat{f}^{(\text{PYP})}$, that is the expected value of the posterior distribution (31). The evaluation of $\hat{f}^{(\text{PYP})}$ follows from two steps:

- i) the estimation of the prior's parameter (α, θ) by means of the minimum Wasserstein distance method;
- ii) the evaluation, with respect to the estimated prior's parameter, of the posterior distribution (31).

Step i) has been described above. Step ii) can be implemented either via the exact representation in (27) or via its limiting behaviour in (28), which is accurate provided that the total number of observed tokens m is large relative to the considered c_n . This is often the case, especially for real world large datasets where applying CMS in any of its variants is most

warranted. In our numerical experiments we consider datasets whose total observed tokens range from 2 millions to almost 1 billion. The evaluation (27) requires the computation of multiple expectations, one for each $l = 0, \dots, c_n$, which we approximate via MC integration. For each MC estimator to be valid it is necessary to sample each K_{c_n-l} independently from K_{m-c_n} in each expectation term. However the MC estimators themselves, one for each $l = 0, \dots, c_n$, can be correlated. One sample for all MC estimators can be thus obtained as follows: i) Algorithm 1 is used to sample the vector $[K_{c_n-l} \mid l = 0, \dots, c_n]$ in one pass with $\mathcal{O}(c_n)$ cost ii) K_{m-c_n} is sampled from the distribution of $[(\theta + (m - c_n))^\alpha - \theta^\alpha] S_{\alpha, \theta}^{-\alpha}$ where $S_{\alpha, \theta}^{-\alpha}$ is a polynomially tilted α -stable random variable. Sampling from $S_{\alpha, \theta}^{-\alpha}$ can be achieved efficiently by using rejection sampling as described in Devroye (2009). The convergence of K_{m-c_n} to its limiting distribution is fast in $m - c_n$, as illustrated in Figure 3. To ensure numerical stability with both (27) and (28) we work in log-space, i.e. compute the (natural) logarithm of each multiplicative term of (31), and exponentiate back only as final computation. Similarly, to avoid underflow/overflow issues, we apply the “log-sum-exp” trick to sums arising from the MC estimators. The denominator of (27) does not need to be evaluated, as it suffices to compute $p_{f_{X_{m+1}}}(l; c_n, \alpha, \theta)$ up to a constant of proportionality and then normalize the masses to sum up to 1. In doing this, the MC variance is additionally reduced.




Figure 3: Exact sampling of K_m via Algorithm 1 and (asymptotic) approximate sampling of K_m via $K_m \sim [(\theta + m)^\alpha - \theta^\alpha] S_{\alpha, \theta}^{-\alpha}$ for different values of m , $\alpha = 0.5$, $\theta = 10$; densities estimated by kernel density estimation.

4. Experiments

We present numerical experiments for the CMS-PYP. We apply the CMS-PYP to synthetic and real textual data, and we compare its performance with respect to some Bayesian and non-Bayesian approaches. With regards to the Bayesian approaches, we consider the CMS-DP of Cai et al. (2018), here denoted by $\hat{f}^{(\text{DP})}$, and the CMS-NIGP of Dolera et al. (2021), here denoted by $\hat{f}^{(\text{NIGP})}$. In particular, $\hat{f}^{(\text{DP})}$ is the mean of the posterior distribution of $f_{X_{m+1}}$, given $\{C_{n, h_n}(X_{m+1})\}_{n \in [N]}$, under the DP prior, whereas $\hat{f}^{(\text{NIGP})}$ is the mean of the posterior distribution of $f_{X_{m+1}}$, given $\{C_{n, h_n}(X_{m+1})\}_{n \in [N]}$, under the NIGP prior. With regards to non-Bayesian approaches, we consider the CMS of Cormode and Muthukrishnan

Algorithm 1 Sampling K_{c-l} for $l = 0, \dots, c$

```

 $K[0] \leftarrow 0;$ 
 $K[1] \leftarrow 1;$ 
 $i \leftarrow 1;$ 
while  $i \leq c$  do
     $Ber \leftarrow$  random sample from Bernoulli  $\left(\frac{\theta + \alpha K[i-1]}{\theta + i}\right);$ 
     $K[i] \leftarrow K[i-1] + Ber;$ 
     $i \leftarrow i + 1;$ 
end while
return  $K[c], K[c-1], \dots, K[1]$ 

```

(2005), the CMM of Goyal et al. (2012) and the BDCM of Ting (2018). The CMS, here denoted by $\hat{f}^{(\text{CMS})}$, has been specified in (1). Both the CMM, here denoted by $\hat{f}^{(\text{CMM})}$, and the BDCM, here denoted by $\hat{f}^{(\text{BDCM})}$, rely on the same summary statistics used in the CMS, i.e. the hashed frequencies $\{\mathbf{C}_n\}_{n \in [N]}$. This facilitates the implementation of a fair comparison among estimators, since the storage requirement and sketch update complexity are unchanged. The CMM estimator subtracts the value of the estimated noise from each of the N counters, and returns minimum between the median of the N residues and the CMS estimator. For a token x the noise corresponding to each counter is given by $(m - C_{n, h_n(x)}) / (J - 1)$. Similarly to the CMM, the BDCM estimator aims at de-biasing the CMS estimator. Specifically, we consider the Algorithm 1 of Ting (2018). An (almost) unbiased estimator is obtained by removing the mean of the CMS estimators computed over the vector counters $\{\mathbf{C}_{n,j}\}_{n \in [N]}$ for $j = 1, \dots, J$. The BDCM estimator is obtained by restricting this unbiased estimator to be non-negative. In all the experiments random hash functions $h_n(x)$ as

$$h_n(x) = ((a_n x + b_n) \bmod LP) \bmod J,$$

where x is the non-negative integer index corresponding to the token of interest, LP denotes a large prime number (here chosen to be $2^{32} - 1$), and a_n and b_n are two random integer numbers that are i.i.d. as a Uniform distribution on $[1, LP]$. The required hash functions are generated once for each numerical experiment and kept fixed while comparing different estimators.

4.1 Estimation of the prior's parameter (α, θ)

We present an empirical study of the likelihood-free estimation approach detailed in Section 3. We start with a scenario where the data generating process (PYP-DGP) is (23). In particular, we generate 10 synthetic datasets of $m = 300000$ tokens each, for different prior's parameter (α, θ) . See Table 1. For each dataset, the estimation of the prior's parameter (α, θ) is performed by means of (32) and (33) with $m' = 100000$ for $R = 25$. The optimization procedure is based on Letham et al. (2019), as implemented by the AX library. See <https://ax.dev/> for details. The stochastic objective function (33) is evaluated a total of 50 times for each dataset. Results from Table 1 support our inferential procedure for (α, θ) . It is also apparent that, for the datasets under consideration, α is more easily identified than θ . We also consider synthetic datasets generated from Zipf's distributions

with (exponent) parameter $c > 1$, i.e. a Zipf’s data generating process with parameter c (\mathcal{Z}_c -DGP). In particular, we recall that the parameter c controls the tail behaviour of the Zipf’s distribution: the smaller c the heavier is the tail of the distribution, i.e., the smaller c the larger the fraction of types with low-frequency tokens. We generate 7 synthetic datasets of $m = 500000$ tokens each, for different parameter c . See Table 2. For each dataset, the estimation of the prior’s parameter (α, θ) is performed by means of (32) and (33) with $m' = 100000$ for $R = 25$. The optimization procedure is still based on the work of Letham et al. (2019). The stochastic objective function (33) is evaluated a total of 50 times for each dataset. The results from Table 2 shows that the PYP prior is able to adapt to different power-law tails behaviours. In particular, we observe that the larger c the smaller $\hat{\alpha}$, which is in agreement with the interpretation of α as the parameter controlling the tail behaviour of the PYP prior.

PYP-DGP		Estimates	
α	θ	$\hat{\alpha}$	$\hat{\theta}$
0.00	25.00	0.02	36.31
0.10	25.00	0.11	21.86
0.20	25.00	0.18	16.78
0.30	25.00	0.26	22.83
0.40	25.00	0.41	17.32
0.50	25.00	0.56	10.69
0.60	25.00	0.56	13.42
0.70	25.00	0.63	24.89
0.80	25.00	0.77	10.21
0.90	25.00	0.88	11.26

Table 1: Prior’s parameter (α, θ) estimates, under PYP-DGP.

\mathcal{Z}_c -DGP	Estimates	
c	$\hat{\alpha}$	$\hat{\theta}$
1.05	0.92	25.37
1.18	0.80	5.56
1.33	0.71	1.53
1.54	0.67	0.61
1.82	0.38	0.49
2.22	0.17	0.11
2.86	0.01	0.23

Table 2: Prior’s parameter (α, θ) estimates, under \mathcal{Z}_c -DGP.

4.2 Applications to synthetic and real data

With regards to synthetic data, we consider datasets generated from Zipf’s distributions with exponent $c = 1.3, 1.6, 1.9, 2.2$. Each dataset consists of $m = 500000$ tokens. We make use of a 2-universal hash family, and then assume the following pairs of hashing parameters: i) $J = 320$ and $N = 2$; ii) $J = 160$ and $N = 4$. Bayesian and non-Bayesian approaches are compared in terms of the MAE (mean absolute error) between true frequencies and their estimates. Table 3 and Table 5 report the MAE for the Bayesian approaches, with respect to the case $J = 320$ and $N = 2$ and the case $J = 160$ and $N = 4$, respectively. From Table 3 and Table 5, it is clear that $\hat{f}^{(\text{PYP})}$ has a remarkable better performance than $\hat{f}^{(\text{DP})}$ in the estimation of low-frequency tokens. In particular, for both Table 3 and Table 5, if we consider the bin of low-frequencies $(0, 256]$ the MAE of $\hat{f}^{(\text{PYP})}$ is always smaller than the MAE of $\hat{f}^{(\text{DP})}$, i.e. $\hat{f}^{(\text{PYP})}$ outperforms $\hat{f}^{(\text{DP})}$. This behaviour becomes more and more evident as the parameter c decreases, that is the heavier is the tail of the distribution the more the estimator $\hat{f}^{(\text{PYP})}$ outperforms the estimator $\hat{f}^{(\text{DP})}$. For any fixed exponent c , the gap between the MAEs of $\hat{f}^{(\text{PYP})}$ and $\hat{f}^{(\text{DP})}$ reduces as v increases, and this reduction is much more evident as c becomes large. For any exponent c we expect a frequency threshold, say $v^*(c)$, such that $\hat{f}^{(\text{PYP})}$ underestimates $f_{x_{m+1}}$ for $v > v^*(c)$. From Table 3 and Table 5, for any two exponents c_1 and c_2 such that $c_1 < c_2$ it will be $v^*(c_1) > v^*(c_2)$. From Table 3, i.e. $J = 320$ and $N = 2$, it emerges that $\hat{f}^{(\text{PYP})}$ has a remarkable better performance than $\hat{f}^{(\text{NIGP})}$ for data with heavier power-law tails ($c = 1.3, 1.6$), whereas from Table 5, i.e. $J = 160$ and $N = 4$, $\hat{f}^{(\text{PYP})}$ is competitive with $\hat{f}^{(\text{NIGP})}$ for data with heavier power-law tails ($c = 1.3, 1.6$). In general, the better performance of $\hat{f}^{(\text{PYP})}$ with respect to $\hat{f}^{(\text{NIGP})}$ is not surprising, as $\hat{f}^{(\text{NIGP})}$ has the critical limitation that it can not be tuned to the power-law degree of the data.

Table 4 and Table 6 report the MAE for the non-Bayesian approaches for the case $J = 320$ and $N = 2$ and the case $J = 160$ and $N = 4$, respectively. From Table 4 and Table 6, it is clear that $\hat{f}^{(\text{PYP})}$ outperforms the $\hat{f}^{(\text{CMS})}$ in the estimation of low-frequency tokens for both the choices of hashing parameters, i.e. the case $J = 320$ and $N = 2$ and the case $J = 160$ and $N = 4$. Moreover, $\hat{f}^{(\text{PYP})}$ outperforms both $\hat{f}^{(\text{CMM})}$ and $\hat{f}^{(\text{BDCM})}$ in the estimation of low-frequency token. The better performance of $\hat{f}^{(\text{PYP})}$ with respect to $\hat{f}^{(\text{CMM})}$ is a remarkable results, as the CMM is known to stand out in the estimation of low-frequency tokens Goyal et al. (2012, Figure 1). In general, a good performance in the estimation of low-frequency tokens is a desirable feature in natural language or textual data, where it is common the power-law behaviour of the data stream of tokens. In such a data, highest frequency events are often of low interest: frequent words are often grammatical, highly polysemous or without any interesting semantics, while low-frequency words are more relevant.

We conclude by presenting an application of the CMS-PYP to some textual datasets, for which the distribution of words is typically a power-law distribution (Clauset et al., 2009). In particular, we consider 4 textual datasets of increasing corpora size: the 20 Newsgroups dataset¹, the Enron dataset², the WikiText-103 dataset³ and the 1 Billion Word

1. <http://qwone.com/~jason/20Newsgroups/>

2. <https://archive.ics.uci.edu/ml/machine-learning-databases/bag-of-words/>

3. <https://blog.salesforceairesearch.com/the-wikitext-long-term-dependency-language-modeling-dataset/>

Language Model Benchmark (1BWLMB) dataset⁴. The 20 Newsgroups dataset consists of $m = 2765300$ tokens with 53975 distinct tokens, whereas the Enron dataset consists of $m = 6412175$ tokens with 28102 distinct tokens. Following the experiments in Cai et al. (2018), we make use of a 2-universal hash family, with the following hashing parameters: i) $J = 12000$ and $N = 2$; ii) $J = 8000$ and $N = 4$. By means the goodness of fit test proposed in Clauset et al. (2009), we found that the 20 Newsgroups and Enron datasets fit with a power-law distribution with exponent $\nu = 2.3$ and $\nu = 2.1$, respectively. The CMS-PYP estimators $\hat{f}^{(PYP)}$ for the 20 Newsgroups and Enron datasets are obtained through the implementation of (27). Table 8 and Table 9 report the MAEs of $\hat{f}^{(CMS)}$, $\hat{f}^{(CMM)}$ and $\hat{f}^{(BDCM)}$, whereas Table 7 reports the MAEs of the estimators $\hat{f}^{(DP)}$ and $\hat{f}^{(PYP)}$. Results of Table 8 and Table 9 and Table 7 confirm the behaviour observed in Zipf’ synthetic data. That is, $\hat{f}^{(PYP)}$ outperforms $\hat{f}^{(DP)}$ for low-frequency tokens. Moreover, a comparison with respect to $\hat{f}^{(CMM)}$ reveals that $\hat{f}^{(PYP)}$ is competitive with $\hat{f}^{(CMM)}$ in the context of the estimation of low-frequency tokens.

Finally, we consider the WikiText-103 and 1BWLMB datasets. The former consists of $m = 82810656$ tokens with 606753 distinct tokens, whereas the latter consists of $m = 658195953$ tokens with 1256524 distinct tokens. The fit test of Clauset et al. (2009) results in power-law distributions with exponent $\nu = 2.15$ and $\nu = 1.5$ respectively. Taking into account the increased corpora sizes we consider the following hashing parameters: i) $J = 50000$ and $N = 2$; ii) $J = 35000$ and $N = 4$ for WikiText-103; i) $J = 140000$ and $N = 2$; ii) $J = 100000$ and $N = 4$ for 1BWLMB. The CMS-PYP estimators $\hat{f}^{(PYP)}$ are obtained through the implementation of (28). Table 10 reports the MAEs of the estimators $\hat{f}^{(DP)}$ and $\hat{f}^{(PYP)}$ applied to the WikiText-103 dataset and to the 1BWLMB dataset. The CMS-PYP estimator offers a competitive performance with respect to both the DP and the CMM estimators. The use of (28) reduces the computational significantly, in which case the time required to compute the CMS-PYP estimators is similar to the time required for DP estimators.

Bins of x_{m+1}	$\mathcal{Z}_{1.3}$			$\mathcal{Z}_{1.6}$			$\mathcal{Z}_{1.9}$			$\mathcal{Z}_{2.2}$		
	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$
(0,1]	1,057.61	231.31	1.12	626.85	134.75	3.36	306.70	65.71	115.15	51.38	12.91	3.80
(1,2]	1,194.67	287.43	2.08	512.43	119.22	2.29	153.57	37.03	31.16	288.27	61.87	93.99
(2,4]	1,105.16	262.18	3.63	472.59	95.78	1.85	2,406.00	353.73	1,237.41	133.31	26.90	17.57
(4,8]	1,272.02	302.89	7.40	783.88	175.10	8.89	457.57	83.30	136.16	117.76	21.58	8.26
(8,16]	1,231.63	257.08	11.83	716.52	136.66	10.00	377.99	66.44	90.41	411.21	77.39	127.69
(16,32]	1,252.18	248.41	22.58	829.17	190.05	14.81	286.98	41.99	65.47	501.00	90.29	178.07
(32,64]	1,309.14	284.12	39.23	780.70	139.52	36.47	413.95	67.30	181.84	216.84	48.00	92.07
(64,128]	1,716.76	312.59	104.03	946.20	125.07	79.94	1,869.23	353.10	1,678.82	63.05	65.91	85.70
(128,256]	1,102.96	97.9	168.34	1,720.49	273.50	342.18	199.87	110.32	98.20	45.98	130.94	136.25

Table 3: Synthetic data: MAE for $\hat{f}^{(DP)}$, $\hat{f}^{(NIGP)}$ and $\hat{f}^{(PYP)}$, case $J = 320, N = 2$.

5. Discussion

At the core of the CMS-DP of Cai et al. (2018) lies the computation of the posterior distribution of a point query, given the hashed data, and then estimates of the point query are obtained as mean functionals of such a posterior distribution. While the CMS-DP has

⁴. <https://www.statmt.org/lm-benchmark/>

Bins of x_{m+1}	$\mathcal{Z}_{1,3}$			$\mathcal{Z}_{1,6}$			$\mathcal{Z}_{1,9}$			$\mathcal{Z}_{2,2}$		
	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$
(0,1]	1,061.3	161.72	99.02	629.40	62.19	57.08	308.11	81.10	77.37	51.65	1.04	27.80
(1,2]	1,197.9	169.74	35.67	514.31	102.42	45.38	154.20	2.00	55.14	289.50	2.04	21.16
(2,4]	1,108.3	116.37	62.74	474.82	52.10	26.18	2,419.51	2,215.85	1156.54	134.05	3.40	24.63
(4,8]	1,275.9	378.04	145.04	786.73	214.46	40.30	460.13	258.90	68.67	118.40	6.44	62.20
(8,16]	1,236.1	230.32	90.54	719.84	232.24	66.08	380.05	139.50	39.40	413.13	129.03	51.54
(16,32]	1,256.8	221.98	172.14	831.70	79.73	62.98	288.59	23.90	263.34	503.60	364.30	32.60
(32,64]	1,312.8	235.87	197.59	783.90	184.99	73.98	415.58	54.82	116.04	217.81	82.92	28.73
(64,128]	1,721.7	766.29	119.64	950.31	304.36	56.60	1,875.50	1,762.20	1,120.00	64.01	64.01	25.79
(128,256]	1,107.7	334.57	121.17	1,727.19	1,488.38	109.26	202.09	163.61	239.54	46.80	46.80	25.39

Table 4: Synthetic data: MAE for $\hat{f}^{(CMS)}$, $\hat{f}^{(CMM)}$ and $\hat{f}^{(BDCM)}$, case $J = 320, N = 2$.

Bins of x_{m+1}	$\mathcal{Z}_{1,3}$			$\mathcal{Z}_{1,6}$			$\mathcal{Z}_{1,9}$			$\mathcal{Z}_{2,2}$		
	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(NIGP)}$	$\hat{f}^{(PYP)}$
(0,1]	2,206.09	0.9	0.77	1,254.85	0.25	1.07	420.76	0.18	0.98	153.20	0.32	28.78
(1,2]	2,333.06	0.5	1.07	1,326.71	0.70	2.13	549.12	0.82	1.93	180.71	1.24	21.60
(2,4]	2,266.35	1.3	1.70	1,267.97	2.47	3.53	482.45	2.53	3.55	182.18	2.66	14.92
(4,8]	2,229.22	4.6	4.54	1,371.27	4.67	6.11	538.91	5.28	6.28	250.32	5.96	40.18
(8,16]	2,207.42	10.5	7.06	1,159.29	10.68	11.68	487.69	10.86	10.64	245.09	10.28	95.33
(16,32]	2,279.80	20.7	11.60	1,211.41	19.21	23.88	529.77	22.08	19.04	293.68	21.57	56.37
(32,64]	2,301.99	42.6	28.56	1,280.17	43.14	43.61	632.45	42.64	40.84	118.26	44.49	29.04
(64,128]	2,241.57	92.2	71.58	1,112.41	94.43	93.50	419.42	95.19	81.83	177.61	95.10	58.47
(128,256]	2,235.40	170.0	114.75	1,133.85	173.87	148.71	522.21	185.83	226.96	128.09	180.41	77.92

Table 5: Synthetic data: MAE for $\hat{f}^{(DP)}$, $\hat{f}^{(NIGP)}$ and $\hat{f}^{(PYP)}$, case $J = 160, N = 4$.

Bins of x_{m+1}	$\mathcal{Z}_{1,3}$			$\mathcal{Z}_{1,6}$			$\mathcal{Z}_{1,9}$			$\mathcal{Z}_{2,2}$		
	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$
(0,1]	2,212.1	590.48	126.11	1,262.0	146.11	36.60	424.80	130.90	53.37	154.70	47.10	17.66
(1,2]	2,339.8	359.57	158.41	1,332.7	63.21	72.87	552.00	65.00	189.88	182.70	2.01	18.76
(2,4]	2,270.9	69.42	54.81	1,277.8	301.89	176.24	487.30	163.55	42.74	184.70	97.15	25.86
(4,8]	2,234.6	339.95	92.71	1,375.7	579.94	98.10	545.20	243.08	26.14	252.50	62.70	21.57
(8,16]	2,213.3	313.37	62.11	1,165.7	152.53	59.33	493.20	196.20	102.04	247.30	29.70	22.56
(16,32]	2,283.0	23.30	111.41	1,217.2	22.94	84.80	535.50	154.30	31.80	295.90	190.92	23.36
(32,64]	2,305.7	133.09	172.81	1,284.6	209.13	63.20	637.80	150.05	37.70	120.60	71.86	24.01
(64,128]	2,244.5	102.43	57.11	1,120.2	118.42	73.93	425.10	198.60	36.43	180.30	113.75	22.57
(128,256]	2,237.4	294.43	118.11	1,141.3	573.12	48.07	525.90	267.15	41.26	129.70	129.70	22.73

Table 6: Synthetic data: MAE for $\hat{f}^{(CMS)}$, $\hat{f}^{(CMM)}$ and $\hat{f}^{(BDCM)}$, case $J = 160, N = 4$.

proved to improve on some aspects of CMS, it has the major drawback that the posterior distribution of a point query is obtained through a “constructive” proof that builds upon arguments that are tailored to the DP prior, namely arguments that are not usable for other nonparametric priors. In this paper, we presented a “Bayesian” proof of the CMS-DP, that is we computed the (regular) conditional distribution of $f_{X_{m+1}}$ given $C_{n,h_n}(X_{m+1})$, and we showed that such a distribution coincides with the posterior distribution obtained in Cai et al. (2018) through the “constructive” proof. Besides strengthening the BNP approach of Cai et al. (2018) through rigorous arguments, our proof improve its flexibility by avoiding the use of properties that are peculiar to the DP, thus paving the way to go beyond the use of the DP prior. This first result led to develop a novel learning-augmented CMS under power-law data streams, referred to as CMS-PYP, which relies on BNP modeling of the

Bins of x_{m+1}	$J = 12000$ and $N = 2$				$J = 8000$ and $N = 4$			
	20 Newsgroups		Enron		20 Newsgroups		Enron	
	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$
(0,1]	46.39	1.22	12.20	0.90	53.39	0.99	70.98	1.18
(1,2]	16.60	1.85	13.80	1.86	30.49	2.10	47.38	2.05
(2,4]	38.40	3.24	61.49	3.60	32.49	3.66	52.49	4.14
(4,8]	59.39	5.04	88.39	7.68	38.69	6.59	53.08	6.13
(8,16]	54.29	10.90	23.40	12.85	25.29	13.17	56.98	11.55
(16,32]	17.80	20.89	55.09	23.97	24.99	22.69	89.98	19.29
(32,64]	40.79	43.93	128.48	48.94	39.69	46.42	108.37	47.61
(64,128]	25.99	77.72	131.08	78.51	22.09	91.15	55.67	70.81
(128,256]	13.59	170.82	50.68	165.28	25.79	191.35	80.76	172.07

 Table 7: 20 Newsgroups and Enron real data: MAE for $\hat{f}^{(PYP)}$ and $\hat{f}^{(DP)}$.

Bins of x_{m+1}	20 Newsgroups			Enron		
	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$
(0,1]	46.4	5.41	18.04	12.2	0.90	29.20
(1,2]	16.6	2.16	56.00	13.8	2.00	28.30
(2,4]	38.4	7.91	23.14	61.5	9.90	26.00
(4,8]	59.4	35.70	31.42	88.4	17.32	24.30
(8,16]	54.3	45.40	22.72	23.4	9.52	118.70
(16,32]	17.8	20.99	27.72	55.1	21.00	21.60
(32,64]	40.8	58.86	48.24	128.5	134.47	24.40
(64,128]	26.0	91.59	23.20	131.1	110.27	27.60
(128,256]	13.6	186.92	28.05	50.7	140.43	110.40

 Table 8: Real data: MAE for $\hat{f}^{(CMS)}$, $\hat{f}^{(CMM)}$ and $\hat{f}^{(BDCM)}$, case $J = 12000$ and $N = 2$.

Bins of x_{m+1}	20 Newsgroups			Enron		
	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$	$\hat{f}^{(CMS)}$	$\hat{f}^{(CMM)}$	$\hat{f}^{(BDCM)}$
(0,1]	53.4	4.50	25.90	71.0	51.00	31.56
(1,2]	30.5	2.00	20.30	47.4	27.20	21.26
(2,4]	32.5	4.80	19.00	52.5	3.90	22.42
(4,8]	38.7	6.23	22.40	53.1	10.50	34.88
(8,16]	25.3	13.50	24.50	57.0	22.20	14.74
(16,32]	25.0	21.60	22.20	90.0	20.60	27.94
(32,64]	39.7	39.22	24.00	108.4	61.38	49.38
(64,128]	22.1	86.32	19.50	55.7	66.50	21.12
(128,256]	25.8	183.96	26.30	80.8	90.20	34.10

 Table 9: Real data: MAE for $\hat{f}^{(CMS)}$, $\hat{f}^{(CMM)}$ and $\hat{f}^{(BDCM)}$, case $J = 8000$ and $N = 4$.

Bins of x_{m+1}	WikiText-103				1BWLMB			
	$J = 50000$ and $N = 2$		$J = 35000$ and $N = 4$		$J = 140000$ and $N = 2$		$J = 100000$ and $N = 4$	
	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$	$\hat{f}^{(DP)}$	$\hat{f}^{(PYP)}$
(0,1]	97.30	43.15	119.59	40.40	702.70	41.16	156.50	34.99
(1,2]	61.30	34.07	145.39	31.30	104.10	35.18	138.10	35.89
(2,4]	157.70	34.29	91.79	31.50	50.10	37.55	65.00	35.50
(4,8]	192.59	35.45	120.49	34.00	552.59	33.92	49.40	37.50
(8,16]	191.59	41.38	111.09	32.20	176.40	34.75	43.70	32.70
(16,32]	195.19	33.10	127.09	46.50	143.40	38.09	97.40	35.09
(32,64]	248.29	34.29	102.09	44.30	600.30	37.31	168.90	40.39
(64,128]	632.19	37.71	208.29	42.40	143.40	44.31	89.90	41.50
(128,256]	107.69	42.09	140.29	59.90	485.29	48.99	58.60	55.49

Table 10: WikiText-103 and 1BWLMB real data: MAE for $\hat{f}^{(PYP)}$ and $\hat{f}^{(DP)}$.

data stream of tokens via a PYP prior. Under this more general framework, we applied the arguments of the ‘‘Bayesian’’ proof of the CMS-DP, suitably adapted to the PYP prior, to compute the posterior distribution of a point query, given the hashed data. Both the CMS-DP and the CMS-PYP have been also investigated with respect to large sample asymptotic behaviours of their corresponding posterior distributions. Applications to synthetic data and real textual data revealed that the CMS-PYP outperforms the CMS and the CMS-DP in estimating low-frequency tokens, and it is competitive with respect to the CMM and the BDCM.

Our ‘‘Bayesian’’ proof of the CMS-DP can be extended to deal with more general queries. Of notable interest is the problem of estimating the overall frequency of $s \geq 1$ tokens in the stream, also referred to as s -range query, which generalizes the point query (Cormode and Yi, 2020, Chapter 3). For $m \geq 1$ let $x_{1:m}$ be a stream of \mathcal{V} -valued tokens, and for positive integers J and N let h_1, \dots, h_N , with $h_n : \mathcal{V} \rightarrow [J]$, be random hash functions that are i.i.d. from a pairwise independent hash family \mathcal{H} . Then, assuming $x_{1:m}$ to be available through the hashed data $\{(C_{n,1}, \dots, C_{n,J})\}_{n \in [N]}$, the goal is to estimate, or recover, the vector of frequencies $(f_{x_{m+1}}, \dots, f_{x_{m+s}})$ of s new tokens $(x_{m+1}, \dots, x_{m+s})$ in $x_{1:m}$, with $f_{x_{m+r}}$ being defined as

$$f_{x_{m+r}} = \sum_{i=1}^m \mathbb{1}_{\{x_i\}}(x_{m+r})$$

for $r = 1, \dots, s$, and hence the s -range query $\bar{f}_s = \sum_{1 \leq r \leq s} f_{x_{m+r}}$. The arguments of the ‘‘constructive’’ proof of Cai et al. (2018) exploit the unidimensional nature of point queries, and therefore they cannot be used for the vector $(f_{x_{m+1}}, \dots, f_{x_{m+s}})$ nor for \bar{f}_s . In Appendix I we show how to adapt our ‘‘Bayesian’’ proof to the problem of computing the posterior distribution of $(f_{x_{m+1}}, \dots, f_{x_{m+s}})$, given hashed data, and, as an illustrative example, we present the posterior distribution of $(f_{x_{m+1}}, f_{x_{m+2}})$. We focus on the DP prior, though the same arguments apply to the PYP prior. Unfortunately, the posterior distribution of $(f_{x_{m+1}}, f_{x_{m+2}})$ has a rather complicated form, and for a large m the computational burden for its evaluation becomes overwhelming. We defer to future work the study of a large sample behaviour of the posterior distribution, with the aim of obtaining a simple approximated version of it.

Our work paves the way to some fruitful directions for future research in the context of the BNP approach to obtain learning-augmented CMSs. Investigating large sample asymptotic properties of the CMS-DP and CMS-PYP would be of interest, especially with

the aim of obtaining simple approximated versions of the posterior distributions (8) and (31). For a single hash function, i.e. $N = 1$, Proposition 2 and Proposition 5, as well as Equation (28), provide results in this direction. However, it would be of greater interest to consider corresponding results for an arbitrary N , that is for the posterior distributions (8) and (31). Our conjecture is that, under suitable assumptions, the large m limiting posterior distribution of a rescaled point query reduces to a distribution that involves only the minimum of the hashed frequencies, i.e. $\min\{c_1, \dots, c_N\}$, thus making a link with the CMS. In this respect, it would be interesting to obtain some form of central limit theorem for the posterior distributions (8) and (31). For $\alpha = 0$, Cai et al. (2018) showed that the posterior mode may recover the CMS estimate of Cormode and Muthukrishnan (2005), while other CMS-DP estimates may be viewed as CMS estimates with shrinkage; it is natural to ask whether there exists a similar interplay between the CMS-PYP and variations of the CMS for power-law data streams, e.g. the CMM. Other directions of interest consist in using the CMS-DP and CMS-PYP for large-scale streaming algorithms, e.g., for large text or streaming graphs applications (Cormode et al., 2012), as well as to accommodate nonlinear update operations, such as the conservative update (Cormode and Yi, 2020, Chapter 3).

It remains open the problem of extending Theorem 3 to other nonparametric priors than the PYP prior, such as Gibbs-type priors and, in general, nonparametric priors arising from the normalization of completely random measures (James, 2002; Prünster, 2002; Pitman, 2003; Regazzini et al., 2003). We expect that an extension to Gibbs-type priors may require a non trivial effort to overcome the lack of the critical property (26), while still exploiting the general strategy underlying the “Bayesian” proof. Instead, an extension to more general priors than Gibbs-type priors might require a different strategy of proof. An extension of Theorem 3 to general classes of nonparametric priors, which include the PYP prior as a special case, may be useful to discover novel peculiar properties of the PYP prior within such classes, as well as “sufficientness” postulates with respect to the hashed data. A further interesting problem consists in extending Theorem 3 with respect to the use of the sampling information. Theorem 3 provides the posterior distribution of $f_{X_{m+1}}$ given $C_{n,h_n(X_{m+1})}$, thus making use of sole sampling information contained in the bucket $C_{n,h_n(X_{m+1})}$ i.e. the bucket where X_{m+1} is hashed. While this approach is in line with the CMS of Cormode and Muthukrishnan (2005), from a statistical perspective it would be more reasonable to make use of the sampling information contained in all the buckets $(C_{n,1}, \dots, C_{n,J})$. In particular, one may consider the computation of the posterior distribution of $f_{X_{m+1}}$ given $(C_{n,1}, \dots, C_{n,J})$ and $h(X_{m+1})$. We conjecture that the DP prior is the sole Gibbs-type prior for which the posterior distribution of $f_{X_{m+1}}$ given $(C_{n,1}, \dots, C_{n,J})$ and $h(X_{m+1})$ coincides with the posterior distribution of $f_{X_{m+1}}$ given $C_{n,h_n(X_{m+1})}$. This would provide a “sufficientness” postulate for the DP prior with respect to hashed data. Work on this is ongoing.

Appendix A. Proof of Proposition 2

The proof of Equation (17) is straightforward, and it follows from Equation (11) by means of the definition of Beta-Binomial distribution (Johnson et al., 2005, Chapter 6). With regards to the proof of Equation (16), for $t \in \mathbb{R}^+$ and $u \in \mathbb{N}_0$, let $(t)_{[u]} := \prod_{0 \leq i \leq u-1} (t - i)$

denote the falling factorial of t of order u , with the proviso that $(t)_{[0]} := 1$. In particular, $(t)_{(u)} = (-1)^u (-t)_{[u]}$. Recall that the (u, v) -th Stirling number of the second type, here denoted by $S(u, v)$, is defined as the v -th coefficient in the expansion of t^u into falling factorials, i.e. $t^u = \sum_{0 \leq v \leq u} S(u, v) (t)_{[v]}$; moreover, it is assumed: $S(0, 0) = 1$, $S(u, 0) = 0$ for $u > 0$ and $S(u, v) = 0$ for $v > u$. Then, for $r \geq 1$

$$\begin{aligned}
 \mathbb{E} \left[\left(\frac{F_{X_{m+1}}}{c_n} \right)^r \right] &= c_n^{-r} \sum_{l=0}^{c_n} l^r \frac{\frac{\theta}{J}}{\frac{\theta}{J} + c_n} \frac{(c_n - l + 1)_{(l)}}{(\frac{\theta}{J} + c_n - l)_{(l)}} \\
 &= c_n^{-r} \sum_{l=0}^{c_n} \left(\sum_{k=0}^r S(r, k) (l)_{[k]} \right) \frac{\frac{\theta}{J}}{\frac{\theta}{J} + c_n} \frac{(c_n - l + 1)_{(l)}}{(\frac{\theta}{J} + c_n - l)_{(l)}} \\
 &= c_n^{-r} \sum_{k=0}^r S(r, k) \sum_{l=0}^{c_n} \frac{\frac{\theta}{J}}{\frac{\theta}{J} + c_n} \frac{(c_n - l + 1)_{(l)}}{(\frac{\theta}{J} + c_n - l)_{(l)}} (l)_{[k]} \\
 &= c_n^{-r} \sum_{k=0}^r S(r, k) \frac{\frac{\theta}{J}}{\frac{\theta}{J} + c_n} \frac{\Gamma(c_n + 1)}{\Gamma(\frac{\theta}{J} + c_n)} k! \sum_{l=0}^{c_n} \binom{c_n - l}{k} \frac{\Gamma(\frac{\theta}{J} + l)}{\Gamma(1 + l)} \\
 &= c_n^{-r} \sum_{k=0}^r S(r, k) \frac{\frac{\theta}{J}}{\frac{\theta}{J} + c_n} \frac{\Gamma(c_n + 1)}{\Gamma(\frac{\theta}{J} + c_n)} k! \frac{\Gamma(1 + \frac{\theta}{J} + c_n)}{\Gamma(c_n + 1 - k) (\frac{\theta}{J})_{(k+1)}} \\
 &= c_n^{-r} \frac{\theta}{J} \sum_{k=0}^r S(r, k) \frac{\Gamma(k + 1) \Gamma(c_n + 1)}{\Gamma(c_n + 1 - k) (\frac{\theta}{J})_{(k+1)}}.
 \end{aligned}$$

By a direct application of Stirling formula for the ratio of Gamma functions, as $c_n \rightarrow +\infty$ it holds

$$\begin{aligned}
 \mathbb{E} \left[\left(\frac{F_{X_{m+1}}}{c_n} \right)^r \right] &\approx c_n^{-r} \frac{\theta}{J} \sum_{k=0}^r S(r, k) \frac{\Gamma(k + 1)}{(\frac{\theta}{J})_{(k+1)}} c_n^k \\
 &\approx \frac{\theta}{J} \frac{\Gamma(r + 1)}{(\frac{\theta}{J})_{(r+1)}} \\
 &= \frac{\Gamma(r + 1) \Gamma(\frac{\theta}{J} + 1)}{\Gamma(\frac{\theta}{J} + r + 1) \Gamma(1)} \\
 &= \mathbb{E}[B_{1, \frac{\theta}{J}}^r].
 \end{aligned}$$

for any $r \geq 1$. This completes the proof of (16), and hence the proof of Proposition 2 is completed.

Appendix B. Proof of Theorem 3

The proof is along lines similar to the ‘‘Bayesian’’ proof of Section 2. To simplify the notation, we remove the subscript n from h_n and c_n . Then, we are interest in computing the posterior distribution

$$\Pr[f_{X_{m+1}} = l \mid C_{h(X_{m+1})} = c] \tag{34}$$

$$\begin{aligned}
 &= \Pr \left[f_{X_{m+1}} = l \mid \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] \\
 &= \frac{\Pr \left[f_{X_{m+1}} = l, \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right]}{\Pr \left[\sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right]}
 \end{aligned}$$

for $l = 0, 1, \dots, m$. The independence between h_n and $X_{1:m}$ allows us to invoke the “freezing lemma” (Baldi, 2017, Lemma 4.1), according to which we can treat h_n as it was fixed, i.e. non-random. We start with the denominator of (34). Uniformity of the hash function h implies that h induces a (fixed) J -partition $\{B_1, \dots, B_J\}$ of \mathcal{V} such that $B_j = \{v \in \mathcal{V} : h(v) = j\}$ and $\nu(B_j) = J^{-1}$ for $j = 1, \dots, J$. Accordingly, we can write the denominator of (34) as

$$\begin{aligned}
 &\Pr \left[\sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] \tag{35} \\
 &= J \binom{m}{c} \mathbb{E}[(P(B_j))^{c+1} (1 - P(B_j))^{m-c}] \\
 &= J \binom{m}{c} \mathbb{E}[(P(B_j))^{c+1} P(\bar{B}_j)^{m-c}] \\
 &= J \binom{m}{c} \sum_{i=0}^{c+1} \sum_{j=0}^{m-c} \frac{\binom{\theta}{\alpha}_{(i+j)}}{(\theta)_{(m+1)}} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c+1, i; \alpha) \mathcal{C}(m-c, j; \alpha),
 \end{aligned}$$

where the last equality follows from Sangalli (2006, Equation 3.3). This completes the study of the denominator of (34). Now, we consider the numerator of (34). Let us define the event $B(m, l) = \{X_1 = \dots = X_l = X_{m+1}, \{X_{l+1}, \dots, X_m\} \cap \{X_{m+1}\} = \emptyset\}$. In particular, we write

$$\begin{aligned}
 &\Pr \left[f_{X_{m+1}} = l, \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] \tag{36} \\
 &= \binom{m}{l} \Pr \left[B(m, l), \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] \\
 &= \binom{m}{l} \Pr \left[B(m, l), \sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right].
 \end{aligned}$$

That is, the distribution of $(f_{X_{m+1}}, C_j)$ is completely determined by the knowledge of the distribution of (X_1, \dots, X_{m+1}) . Let $\Pi(s, k)$ denote the set of all possible partitions of the set $\{1, \dots, s\}$ into k disjoint subsets π_1, \dots, π_k such that n_i is the cardinality of π_i . In particular, from Sangalli (2006, Equation 3.5), for any measurable A_1, \dots, A_{m+1} we have that

$$\Pr[X_1 \in A_1, \dots, X_{m+1} \in A_{m+1}] = \sum_{k=1}^{m+1} \frac{\prod_{i=0}^{k-1} (\theta + i\alpha)}{(\theta)_{(m+1)}}$$

$$\times \sum_{(\pi_1, \dots, \pi_k) \in \Pi(m+1, k)} \prod_{i=1}^k (1 - \alpha)_{(n_i-1)} \nu(\cap_{r \in \pi_i} A_r)$$

for $m \geq 1$. Let \mathcal{V} be the Borel σ -algebra of \mathcal{V} . Let $\nu_{\pi_1, \dots, \pi_k}$ be a probability measure on $(\mathcal{V}^{m+1}, \mathcal{V}^{m+1})$ defined as

$$\nu_{\pi_1, \dots, \pi_k}(A_1 \times \dots \times A_{m+1}) = \prod_{1 \leq i \leq k} \nu(\cap_{r \in \pi_i} A_r),$$

and attaching to $B(m, l)$ a value that is either 0 or 1. In particular, $\nu_{\pi_1, \dots, \pi_k}(B(m, l)) = 1$ if and only if one of the π_i 's is equal to the set $\{1, \dots, l, m+1\}$. Hence, based on the measure $\nu_{\pi_1, \dots, \pi_k}$, we write

$$\begin{aligned} & \Pr \left[B(m, l), \sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right] \\ &= \sum_{k=2}^{m-l+1} \frac{\prod_{i=0}^{k-1} (\theta + i\alpha)}{(\theta)_{(m+1)}} \\ & \quad \times \sum_{(\pi_1, \dots, \pi_{k-1}) \in \Pi(m-l, k-1)} (1 - \alpha)_{(l)} \prod_{i=1}^{k-1} (1 - \alpha)_{(n_i-1)} \nu_{\pi_1, \dots, \pi_k} \left(\sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right) \\ &= \theta \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)_{(l)} \\ & \quad \times \sum_{r=1}^{m-l} \frac{\prod_{i=0}^{r-1} (\theta + \alpha + i\alpha)}{(\theta + \alpha)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (1 - \alpha)_{(n_i-1)} \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=1}^{m-l} \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right). \end{aligned}$$

Now,

$$\sum_{r=1}^{m-l} \frac{\prod_{i=0}^{r-1} (\theta + \alpha + i\alpha)}{(\theta + \alpha)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (1 - \alpha)_{(n_i-1)} \nu_{\pi_1, \dots, \pi_r}(\cdot)$$

is the distribution of a random sample (X_1, \dots, X_{m-l}) from $P \sim \text{PYP}(\alpha, \theta + \alpha; \nu)$. This a consequence of the quasi-conjugacy property of the PYP prior (Lijoi et al., 2008). Again, the distribution of (X_1, \dots, X_{m-l}) is given in Sangalli (2006, Equation 3.5). In particular, we write

$$\begin{aligned} & \Pr \left[B(m, l), \sum_{i=l+1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right] \\ &= \theta \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)_{(l)} \\ & \quad \times \sum_{r=1}^{m-l} \frac{\prod_{i=0}^{r-1} (\theta + \alpha + i\alpha)}{(\theta + \alpha)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (1 - \alpha)_{(n_i-1)} \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=1}^{m-l} \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c - l \right) \\ &= \theta \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)_{(l)} \binom{m-l}{c-l} \mathbb{E}[(P(B_j))^{c-l} (1 - P(B_j))^{m-c}] \end{aligned}$$

$$\begin{aligned}
 &= \theta \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)_{(l)} \binom{m-l}{c-l} \mathbb{E}[(P(B_j))^{c-l} P(\bar{B}_j)^{m-c}] \\
 &= \theta \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)_{(l)} \binom{m-l}{c-l} \sum_{i=0}^{c-l} \sum_{j=0}^{m-c} \frac{\binom{\theta+\alpha}{\alpha}_{(i+j)}}{(\theta + \alpha)_{(m-l)}} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c-l, i; \alpha) \mathcal{C}(m-c, j; \alpha),
 \end{aligned}$$

where the second identity and the last identity follow from an application of Sangalli (2006, Proposition 3.1) and Sangalli (2006, Equation 3.3), respectively, under the PYP prior; see also the formule displayed at page 469 of Sangalli (2006)). Accordingly, from (36) we can write that

$$\begin{aligned}
 \Pr \left[f_{X_{m+1}} = l, \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] & \tag{37} \\
 &= \binom{m}{l} \theta \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)_{(l)} \binom{m-l}{c-l} \\
 &\quad \times \sum_{i=0}^{c-l} \sum_{j=0}^{m-c} \frac{\binom{\theta+\alpha}{\alpha}_{(i+j)}}{(\theta + \alpha)_{(m-l)}} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c-l, i; \alpha) \mathcal{C}(m-c, j; \alpha).
 \end{aligned}$$

This completes the study of the numerator of (34). By combining (34) with (35) and (37) we obtain

$$\begin{aligned}
 \Pr \left[f_{X_{m+1}} = l \mid \sum_{i=1}^m \mathbb{1}_{\{h(X_i)\}}(h(X_{m+1})) = c \right] & \tag{38} \\
 &= \frac{\theta}{J} \binom{c}{l} \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)_{(l)} \\
 &\quad \times \frac{\sum_{i=0}^{c-l} \sum_{j=0}^{m-c} \frac{\binom{\theta+\alpha}{\alpha}_{(i+j)}}{(\theta+\alpha)_{(m-l)}} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c-l, i; \alpha) \mathcal{C}(m-c, j; \alpha)}{\sum_{i=0}^{c+1} \sum_{j=0}^{m-c} \frac{\binom{\theta}{\alpha}_{(i+j)}}{(\theta)_{(m+1)}} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c+1, i; \alpha) \mathcal{C}(m-c, j; \alpha)} \\
 &= \frac{\theta}{J} \binom{c}{l} (1 - \alpha)_{(l)} \frac{\sum_{i=0}^{c-l} \sum_{j=0}^{m-c} \frac{\binom{\theta+\alpha}{\alpha}_{(i+j)}}{(\theta+\alpha)_{(m-l)}} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c-l, i; \alpha) \mathcal{C}(m-c, j; \alpha)}{\sum_{i=0}^{c+1} \sum_{j=0}^{m-c} \frac{\binom{\theta}{\alpha}_{(i+j)}}{(\theta)_{(m+1)}} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c+1, i; \alpha) \mathcal{C}(m-c, j; \alpha)}.
 \end{aligned}$$

for $l = 0, 1, \dots, c$. By an application of Charalambides (2005, Equation 2.56 and Equation 2.60) it is easy to show that (38) is a proper distribution on $\{0, 1, \dots, c\}$. The proof is completed.

Appendix C. Theorem 1 from Theorem 3 with $\alpha = 0$

We show how Theorem 3 reduces to Theorem 1 by setting $\alpha = 0$. First, we recall two identities involving the generalized factorial coefficient $\mathcal{C}(m, k; \alpha)$ and the signless Stirling number of the first type. See Charalambides (2005, Chapter 2) for details. In particular, it holds

$$\sum_{k=0}^m a^k |s(m, k)| = (a)_{(m)} \tag{39}$$

for $a > 0$, and

$$\lim_{\alpha \rightarrow 0} \frac{\mathcal{C}(m, k; \alpha)}{\alpha^k} = |s(m, k)|. \quad (40)$$

Hereafter, we apply the identities (39) and (40) in order to show that Theorem 3 reduces to Theorem 1 by setting $\alpha = 0$. In this respect, we rewrite the posterior distribution (24) as follows

$$\begin{aligned} & \Pr[f_{X_{m+1}} = l \mid C_{n, h_n}(X_{m+1}) = c_n] \\ &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n - l, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n + 1, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)}. \end{aligned}$$

Then,

$$\begin{aligned} & \lim_{\alpha \rightarrow 0} \Pr[f_{X_{m+1}} = l \mid C_{n, h_n}(X_{m+1}) = c_n] \\ &= \lim_{\alpha \rightarrow 0} \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \\ & \quad \times \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \alpha^{i+j} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \frac{\mathcal{C}(c_n-l, i; \alpha)}{\alpha^i} \frac{\mathcal{C}(m-c_n, j; \alpha)}{\alpha^j}}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \alpha^{i+j} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \frac{\mathcal{C}(c_n+1, i; \alpha)}{\alpha^i} \frac{\mathcal{C}(m-c_n, j; \alpha)}{\alpha^j}} \\ & \text{[by the identity (40)]} \\ &= \frac{\theta}{J} \binom{c_n}{l} l! \\ & \quad \times \frac{\sum_{i=0}^{c_n-l} \left(\frac{\theta}{J}\right)^i |s(c_n - l, i)| \sum_{j=0}^{m-c_n} \left(\theta \left(1 - \frac{1}{J}\right)\right)^j |s(m - c_n, j)|}{\sum_{i=0}^{c_n+1} \left(\frac{\theta}{J}\right)^i |s(c_n + 1, i)| \sum_{j=0}^{m-c_n} \left(\theta \left(1 - \frac{1}{J}\right)\right)^j |s(m - c_n, j)|} \\ & \text{[by the identity (39)]} \\ &= \frac{\theta}{J} \binom{c_n}{l} l! \frac{\left(\frac{\theta}{J}\right)_{(c_n-l)} \left(\theta \left(1 - \frac{1}{J}\right)\right)_{(m-c_n)}}{\left(\frac{\theta}{J}\right)_{(c_n-l)} \left(\theta \left(1 - \frac{1}{J}\right)\right)_{(m-c_n)}} \\ &= \frac{\theta}{J} \frac{\Gamma(c_n + 1) \Gamma(c_n - l + \frac{\theta}{J})}{\Gamma(c_n - l + 1) \Gamma(\frac{\theta}{J} + c_n + 1)} \\ &= \frac{\frac{\theta}{J}}{\frac{\theta}{J} + c_n} \frac{(c_n - l + 1)^{(l)}}{(\frac{\theta}{J} + c_n - l)^{(l)}}, \end{aligned}$$

which is the expression for the posterior distribution stated in Theorem 1. The proof is completed.

Appendix D. Proof of Equation (27)

Let $X_{1:m}$ be a random sample from $P \sim \text{PYP}(\alpha, \theta; \nu)$, with $\alpha \in [0, 1)$ and $\theta > -\alpha$, and let K_m be the number of distinct types in $X_{1:m}$. We recall from (19) that for $k = 1, \dots, m$ it holds

$$\Pr[K_m = k] = \frac{\left(\frac{\theta}{\alpha}\right)^{(k)}}{(\theta)^{(m)}} \mathcal{C}(m, k; \alpha).$$

Now, assuming $c_n > 0$ and $m - c_n > 0$, we rewrite the posterior distribution of Theorem 3 in terms of the distribution of K_m . In particular, for any $l = 0, 1, \dots, c_n - 1$ we can write that

$$\begin{aligned}
 & \Pr[fX_{m+1} = l \mid C_{n,h_n}(X_{m+1}) = c_n] \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n - l, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n + 1, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)} \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{(\theta)_{(c_n-l)} (\theta)_{(m-c_n)}}{(\theta)_{(c_n+1)} (\theta)_{(m-c_n)}} \\
 &\quad \times \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \left(\frac{\theta}{\alpha}\right)_{(i)} \left(\frac{\theta}{\alpha}\right)_{(j)} (\theta)_{(c_n-l)} \mathcal{C}(c_n - l, i; \alpha) \frac{(\frac{\theta}{\alpha})_{(j)}}{(\theta)_{(m-c_n)}} \mathcal{C}(m - c_n, j; \alpha)}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \left(\frac{\theta}{\alpha}\right)_{(i)} \left(\frac{\theta}{\alpha}\right)_{(j)} (\theta)_{(c_n+1)} \mathcal{C}(c_n + 1, i; \alpha) \frac{(\frac{\theta}{\alpha})_{(j)}}{(\theta)_{(m-c_n)}} \mathcal{C}(m - c_n, j; \alpha)} \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{(\theta)_{(c_n-l)}}{(\theta)_{(c_n+1)}} \\
 &\quad \times \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \Pr[K_{c_n-l} = i] \Pr[K_{m-c_n} = j]}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \Pr[K_{c_n+1} = i] \Pr[K_{m-c_n} = j]} \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{(\theta)_{(c_n-l)} \mathbb{E} \left[\frac{(\frac{\theta+\alpha}{\alpha})_{(K_{c_n-l} + K_{m-c_n})} \left(\frac{1}{J}\right)^{K_{c_n-l}} \left(1 - \frac{1}{J}\right)^{K_{m-c_n}}}{(\frac{\theta}{\alpha})_{(K_{c_n-l})} (\frac{\theta}{\alpha})_{(K_{m-c_n})}} \right]}{(\theta)_{(c_n+1)} \mathbb{E} \left[\frac{(\frac{\theta}{\alpha})_{(K_{c_n+1} + K_{m-c_n})} \left(\frac{1}{J}\right)^{K_{c_n+1}} \left(1 - \frac{1}{J}\right)^{K_{m-c_n}}}{(\frac{\theta}{\alpha})_{(K_{c_n+1})} (\frac{\theta}{\alpha})_{(K_{m-c_n})}} \right]} \\
 &= \frac{\frac{\alpha}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \mathbb{E} \left[\frac{\Gamma(\frac{\theta+\alpha}{\alpha} + K_{c_n-l} + K_{m-c_n})}{\Gamma(\frac{\theta}{\alpha} + K_{c_n-l}) \Gamma(\frac{\theta}{\alpha} + K_{m-c_n})} \left(\frac{1}{J}\right)^{K_{c_n-l}} \left(1 - \frac{1}{J}\right)^{K_{m-c_n}} \right]}{(\theta + c_n - l)^{(l+1)} \mathbb{E} \left[\frac{\Gamma(\frac{\theta}{\alpha} + K_{c_n+1} + K_{m-c_n})}{\Gamma(\frac{\theta}{\alpha} + K_{c_n+1}) \Gamma(\frac{\theta}{\alpha} + K_{m-c_n})} \left(\frac{1}{J}\right)^{K_{c_n+1}} \left(1 - \frac{1}{J}\right)^{K_{m-c_n}} \right]},
 \end{aligned}$$

where K_{c_n-l} and K_{m-c_n} in the numerator are independent random variables for any $l = 0, 1, \dots, c_n - 1$, and K_{c_n+1} and K_{m-c_n} in the denominator are independent random variables. For $l = c_n$

$$\begin{aligned}
 & \Pr[fX_{m+1} = c_n \mid C_{n,h_n}(X_{m+1}) = c_n] \\
 &= \frac{\theta}{J} (1 - \alpha)^{(c_n)} \frac{\sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(j)} \left(1 - \frac{1}{J}\right)^j \mathcal{C}(m - c_n, j; \alpha)}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n + 1, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)} \\
 &= \frac{\theta}{J} (1 - \alpha)^{(c_n)} \frac{(\theta)_{(m-c_n)}}{(\theta)_{(c_n+1)} (\theta)_{(m-c_n)}}
 \end{aligned}$$

$$\begin{aligned}
 & \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(j)} \frac{\left(\frac{1-\frac{1}{J}}{J}\right)^j \left(\frac{\theta}{\alpha}\right)_{(j)}}{\left(\frac{\theta}{\alpha}\right)_{(j)} (\theta)_{(m-c_n)}} \mathcal{C}(m-c_n, j; \alpha) \\
 & \times \frac{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(\frac{1-\frac{1}{J}}{J}\right)^j \left(\frac{\theta}{\alpha}\right)_{(i)} \frac{\mathcal{C}(c_n+1, i; \alpha)}{(\theta)_{(c_n+1)}} \frac{\left(\frac{\theta}{\alpha}\right)_{(j)} \mathcal{C}(m-c_n, j; \alpha)}{(\theta)_{(m-c_n)}}}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(\frac{1-\frac{1}{J}}{J}\right)^j \frac{\Pr[K_{c_n+1} = i] \Pr[K_{m-c_n} = j]}{\left(\frac{\theta}{\alpha}\right)_{(i)} \left(\frac{\theta}{\alpha}\right)_{(j)}}} \\
 & = \frac{\theta}{J} (1-\alpha)_{(c_n)} \frac{1}{(\theta)_{(c_n+1)}} \\
 & \times \frac{\sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(j)} \frac{\left(\frac{1-\frac{1}{J}}{J}\right)^j \Pr[K_{m-c_n} = j]}{\left(\frac{\theta}{\alpha}\right)_{(j)}}}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(\frac{1-\frac{1}{J}}{J}\right)^j \Pr[K_{c_n+1} = i] \Pr[K_{m-c_n} = j]} \\
 & = \frac{\theta}{J} (1-\alpha)_{(c_n)} \frac{\mathbb{E} \left[\frac{\left(\frac{\theta+\alpha}{\alpha}\right)_{(K_{m-c_n})} \left(\frac{1-\frac{1}{J}}{J}\right)^{K_{m-c_n}}}{\left(\frac{\theta}{\alpha}\right)_{(K_{m-c_n})}} \right]}{(\theta)_{(c_n+1)} \mathbb{E} \left[\frac{\left(\frac{\theta}{\alpha}\right)_{(K_{c_n+1}+K_{m-c_n})} \left(\frac{1}{J}\right)^{K_{c_n+1}} \left(\frac{1-\frac{1}{J}}{J}\right)^{K_{m-c_n}}}{\left(\frac{\theta}{\alpha}\right)_{(K_{c_n+1})} \left(\frac{\theta}{\alpha}\right)_{(K_{m-c_n})}} \right]} \\
 & = \frac{\frac{\alpha}{J} (1-\alpha)_{(c_n)}}{\Gamma(\theta/\alpha) (\theta)_{(c_n+1)}} \frac{\mathbb{E} \left[\frac{\Gamma\left(\frac{\theta+\alpha}{\alpha}+K_{m-c_n}\right)}{\Gamma\left(\frac{\theta}{\alpha}+K_{m-c_n}\right)} \left(1-\frac{1}{J}\right)^{K_{m-c_n}} \right]}{\mathbb{E} \left[\frac{\Gamma\left(\frac{\theta}{\alpha}+K_{c_n+1}+K_{m-c_n}\right)}{\Gamma\left(\frac{\theta}{\alpha}+K_{c_n+1}\right) \Gamma\left(\frac{\theta}{\alpha}+K_{m-c_n}\right)} \left(\frac{1}{J}\right)^{K_{c_n+1}} \left(1-\frac{1}{J}\right)^{K_{m-c_n}} \right]},
 \end{aligned}$$

where K_{c_n+1} and K_{m-c_n} are independent random variables. This completes the proof of Equation (27).

Appendix E. An alternative expression for Equation (24)

For any $\alpha \in (0, 1)$, an alternative expression for (24) may be given in terms of the distribution of exponentially tilted α -stable random variables (Zolotarev, 1986). In particular, if g_α denotes the density function of a positive α -stable distribution, then for any $c > 0$ an exponentially tilted α -stable random variable is defined as the random variable $T_{\alpha,c}$ whose distribution has density function $f_{S_{\alpha,c}}(x) \propto \exp\{-c^{1/\alpha}x\}g_\alpha(x)\mathbb{1}_{\mathbb{R}^+}(x)$. If $c_n > 0$, then for $l = 0, 1, \dots, c_n$

$$\begin{aligned}
 & \Pr[f_{X_{m+1}} = l \mid C_{n,h_n}(X_{m+1}) = c_n] \\
 & = \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)_{(l)} \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1-\frac{1}{J}\right)^j \mathcal{C}(c_n-l, i; \alpha) \mathcal{C}(m-c_n, j; \alpha)}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1-\frac{1}{J}\right)^j \mathcal{C}(c_n+1, i; \alpha) \mathcal{C}(m-c_n, j; \alpha)} \\
 & = \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)_{(l)} \\
 & \times \frac{\frac{1}{\Gamma\left(\frac{\theta+\alpha}{\alpha}\right)} \int_0^{+\infty} x^{\frac{\theta+\alpha}{\alpha}-1} e^{-x} \left(\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{x}{J}\right)^i \left(x\left(1-\frac{1}{J}\right)\right)^j \mathcal{C}(c_n-l, i; \alpha) \mathcal{C}(m-c_n, j; \alpha)\right) dx}{\frac{1}{\Gamma\left(\frac{\theta}{\alpha}\right)} \int_0^{+\infty} x^{\frac{\theta}{\alpha}-1} e^{-x} \left(\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{x}{J}\right)^i \left(x\left(1-\frac{1}{J}\right)\right)^j \mathcal{C}(c_n+1, i; \alpha) \mathcal{C}(m-c_n, j; \alpha)\right) dx}
 \end{aligned}$$

By means of Favaro et al. (2015, Equation 13) we can write the numerator and the denominator of the previous expression in terms of the distribution of $T_{\alpha,c}$, for suitable choices of

c. That is,

$$\begin{aligned}
 & \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)^{(l)} \\
 & \times \frac{\frac{1}{\Gamma(\frac{\theta+\alpha}{\alpha})} \int_0^{+\infty} x^{\frac{\theta+\alpha}{\alpha}-1} e^{-x} \left(\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{x}{J}\right)^i \left(x(1-\frac{1}{J})\right)^j \mathcal{C}(c_n-l, i; \alpha) \mathcal{C}(m-c_n, j; \alpha) \right) dx}{\frac{1}{\Gamma(\frac{\theta}{\alpha})} \int_0^{+\infty} x^{\frac{\theta}{\alpha}-1} e^{-x} \left(\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{x}{J}\right)^i \left(x(1-\frac{1}{J})\right)^j \mathcal{C}(c_n+1, i; \alpha) \mathcal{C}(m-c_n, j; \alpha) \right) dx} \\
 & = \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)^{(l)} \\
 & \times \frac{\frac{1}{\Gamma(\frac{\theta+\alpha}{\alpha})} \int_0^{+\infty} x^{\frac{\theta+\alpha}{\alpha}-1} e^{-x} \left(\left(\frac{x}{J}\right)^{\frac{c_n-l}{\alpha}} \mathbb{E} \left[T_{\alpha, \frac{x}{J}}^{c_n-l} \right] \left(x(1-\frac{1}{J})\right)^{\frac{m-c_n}{\alpha}} \mathbb{E} \left[T_{\alpha, x(1-\frac{1}{J})}^{m-c_n} \right] \right) dx}{\frac{1}{\Gamma(\frac{\theta}{\alpha})} \int_0^{+\infty} x^{\frac{\theta}{\alpha}-1} e^{-x} \left(\left(\frac{x}{J}\right)^{\frac{c_n+1}{\alpha}} \mathbb{E} \left[T_{\alpha, \frac{x}{J}}^{c_n+1} \right] \left(x(1-\frac{1}{J})\right)^{\frac{m-c_n}{\alpha}} \mathbb{E} \left[T_{\alpha, x(1-\frac{1}{J})}^{m-c_n} \right] \right) dx} \\
 & = \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)^{(l)} \\
 & \times \frac{\frac{(\frac{1}{J})^{\frac{c_n-l}{\alpha}} (1-\frac{1}{J})^{\frac{m-c_n}{\alpha}}}{\Gamma(\frac{\theta+\alpha}{\alpha})} \int_0^{+\infty} x^{\frac{\theta+\alpha+m-l}{\alpha}-1} e^{-x} \left(\mathbb{E} \left[T_{\alpha, \frac{x}{J}}^{c_n-l} \right] \mathbb{E} \left[T_{\alpha, x(1-\frac{1}{J})}^{m-c_n} \right] \right) dx}{\frac{(\frac{1}{J})^{\frac{c_n+1}{\alpha}} (1-\frac{1}{J})^{\frac{m-c_n}{\alpha}}}{\Gamma(\frac{\theta}{\alpha})} \int_0^{+\infty} x^{\frac{\theta+m+1}{\alpha}-1} e^{-x} \left(\mathbb{E} \left[T_{\alpha, \frac{x}{J}}^{c_n+1} \right] \mathbb{E} \left[T_{\alpha, x(1-\frac{1}{J})}^{m-c_n} \right] \right) dx} \\
 & = \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)^{(l)} \tag{41} \\
 & \times \frac{\left(\frac{1}{J}\right)^{\frac{c_n-l}{\alpha}} \left(\frac{\theta+\alpha}{\alpha}\right)_{(m-l)} \int_0^{+\infty} \mathbb{E} \left[T_{\alpha, \frac{x}{J}}^{c_n-l} \right] \mathbb{E} \left[T_{\alpha, x(1-\frac{1}{J})}^{m-c_n} \right] f_{G_{\frac{\theta+\alpha+m-l}{\alpha}, 1}}(x) dx}{\left(\frac{1}{J}\right)^{\frac{c_n+1}{\alpha}} \left(\frac{\theta}{\alpha}\right)_{(m+1)} \int_0^{+\infty} \mathbb{E} \left[T_{\alpha, \frac{x}{J}}^{c_n+1} \right] \mathbb{E} \left[T_{\alpha, x(1-\frac{1}{J})}^{m-c_n} \right] f_{G_{\frac{\theta+m+1}{\alpha}, 1}}(x) dx},
 \end{aligned}$$

where $f_{G_{a,b}}$ is the density function of a Gamma distribution with parameter (a, b) . Equation (41) allows for an MC evaluation of (24), which requires to sample from a Gamma distribution and to sample $T_{\alpha,c}$, for suitable choices of c . See Devroye (2009) and references therein.

Appendix F. Proof Equation (28) and Equation (29)

Under the setting of Theorem 3, we consider $m \rightarrow +\infty$, while c_n is fixed. For any $l = 0, 1, \dots, c_n$

$$\begin{aligned}
 & \Pr[f_{X_{m+1}} = l \mid C_{n, h_n(X_{m+1})} = c_n] \\
 & = \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)^{(l)} \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \mathcal{C}(c_n-l, i; \alpha) \frac{(1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}{\sum_{j=0}^{m-c_n} (1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \mathcal{C}(c_n+1, i; \alpha) \frac{(1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}{\sum_{j=0}^{m-c_n} (1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}}
 \end{aligned}$$

$$\begin{aligned}
 &= \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)_{(l)} \frac{\sum_{i=0}^{c_n-l} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i)} \left(\frac{1}{J}\right)^i \mathcal{C}(c_n-l, i; \alpha) \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha} + i\right)_{(j)} \frac{(1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}{\sum_{j=0}^{m-c_n} (1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}}{\sum_{i=0}^{c_n+1} \left(\frac{\theta}{\alpha}\right)_{(i)} \left(\frac{1}{J}\right)^i \mathcal{C}(c_n+1, i; \alpha) \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha} + i\right)_{(j)} \frac{(1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}{\sum_{j=0}^{m-c_n} (1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}}. \\
 &\tag{42}
 \end{aligned}$$

Now, consider the numerator of the last member in (42). From Dolera and Favaro (2020a, Lemma 2), as $m \rightarrow +\infty$

$$\begin{aligned}
 &\lim_{m \rightarrow +\infty} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha} + i\right)_{(j)} \frac{(1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}{\sum_{j \geq 1} (1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)} \\
 &= \sum_{j \geq 1} \left(\frac{\theta+\alpha}{\alpha} + i\right)_{(j)} e^{-(1-\frac{1}{J})} \frac{(1-\frac{1}{J})^{j-1}}{(j-1)!} \\
 &= e^{-(1-\frac{1}{J})} \frac{\left(\frac{\theta+\alpha}{\alpha} + i\right)}{\left(\frac{1}{J}\right)^{\frac{\theta+\alpha}{\alpha} + i + 1}}. \\
 &\tag{43}
 \end{aligned}$$

Now, consider the denominator of the last member in (42). Again, from Dolera and Favaro (2020a, Lemma 2), as $m \rightarrow +\infty$

$$\begin{aligned}
 &\lim_{m \rightarrow +\infty} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha} + i\right)_{(j)} \frac{(1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)}{\sum_{j=0}^{m-c_n} (1-\frac{1}{J})^j \mathcal{C}(m-c_n, j; \alpha)} \\
 &= \sum_{j \geq 1} \left(\frac{\theta}{\alpha} + i\right)_{(j)} e^{-(1-\frac{1}{J})} \frac{(1-\frac{1}{J})^{j-1}}{(j-1)!} \\
 &= e^{-(1-\frac{1}{J})} \frac{\left(\frac{\theta}{\alpha} + i\right)}{\left(\frac{1}{J}\right)^{\frac{\theta}{\alpha} + i + 1}}. \\
 &\tag{44}
 \end{aligned}$$

Then, by combining (42) with (43) and (44), for any $l = 0, 1, \dots, c_n$, as $m \rightarrow +\infty$ we can write

$$\begin{aligned}
 &\lim_{m \rightarrow +\infty} \Pr[f_{X_{m+1}} = l \mid C_{n, h_n}(X_{m+1}) = c_n] \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)_{(l)} \frac{\sum_{i=0}^{c_n-l} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i)} \left(\frac{1}{J}\right)^i \mathcal{C}(c_n-l, i; \alpha) e^{-(1-\frac{1}{J})} \frac{\left(\frac{\theta+\alpha}{\alpha} + i\right)}{\left(\frac{1}{J}\right)^{\frac{\theta+\alpha}{\alpha} + i + 1}}}{\sum_{i=0}^{c_n+1} \left(\frac{\theta}{\alpha}\right)_{(i)} \left(\frac{1}{J}\right)^i \mathcal{C}(c_n+1, i; \alpha) e^{-(1-\frac{1}{J})} \frac{\left(\frac{\theta}{\alpha} + i\right)}{\left(\frac{1}{J}\right)^{\frac{\theta}{\alpha} + i + 1}}} \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)_{(l)} \frac{\left(\frac{1}{J}\right)^{-\frac{\theta+\alpha}{\alpha}-1} (\theta+\alpha)_{(c_n-l)} \sum_{i=0}^{c_n-l} \frac{\left(\frac{\theta+\alpha}{\alpha}\right)_{(i)}}{(\theta+\alpha)_{(c_n-l)}} \mathcal{C}(c_n-l, i; \alpha) \left(\frac{\theta+\alpha}{\alpha} + i\right)}{\left(\frac{1}{J}\right)^{-\frac{\theta}{\alpha}-1} (\theta)_{(c_n+1)} \sum_{i=0}^{c_n+1} \frac{\left(\frac{\theta}{\alpha}\right)_{(i)}}{(\theta)_{c_n+1}} \mathcal{C}(c_n+1, i; \alpha) \left(\frac{\theta}{\alpha} + i\right)} \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)_{(l)} \frac{\left(\frac{1}{J}\right)^{-\frac{\theta+\alpha}{\alpha}-1} (\theta+\alpha)_{(c_n-l)} \sum_{i=0}^{c_n-l} \frac{\left(\frac{\theta+\alpha}{\alpha}\right)_{(i)}}{(\theta+\alpha)_{(c_n-l)}} \mathcal{C}(c_n-l, i; \alpha) \left(\frac{\theta+\alpha}{\alpha} + i\right)}{\left(\frac{1}{J}\right)^{-\frac{\theta}{\alpha}-1} (\theta)_{(c_n+1)} \sum_{i=0}^{c_n+1} \frac{\left(\frac{\theta}{\alpha}\right)_{(i)}}{(\theta)_{c_n+1}} \mathcal{C}(c_n+1, i; \alpha) \left(\frac{\theta}{\alpha} + i\right)} \\
 &\tag{45}
 \end{aligned}$$

[by Pitman (2006, Equation 3.13)]

$$\begin{aligned}
 &= \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)^{(l)} \frac{\left(\frac{1}{J}\right)^{-\frac{\theta+\alpha}{\alpha}-1} (\theta+\alpha)_{(c_n-l)}}{\left(\frac{1}{J}\right)^{-\frac{\theta}{\alpha}-1} (\theta)_{(c_n+1)}} \frac{\frac{\theta+\alpha}{\alpha} + \frac{(\theta+2\alpha)_{(c_n-l)}}{\alpha(\theta+\alpha+1)_{(c_n-l-1)}} - \frac{\theta+\alpha}{\alpha}}{\frac{\theta}{\alpha} + \frac{(\theta+\alpha)_{(c_n+1)}}{\alpha(\theta+1)_{(c_n+1-1)}} - \frac{\theta}{\alpha}} \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1-\alpha)^{(l)} \left(\frac{1}{J}\right)^{-1} \frac{(\theta+\alpha)_{(c_n-l)}}{(\theta)_{(c_n+1)}} \frac{\frac{(\theta+2\alpha)_{(c_n-l)}}{(\theta+\alpha+1)_{(c_n-l-1)}}}{\frac{(\theta+\alpha)_{(c_n+1)}}{(\theta+1)_{(c_n+1-1)}}}} \\
 &= \binom{c_n}{l} (1-\alpha)^{(l)} \frac{(\theta+2\alpha)_{(c_n-l)}}{(\theta+\alpha+1)_{(c_n)}}.
 \end{aligned}$$

This completes the proof of Equation (28). Equation (29) follows by a direct calculation from (28).

Appendix G. Proof of Proposition 5

Let $B_{a,b}$ be a Beta random variable with parameter (a, b) , and denote by $f_{B_{a,b}}$ the density function of the distribution of $B_{a,b}$. We start by some considerations on the distribution of $B_{a,b}$:

i)

$$\frac{\Gamma(\theta+\alpha+m-l)\Gamma(1-\alpha+l)}{\Gamma(\theta+m+1)} = \int_0^1 t^{\theta+\alpha+m-l-1} (1-t)^{l-\alpha} dt;$$

ii)

$$\frac{\Gamma(\theta+\alpha+m-l)\Gamma(1-\alpha+l)}{\Gamma(\theta+m+1)} = \frac{\Gamma(\theta+\alpha)\Gamma(1-\alpha)}{\Gamma(\theta+1)} \mathbb{E}[B_{\theta+\alpha, 1-\alpha}^{m-l} (1-B_{\theta+\alpha, 1-\alpha}^l)].$$

Moreover, we observe that we can rewrite the numerator and the denominator of (27) as follows

$$\begin{aligned}
 &\int_0^{+\infty} \int_0^{+\infty} g_\alpha(h) g_\alpha(x) x^{-\theta-\alpha} \frac{\left(\frac{h}{x}(J-1)^{\frac{1}{\alpha}}\right)^{m-c}}{\left(\frac{h}{x}(J-1)^{\frac{1}{\alpha}}+1\right)^{\theta+m-l+\alpha}} dx dh \\
 &= \frac{\Gamma(2+\theta/\alpha)}{\Gamma(1+\theta+\alpha)} \int_0^{+\infty} \int_0^{+\infty} f_{S_{\alpha,0}}(h) f_{S_{\alpha,\theta+\alpha}}(x) \frac{\left(\frac{h}{x}(J-1)^{\frac{1}{\alpha}}\right)^{m-c}}{\left(\frac{h}{x}(J-1)^{\frac{1}{\alpha}}+1\right)^{\theta+m-l+\alpha}} dx dh \\
 &= \frac{\Gamma(2+\theta/\alpha)}{\Gamma(1+\theta+\alpha)} \int_0^{+\infty} \left[\frac{x^{m-c}}{(x+1)^{\theta+m-l+\alpha}} \right] f_{Z_{\alpha,\theta+\alpha}}(x) dx
 \end{aligned}$$

and

$$\int_0^{+\infty} \int_0^{+\infty} g_\alpha(h) g_\alpha(x) x^{-\theta} \frac{\left(\frac{h}{x}(J-1)^{\frac{1}{\alpha}}\right)^{m-c}}{\left(\frac{h}{x}(J-1)^{\frac{1}{\alpha}}+1\right)^{\theta+m+1}} dx dh$$

$$\begin{aligned}
 &= \frac{\Gamma(1 + \theta/\alpha)}{\Gamma(1 + \theta)} \int_0^{+\infty} \int_0^{+\infty} f_{S_{\alpha,0}}(h) f_{S_{\alpha,\theta}}(x) \frac{\left(\frac{h}{x}(J-1)\frac{1}{\alpha}\right)^{m-c}}{\left(\frac{h}{x}(J-1)\frac{1}{\alpha} + 1\right)^{\theta+m+1}} dx dh \\
 &= \frac{\Gamma(1 + \theta/\alpha)}{\Gamma(1 + \theta)} \int_0^{+\infty} \left[\frac{x^{m-c}}{(x+1)^{\theta+m+1}} \right] f_{W_{\alpha,\theta}}(x) dx,
 \end{aligned}$$

respectively. First, we prove that the distribution $\Pr[f_{X_{m+1}} \in \cdot | C_{n,h_n}(X_{m+1}) = c_n]$ admits a representation in terms of a suitable mixture of Binomial distribution. In particular, we write

$$\begin{aligned}
 &\Pr[f_{X_{m+1}} = l | C_{n,h_n}(X_{m+1}) = c_n] \\
 &= \frac{\theta}{J} \binom{c_n}{l} (1 - \alpha)^{(l)} \frac{\sum_{i=0}^{c_n-l} \sum_{j=0}^{m-c_n} \left(\frac{\theta+\alpha}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n - l, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)}{\sum_{i=0}^{c_n+1} \sum_{j=0}^{m-c_n} \left(\frac{\theta}{\alpha}\right)_{(i+j)} \left(\frac{1}{J}\right)^i \left(1 - \frac{1}{J}\right)^j \mathcal{C}(c_n + 1, i; \alpha) \mathcal{C}(m - c_n, j; \alpha)} \\
 &= \frac{\theta}{J} \binom{c_n}{l} \frac{(\theta + \alpha)_{(m-l)}}{(\theta)_{(m+1)}} (1 - \alpha)^{(l)} \\
 &\quad \times \frac{\sum_{i=0}^{m-c_n} \binom{m-c}{i} (-1)^{m-c_n-i} \sum_{k=0}^{m-l-i} \frac{\left(\frac{\theta+\alpha}{\alpha}\right)_{(k)}}{(\theta+\alpha)_{(m-l-i)}} \frac{1}{J^k} \mathcal{C}(m-l-i, k; \alpha)}{\sum_{i=0}^{m-c_n} \binom{m-c_n}{i} (-1)^{m-c_n-i} \sum_{k=0}^{m-i+1} \frac{\left(\frac{\theta}{\alpha}\right)_{(k)}}{(\theta)_{(m-i+1)}} \frac{1}{J^k} \mathcal{C}(m-i+1, k; \alpha)} \\
 &= \alpha \binom{c_n}{l} \frac{\Gamma(\theta + \alpha + m - l)}{\Gamma(\theta + m + 1)} (1 - \alpha)^{(l)} \frac{\int_{(0,+\infty)^2} g_{\alpha}(h) g_{\alpha}(x) x^{-\theta-\alpha} \frac{\left(\frac{h}{x}(J-1)\frac{1}{\alpha}\right)^{m-c_n}}{\left(\frac{h}{x}(J-1)\frac{1}{\alpha} + 1\right)^{\theta+m-l+\alpha}} dx dh}{\int_{(0,+\infty)^2} g_{\alpha}(h) g_{\alpha}(x) x^{-\theta} \frac{\left(\frac{h}{x}(J-1)\frac{1}{\alpha}\right)^{m-c_n}}{\left(\frac{h}{x}(J-1)\frac{1}{\alpha} + 1\right)^{\theta+m+1}} dx dh} \\
 &= \binom{c_n}{l} \int_0^1 t^{m-l} (1-t)^l \frac{\int_0^{+\infty} \left[\frac{x^{m-c_n}}{(x+1)^{\theta+m-l+\alpha}} \right] f_{Z_{\alpha,\theta+\alpha}}(x) dx}{\int_0^{+\infty} \left[\frac{x^{m-c_n}}{(x+1)^{\theta+m+1}} \right] f_{W_{\alpha,\theta}}(x) dx} f_{B_{\theta+\alpha,1-\alpha}}(t) dt \\
 &= \frac{1}{D(m, c_n; \alpha, \theta, J)} \int_0^1 \int_0^{+\infty} \left[\binom{c_n}{l} (1-t)^l \left(\frac{t}{x+1}\right)^{c_n-l} \right] \frac{\left(\frac{tx}{x+1}\right)^{m-c_n}}{(x+1)^{\theta+\alpha}} f_{B_{\theta+\alpha,1-\alpha}}(t) f_{Z_{\alpha,\theta+\alpha}}(x) dt dx \\
 &= \frac{1}{D(m, c_n; \alpha, \theta, J)} \int_0^1 \int_0^{+\infty} \left[\binom{c_n}{l} (1-t)^l \left(\frac{t}{x+1}\right)^{c_n-l} \right] \frac{\left(\frac{tx}{x+1}\right)^{m-c_n}}{(x+1)^{\theta+\alpha}} f_{B_{\theta+\alpha,1-\alpha}}(t) f_{Z_{\alpha,\theta+\alpha}}(x) dt dx \\
 &= \frac{\frac{\Gamma(\theta+1)\Gamma(m-l+\theta+\alpha)\Gamma(l+1-\alpha)}{\Gamma(\theta+\alpha)\Gamma(1-\alpha)\Gamma(m+\theta+1)}}{D(m, c_n; \alpha, \theta, J)} \int_0^{+\infty} \binom{c}{l} (x+1)^{l-c_n} \frac{\left(\frac{x}{x+1}\right)^{m-c_n}}{(x+1)^{\theta+\alpha}} f_{Z_{\alpha,\theta+\alpha}}(x) dx \\
 &= \frac{\frac{\Gamma(\theta+1)\Gamma(m-l+\theta+\alpha)\Gamma(l+1-\alpha)}{\Gamma(\theta+\alpha)\Gamma(1-\alpha)\Gamma(m+\theta+1)}}{D(m, c_n; \alpha, \theta, J)} \binom{c_n}{l} \int_0^{+\infty} x^{m-c_n} (x+1)^{l-m-\theta-\alpha} f_{Z_{\alpha,\theta+\alpha}}(x) dx, \tag{45}
 \end{aligned}$$

where

$$D(m, c_n; \alpha, \theta, J) = \int_0^{+\infty} \left[\frac{x^{m-c_n}}{(x+1)^{\theta+m+1}} \right] f_{W_{\alpha,\theta}}(x) dx.$$

It is easy to show that (45) is mixture of Binomial distributions. In particular, from (45) we write

$$\begin{aligned} & \Pr[f_{X_{m+1}} = l \mid C_{n, h_n}(X_{m+1}) = c_n] \\ &= \frac{1}{D(m, c_n; \alpha, \theta, J)} \\ & \times \int_0^1 \int_0^{+\infty} \left[\binom{c_n}{l} \left(\frac{1-t}{\frac{x+1-xt}{x+1}} \right)^l \left(\frac{\frac{t}{x+1}}{\frac{x+1-xt}{x+1}} \right)^{c_n-l} \right] \frac{\left(\frac{x+1-xt}{x+1} \right)^{c_n} \left(\frac{tx}{x+1} \right)^{m-c_n}}{(x+1)^{\theta+\alpha}} f_{B_{\theta+\alpha, 1-\alpha}}(t) f_{Z_{\alpha, \theta+\alpha}}(x) dt dx, \end{aligned} \quad (46)$$

Now, $F_{X_{m+1}}$ be a random variable with distribution (46) and compute the moment of order r of $F_{X_{m+1}}$. From the representation of the distribution of $F_{X_{m+1}}$ as a mixture of Binomial distribution,

$$\begin{aligned} & \mathbb{E}[(F_{X_{m+1}})^r] \\ &= \mathbb{E}[f_{X_{m+1}}^r \mid C_{n, h_n}(X_{m+1}) = c_n] \\ &= \frac{1}{D(m, c_n; \alpha, \theta, J)} \\ & \times \int_0^1 \int_0^{+\infty} \left[\sum_{l=0}^{c_n} \binom{c_n}{l} l^r (1-t)^l \left(\frac{t}{x+1} \right)^{c_n-l} \right] \frac{\left(\frac{tx}{x+1} \right)^{m-c_n}}{(x+1)^{\theta+\alpha}} f_{B_{\theta+\alpha, 1-\alpha}}(t) f_{Z_{\alpha, \theta+\alpha}}(x) dt dx. \end{aligned}$$

Now, in the previous expression, we consider the summation within brackets. Recall that the (u, v) -th Stirling number of the second type, here denoted by $S(u, v)$, is defined as the v -th coefficient in the expansion of t^u into falling factorials, i.e. $t^u = \sum_{0 \leq v \leq u} S(u, v) (t)_{[v]}$; moreover, it is assumed: $S(0, 0) = 1$, $S(u, 0) = 0$ for $u > 0$ and $S(u, v) = 0$ for $v > u$. Then, we write

$$\begin{aligned} & \sum_{l=0}^{c_n} \binom{c_n}{l} l^r (1-t)^l \left(\frac{t}{x+1} \right)^{c_n-l} \\ &= \sum_{l=0}^{c_n} \left(\sum_{k=0}^r S(r, k) k! \binom{l}{k} \right) \binom{c_n}{l} (1-t)^l \left(\frac{t}{x+1} \right)^{c_n-l} \\ &= \sum_{k=0}^r S(r, k) \frac{c_n!}{(c_n-k)!} (1-t)^k \sum_{l=k}^{c_n} \binom{c_n-k}{l-k} (1-t)^{l-k} \left(\frac{t}{x+1} \right)^{c_n-l} \\ &= \sum_{k=0}^r S(r, k) \frac{c_n!}{(c_n-k)!} (1-t)^k \sum_{j=0}^{c_n-k} \binom{c_n-k}{j} (1-t)^j \left(\frac{t}{x+1} \right)^{c_n-k-j} \\ &= \sum_{k=0}^r S(r, k) \frac{c_n!}{(c_n-k)!} (1-t)^k \left(1-t + \frac{t}{x+1} \right)^{c_n-k} \\ &= c_n^r (1-t)^r \left(1-t + \frac{t}{x+1} \right)^{c_n-r} + O(c_n^{r-1}), \end{aligned}$$

where $O(c_n^{r-1})$ in the last identity is intended as $c_n \rightarrow +\infty$. Accordingly, we can write the following

$$\begin{aligned}
 & \mathbb{E} \left[\left(\frac{F_{X_{m+1}}}{c_n} \right)^r \right] \\
 &= \mathbb{E} \left[\left(\frac{f_{X_{m+1}}}{c_n} \right)^r \mid C_{n, h_n(X_{m+1})} = c_n \right] \\
 &= \frac{1}{D(m, c_n; \alpha, \theta, J)} \\
 & \quad \times \int_0^1 \int_0^{+\infty} (1-t)^r \left(1-t + \frac{t}{x+1} \right)^{c_n-r} \frac{\left(\frac{tx}{x+1} \right)^{m-c_n}}{(x+1)^{\theta+\alpha}} f_{B_{\theta+\alpha, 1-\alpha}}(t) f_{Z_{\alpha, \theta+\alpha}}(x) dt dx + O\left(\frac{1}{c_n} \right).
 \end{aligned} \tag{47}$$

Now, the double integral on the right-hand side of (47) can be rewritten by means of the following change of variable: $y = (1+x)/((1+x)(1-t)+t) \in (1, \frac{1}{1-t})$. In particular, we can write

$$\begin{aligned}
 & \int_0^1 \int_0^{+\infty} (1-t)^r \left(1-t + \frac{t}{x+1} \right)^{c_n-r} \frac{\left(\frac{tx}{x+1} \right)^{m-c_n}}{(x+1)^{\theta+\alpha}} f_{B_{\theta+\alpha, 1-\alpha}}(t) f_{Z_{\alpha, \theta+\alpha}}(x) dt dx \\
 &= \int_0^1 (1-t)^r \\
 & \quad \times \left[\int_1^{\frac{1}{1-t}} y^{r-c_n} \left(\frac{y-1}{y} \right)^{m-c_n} \left(\frac{1-y(1-t)}{yt} \right)^{\theta+\alpha} \frac{t f_{Z_{\alpha, \theta+\alpha}}\left(\frac{y-1}{1-y(1-t)} \right)}{[1-y(1-t)]^2} dy \right] f_{B_{\theta+\alpha, 1-\alpha}}(t) dt.
 \end{aligned} \tag{48}$$

We develop a large m asymptotic analysis of (48), as well as of $D(m, c_n; \alpha, \theta, J)$, under the large m asymptotic regime $c_n = \lambda m$. We start from the term $D(m, c_n; \alpha, \theta, J)$, which we rewrite as

$$D(m, c_n; \alpha, \theta, J) = \int_0^{+\infty} \left[\frac{x^\beta}{x+1} \right]^m \varphi(x) dx \tag{49}$$

where $\beta := 1 - \lambda$ and $\varphi(x) := f_{W_{\alpha, \theta}}(x)/(1+x)^{\theta+1}$. The function $\psi : x \mapsto x^\beta/(x+1)$ has a unique maximum point $\bar{x} := \frac{\beta}{1-\beta} = \frac{1-\lambda}{\lambda}$. Moreover, straightforward computations show that

$$\psi'(x) = \frac{\beta x^{\beta-1} - (1-\beta)x^\beta}{(1+x)^2}$$

and

$$\psi''(x) = \frac{-\beta(1-\beta)x^{\beta-2} - 2\beta(2-\beta)x^{\beta-1} + (1-\beta)(2-\beta)x^\beta}{(1+x)^3}.$$

Then, $\psi''(\bar{x}) = -\bar{x}^{\beta-1}/(1+\bar{x})^3$ and the Laplace method leads to the following large m behaviour

$$D(m, c_n; \alpha, \theta, J) \sim \frac{1}{\sqrt{m}} \varphi(\bar{x}) \left[\frac{\bar{x}^\beta}{1+\bar{x}} \right]^{m+\frac{1}{2}} \sqrt{\frac{2\pi(1+\bar{x})^3}{\bar{x}^{\beta-1}}}. \tag{50}$$

We consider (48), i.e. the integral within brackets on the right-hand side of (48), which we rewrite as

$$\begin{aligned} & \int_1^{\frac{1}{1-t}} y^r \left[\frac{(y-1)^\beta}{y} \right]^m \left(\frac{1-y(1-t)}{yt} \right)^{\theta+\alpha} \frac{t f_{Z_{\alpha, \theta+\alpha}} \left(\frac{y-1}{1-y(1-t)} \right)}{[1-y(1-t)]^2} dy \\ &= \int_0^{\frac{t}{1-t}} (1+x)^r \left[\frac{x^\beta}{1+x} \right]^m \varphi_t(x) dx, \end{aligned}$$

where

$$\varphi_t(x) := \left(\frac{1-(x+1)(1-t)}{(x+1)t} \right)^{\theta+\alpha} \frac{t f_{Z_{\alpha, \theta+\alpha}} \left(\frac{x}{1-(x+1)(1-t)} \right)}{[1-(x+1)(1-t)]^2}.$$

To apply the Laplace method, we note that $(0, 1) \ni t \mapsto t/(1-t) \in (0, +\infty)$ is a strictly monotonically increasing function. Thus, $\beta < t$ entails $\bar{x} := \frac{\beta}{1-\beta} < \frac{t}{1-t}$ and, for such t , it holds

$$\int_0^{\frac{t}{1-t}} (1+x)^r \left[\frac{x^\beta}{1+x} \right]^m \varphi_t(x) dx \sim \frac{(1+\bar{x})^r}{\sqrt{m}} \varphi_t(\bar{x}) \left[\frac{\bar{x}^\beta}{1+\bar{x}} \right]^{m+\frac{1}{2}} \sqrt{\frac{2\pi(1+\bar{x})^3}{\bar{x}^{\beta-1}}} \quad (51)$$

for large m . On the other hand, $\beta > t$ entails $\bar{x} := \frac{\beta}{1-\beta} > \frac{t}{1-t}$ and, for such t , there holds a similar large m asymptotic expansion. Now, by exploiting the fact that $\psi : x \mapsto x^\beta/(x+1)$ is a strictly monotonically increasing function for $x \in (0, t/(1-t))$, then we can write the following

$$\int_0^{\frac{t}{1-t}} (1+x)^r \left[\frac{x^\beta}{1+x} \right]^m \varphi_t(x) dx \sim \frac{1}{m^{\theta+2\alpha}} \left[\frac{\left(\frac{t}{1-t} \right)^\beta}{1 + \frac{t}{1-t}} \right]^m \rho(t) = \frac{1}{m^{\theta+2\alpha}} [t^\beta(1-t)^{1-\beta}]^m \rho(t)$$

for large m , where ρ is a suitable function independent of m . Accordingly, we can write that

$$\frac{1}{m^{\theta+2\alpha}} \int_0^\beta (1-t)^r [t^\beta(1-t)^{1-\beta}]^m \rho(t) f_{B_{\theta+\alpha, 1-\alpha}}(t) dt \sim \frac{C(\beta)}{m^{1+\theta+2\alpha}} \left[\frac{\bar{x}^\beta}{1+\bar{x}} \right]^m \quad (52)$$

as $m \rightarrow +\infty$. Then, starting from Equation (47) and then gathering (51) and (52) we can write

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{F_{X_{m+1}}}{c_n} \right)^r \right] \\ & \sim \frac{1}{D(m, c_n; \alpha, \theta, J)} \left\{ \frac{1}{m^{\theta+2\alpha}} \int_0^\beta (1-t)^r [t^\beta(1-t)^{1-\beta}]^m \rho(t) f_T(t) dt \right. \\ & \quad \left. + \frac{(1+\bar{x})^r}{\sqrt{m}} \sqrt{\frac{2\pi(1+\bar{x})^3}{\bar{x}^{\beta-1}}} \left[\frac{\bar{x}^\beta}{1+\bar{x}} \right]^{m+\frac{1}{2}} \int_\beta^1 (1-t)^r \varphi_t(\bar{x}) f_T(t) dt \right\}. \end{aligned} \quad (53)$$

as $m \rightarrow +\infty$. According to (50) the first term in the right-hand side of (53) is negligible, and hence

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{F_{X_{m+1}}}{c_n} \right)^r \right] \\ & \sim \frac{(1 + \bar{x})^r}{\varphi(\bar{x})} \int_{\beta}^1 (1-t)^r \varphi_t(\bar{x}) f_{B_{\theta+\alpha, 1-\alpha}}(t) dt \\ & = \left(\frac{m}{c_n} \right)^r \frac{1}{\varphi(\bar{x})} \int_0^\lambda \tau^r \varphi_{1-\tau}(\bar{x}) f_{B_{\theta+\alpha, 1-\alpha}}(1-\tau) d\tau \end{aligned}$$

as $m \rightarrow +\infty$, which, because of the large m asymptotic regime $c_n = \lambda m$, completes the proof.

Appendix H. Proof of Equation (31)

Because of the independence assumption of \mathcal{H} , i.e. the hash functions h_n 's are i.i.d. according to the strong universal family \mathcal{H} , and by an application of Bayes theorem, we can write

$$\begin{aligned} & \Pr[f_{X_{m+1}} = l \mid \{C_{n, h_n(X_{m+1})}\}_{n \in [N]} = \{c_n\}_{n \in [N]}] \\ & = \frac{1}{\Pr[\{C_{n, h_n(X_{m+1})}\}_{n \in [N]} = \{c_n\}_{n \in [N]}]} \Pr[f_{X_{m+1}} = l] \prod_{n=1}^N \Pr[C_{n, h_n(X_{m+1})} = c_n \mid f_{X_{m+1}} = l] \\ & = \frac{1}{\Pr[\{C_{n, h_n(X_{m+1})}\}_{n \in [N]} = \{c_n\}_{n \in [N]}]} \Pr[f_{X_{m+1}} = l] \prod_{n=1}^N \frac{\Pr[C_{n, h_n(X_{m+1})} = c_n, f_{X_{m+1}} = l]}{\Pr[f_{X_{m+1}} = l]} \\ & = \frac{1}{\Pr[\{C_{n, h_n(X_{m+1})}\}_{n \in [N]} = \{c_n\}_{n \in [N]}]} (\Pr[f_{X_{m+1}} = l])^{1-N} \\ & \quad \times \prod_{n=1}^N \Pr[C_{n, h_n(X_{m+1})} = c_n] \Pr[f_{X_{m+1}} = l \mid C_{n, h_n(X_{m+1})} = c_n] \\ & = (\Pr[f_{X_{m+1}} = l])^{1-N} \prod_{n=1}^N \Pr[f_{X_{m+1}} = l \mid C_{n, h_n(X_{m+1})} = c_n] \end{aligned} \tag{54}$$

for $l = 0, 1, \dots, \min\{c_1, \dots, c_N\}$, where $\Pr[f_{X_{m+1}} = l \mid C_{n, h_n(X_{m+1})} = c_n]$ is precisely the posterior distribution computed in Theorem 3 with respect to the hash function h_n , whereas

$$\begin{aligned} \Pr[f_{X_{m+1}} = l] & = \sum_{\mathbf{m} \in \mathcal{M}_{m, k}} \Pr[\mathbf{M}_m = \mathbf{m}] \Pr[X_{m+1} \in \mathbf{v}_l \mid X_{1:m}] \\ & = \sum_{\mathbf{m} \in \mathcal{M}_{m, k}} \Pr[\mathbf{M}_m = \mathbf{m}] \Pr[X_{m+1} \in \mathbf{v}_l \mid X_{1:m}], \end{aligned}$$

where $\Pr[\mathbf{M}_m = \mathbf{m}]$ is in Equation (18) and $\Pr[X_{m+1} \in \mathbf{v}_l \mid X_{1:m}]$ is in Equation (20). That is,

$$\Pr[f_{X_{m+1}} = l] = \sum_{\mathbf{m} \in \mathcal{M}_{m, k}} m! \frac{\binom{\theta}{\alpha}^{(k)}}{(\theta)^{(m)}} \prod_{i=1}^m \left(\frac{\alpha(1-\alpha)^{(i-1)}}{i!} \right)^{m_i} \frac{1}{m_i!} \begin{cases} \frac{\theta+k\alpha}{\theta+m} & \text{if } l = 0 \\ \frac{m_l(l-\alpha)}{\theta+m} & \text{if } l \geq 1. \end{cases}$$

For $l = 0$

$$\begin{aligned}
 \Pr[f_{X_{m+1}} = 0] &= \sum_{\mathbf{m} \in \mathcal{M}_{m,k}} m! \frac{\left(\frac{\theta}{\alpha}\right)_{(k)}}{(\theta)_{(m)}} \prod_{i=1}^m \left(\frac{\alpha(1-\alpha)_{(i-1)}}{i!} \right)^{m_i} \frac{1}{m_i!} \frac{\theta + k\alpha}{\theta + m} \\
 &= \frac{\theta}{\theta + m} + \frac{\alpha}{\theta + m} \mathbb{E}[K_m] \\
 &= \frac{\theta}{\theta + m} + \frac{\alpha}{\theta + m} \left(\frac{(\theta + \alpha)_{(m)}}{\alpha(\theta + 1)_{(m-1)}} - \frac{\theta}{\alpha} \right),
 \end{aligned}$$

where the last equality follows from Pitman (2006, Equation 3.13). Accordingly, we can write that

$$\Pr[f_{X_{m+1}} = 0] = \frac{(\theta + \alpha)_{(m)}}{(\theta + 1)_{(m)}}$$

For $l \geq 1$

$$\begin{aligned}
 \Pr[f_{X_{m+1}} = l] &= \sum_{\mathbf{m} \in \mathcal{M}_{m,k}} m! \frac{\left(\frac{\theta}{\alpha}\right)_{(k)}}{(\theta)_{(m)}} \prod_{i=1}^m \left(\frac{\alpha(1-\alpha)_{(i-1)}}{i!} \right)^{m_i} \frac{1}{m_i!} \frac{m_l(l-\alpha)}{\theta + m} \\
 &= \frac{l-\alpha}{\theta + m} \mathbb{E}[M_{l,m}] \\
 &= \frac{l-\alpha}{\theta + m} \frac{(1-\alpha)_{(l-1)}}{l!} (m)_{[l]} \frac{(\theta + \alpha)_{(m-l)}}{(\theta + 1)_{(m-1)}}
 \end{aligned}$$

where the last equality follows from Favaro et al. (2013, Proposition 1). Accordingly, for $l = 1, \dots, m$,

$$\Pr[f_{X_{m+1}} = l] = \frac{(1-\alpha)_{(l)}}{l!} (m)_{[l]} \frac{(\theta + \alpha)_{(m-l)}}{(\theta + 1)_{(m)}}. \quad (55)$$

Equation (31) follows by combining the distribution (54) with (55). This completes the proof.

Appendix I. CMS for range queries under DP priors

We assume that the stream $x_{1:m}$ is modeled as a random sample $X_{1:m}$ from an unknown discrete distribution P , which is endowed with a DP prior, i.e. $P \sim \text{DP}(\theta; \nu)$. Let h_1, \dots, h_N be a collection of random hash functions that are i.i.d. from the strong universal family \mathcal{H} , and assume that h_1, \dots, h_N are independent of $X_{1:m}$ for any $m \geq 1$; in particular, by de Finetti's representation theorem, h_1, \dots, h_N are independent of $P \sim \text{DP}(\theta; \nu)$. Under this BNP framework, a s -range query induces the posterior distribution of the frequencies $(f_{x_{m+1}}, \dots, f_{x_{m+s}})$ given $\{(C_{n,h_n(v_1)}, \dots, C_{n,h_n(v_s)})\}_{n \in [N]}$, for arbitrary $\{x_{m+1}, \dots, x_{m+s}\} \in \mathcal{V}$. This posterior distribution, in turn, induces the posterior distribution of the s -range query \bar{f}_s given $\{(C_{n,h_n(v_1)}, \dots, C_{n,h_n(v_s)})\}_{n \in [N]}$. CMS-DP estimates of \bar{f}_s are obtained as functionals of the posterior distribution of \bar{f}_s given $\{(C_{n,h_n(v_1)}, \dots, C_{n,h_n(v_s)})\}_{n \in [N]}$. To compute the posterior distribution of $(f_{x_{m+1}}, \dots, f_{x_{m+s}})$ given $\{(C_{n,h_n(v_1)}, \dots, C_{n,h_n(v_s)})\}_{n \in [N]}$, it is natural to consider s additional random samples $(X_{m+1}, \dots, X_{m+s})$. In particular, for

any $r = 1, \dots, s$ let $f_{X_{m+r}}$ be the frequency of X_{m+r} in $X_{1:m}$, i.e.,

$$f_{X_{m+r}} = \sum_{i=1}^m \mathbb{1}_{\{X_i\}}(X_{m+r})$$

and let $C_{n, h_n(X_{m+r})}$ be the hashed frequency of all X_i 's, for $i = 1, \dots, m$, such that $h_n(X_i) = h_n(X_{m+r})$, i.e.,

$$C_{n, h_n(X_{m+r})} = \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(h(X_{m+r})).$$

Now, let $\mathbf{X}_s = (X_{m+1}, \dots, X_{m+s})$ and for $n \in [N]$ let $\mathbf{f}_{\mathbf{X}_s} = (f_{X_{m+1}}, \dots, f_{X_{m+s}})$. For $n \in [N]$ let $\mathbf{C}_{n, h_n(\mathbf{X}_s)} = (C_{n, h_n(X_{m+1})}, \dots, C_{n, h_n(X_{m+s})})$. For each h_n we are interested in the posterior distribution

$$\Pr[\mathbf{f}_{\mathbf{X}_s} = \mathbf{1}_s \mid \mathbf{C}_{n, h_n(\mathbf{X}_s)} = \mathbf{c}_n] = \frac{\Pr[\mathbf{f}_{\mathbf{X}_s} = \mathbf{1}_s, \mathbf{C}_{n, h_n(\mathbf{X}_s)} = \mathbf{c}_n]}{\Pr[\mathbf{C}_{n, h_n(\mathbf{X}_s)} = \mathbf{c}_n]}. \quad (56)$$

for $\mathbf{1}_s \in \{0, 1, \dots, m\}^s$. For the collection of hash functions h_1, \dots, h_N , the posterior distribution of $\mathbf{f}_{\mathbf{X}_s}$ given $\{\mathbf{C}_{n, h_n(\mathbf{X}_s)}\}_{n \in [N]}$ follows from the posterior distribution (56) by the assumption that the h_n 's are i.i.d. according to the strong universal family \mathcal{H} , and Bayes theorem.

Hereafter we show that the ‘‘Bayesian’’ proof of Section 2 can be readily extended to the computation of the posterior distribution (56). We outline this extension for any range $s \geq 1$, and then we present an explicit example for $s = 2$. To simplify the notation, we remove the subscript n from h_n and \mathbf{c}_n . Then, we are interested in computing the posterior distribution

$$\Pr[\mathbf{f}_{\mathbf{X}_s} = \mathbf{1}_s \mid \mathbf{C}_{h(\mathbf{X}_s)} = \mathbf{c}] = \frac{\Pr[\mathbf{f}_{\mathbf{X}_s} = \mathbf{1}_s, \mathbf{C}_{h(\mathbf{X}_s)} = \mathbf{c}]}{\Pr[\mathbf{C}_{h(\mathbf{X}_s)} = \mathbf{c}]}. \quad (57)$$

For $s = 1$ the posterior distribution (57) reduces to (10). The independence between h_n and $X_{1:m}$ allows us to invoke the ‘‘freezing lemma’’ (Baldi, 2017, Lemma 4.1), according to which we can treat h_n as it was fixed, i.e. non-random. We analyze the posterior distribution (57) starting from its denominator. In particular, the denominator of (57) can be written as follows

$$\begin{aligned} \Pr[\mathbf{C}_{h(\mathbf{X}_s)} = \mathbf{c}] &= \sum_{(j_1, \dots, j_s) \in [J]^s} \Pr[\mathbf{C}_{h(\mathbf{X}_s)} = \mathbf{c}, h(X_{m+1}) = j_1, \dots, h(X_{m+s}) = j_s] \\ &= \sum_{(j_1, \dots, j_s) \in [J]^s} \Pr \left[\sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_1) = c_1, \dots, \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_s) = c_s, \right. \\ &\quad \left. h(X_{m+1}) = j_1, \dots, h(X_{m+s}) = j_s \right]. \end{aligned}$$

To evaluate

$$\Pr \left[\sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_1) = c_1, \dots, \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_s) = c_s, h(X_{m+1}) = j_1, \dots, h(X_{m+s}) = j_s \right], \quad (58)$$

we split the sum over $[J]^s$ and we organize the summands as follows. First, we introduce a variable k which counts how many distinct object there are in each vector (j_1, \dots, j_s) , so that $k \in \{1, 2, \dots, \min\{s, J\}\}$. Second, we consider the vector (r_1, \dots, r_k) of frequencies of the distinct k objects. Third, we consider the vector (j_1^*, \dots, j_k^*) of distinct objects with $\{j_1^*, \dots, j_k^*\} \subseteq \{1, \dots, J\}$. Then, we evaluate the probability (58) in the distinguishing case that

$$\begin{cases} j_1 = \dots = j_{r_1} =: j_1^* \\ j_{r_1+1} = \dots = j_{r_1+r_2} =: j_2^* \\ \dots \\ j_{r_1+\dots+r_{k-1}+1} = \dots = j_{r_1+\dots+r_k} =: j_k^* \end{cases}$$

such that the probability (58) of interest is different from zero if and only if the following holds true

$$\begin{cases} c_1 = \dots = c_{r_1} =: c_1^* \\ c_{r_1+1} = \dots = c_{r_1+r_2} =: c_2^* \\ \dots \\ c_{r_1+\dots+r_{k-1}+1} = \dots = c_{r_1+\dots+r_k} =: c_k^*. \end{cases}$$

That is,

$$\Pr \left[\sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_1^*) = c_1^*, \dots, \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_k^*) = c_k^*, h(X_{m+1}) = \dots = h(X_{m+r_1}) = j_1^*, \dots \right. \\ \left. \dots, h(X_{m+r_1+\dots+r_{k-1}+1}) = \dots = h(X_{m+r_1+\dots+r_k}) = j_k^* \right].$$

Now, we set $B_r^* := \{x \in \mathcal{V} : h(x) = j_r^*\}$ for any $r \in \{1, \dots, k\}$ and we set $B_{k+1}^* = (\cup_{r=1}^k B_r^*)^C$. Thus, $\{B_1^*, \dots, B_{k+1}^*\}$ is a finite partition of \mathcal{V} . If $k = J$, then $B_{k+1}^* = \emptyset$ and in such case we intend that $\{B_1^*, \dots, B_{k+1}^*\}$ is replaced by $\{B_1^*, \dots, B_k^*\}$. Accordingly, we can write the identity

$$\Pr \left[\sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_1^*) = c_1^*, \dots, \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_k^*) = c_k^*, h(X_{m+1}) = \dots = h(X_{m+r_1}) = j_1^*, \dots \right. \\ \left. \dots, h(X_{m+r_1+\dots+r_{k-1}+1}) = \dots = h(X_{m+r_1+\dots+r_k}) = j_k^* \right] \\ = \binom{m}{c_1^*, \dots, c_k^*} \int_{\Delta_k} \left(\prod_{i=1}^k p_i^{c_i^* + r_i} \right) (1 - p_1 - \dots - p_k)^{m - \sum_{i=1}^k c_i^*} \mu_{B_1^*, \dots, B_{k+1}^*} (dp_1 \dots dp_k)$$

where $\mu_{B_1^*, \dots, B_{k+1}^*}$ is the distribution of $(P(B_1^*), \dots, P(B_k^*))$ which, by the finite-dimensional projective property of the DP, is a Dirichlet distribution with parameter $(\theta/J, \dots, \theta/J)$ on Δ_k . If $k < J$

$$\Pr \left[\sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_1^*) = c_1^*, \dots, \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_k^*) = c_k^*, h(X_{m+1}) = \dots = h(X_{m+r_1}) = j_1^*, \dots \right. \\ \left. \dots, h(X_{m+r_1+\dots+r_{k-1}+1}) = \dots = h(X_{m+r_1+\dots+r_k}) = j_k^* \right]$$

$$= \frac{\Gamma(\theta)}{[\Gamma(\frac{\theta}{J})]^k \Gamma((J-k)\frac{\theta}{J})} \frac{\left[\prod_{i=1}^k \Gamma(\frac{\theta}{J} + c_i^* + r_i) \right] \Gamma((J-k)\frac{\theta}{J} + m - \sum_{i=1}^k c_i^*)}{\Gamma(\theta + m + s)},$$

and if $k = J$

$$\Pr \left[\sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_1^*) = c_1^*, \dots, \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_k^*) = c_k^*, h(X_{m+1}) = \dots = h(X_{m+r_1}) = j_1^*, \dots \right. \\ \left. \dots, h(X_{m+r_1+\dots+r_{k-1}+1}) = \dots = h(X_{m+r_1+\dots+r_k}) = j_k^* \right] \\ = \frac{\Gamma(\theta)}{(\Gamma(\frac{\theta}{J}))^k} \frac{\prod_{i=1}^k \Gamma(\frac{\theta}{J} + c_i^* + r_i)}{\Gamma(\theta + m + s)}.$$

Upon denoting by $I_k(c_{n,1}^*, \dots, c_{n,k}^*; r_1, \dots, r_k)$ the right expression of the integral, we conclude that

$$\Pr[\mathbf{C}_{h(\mathbf{X}_s)} = \mathbf{c}] \tag{59} \\ = \sum_{k=1}^{\min\{s, J\}} \frac{J!}{(J-k)!} \\ \times \sum_{(\pi_1, \dots, \pi_k) \in \Pi(s, k)} \Delta(\pi_1, \dots, \pi_k; c_1, \dots, c_s) \binom{m}{c_1^*, \dots, c_k^*} I_k(c_1^*, \dots, c_k^*; |\pi_1|, \dots, |\pi_k|),$$

where: i) $\Pi(s, k)$ denotes the set of all possible partitions of the set $\{1, \dots, s\}$ into k disjoint subsets π_1, \dots, π_k ; $|\pi_i|$ stands for the cardinality of the subset π_i ; ii) $\Delta(\pi_1, \dots, \pi_k; c_1, \dots, c_s)$ is either 0 or 1 with the proviso that it equals 1 if and only if, for all $z \in \{1, \dots, k\}$ for which $|\pi_z| \geq 2$, all the integers c_i with $i \in \pi_z$ are equal; for any $i \in \{1, \dots, k\}$, c_i represents the common integer associated to π_i . Formula (59) simplifies remarkably for small values of s . For instance,

i) for $s = 1$

$$\Pr[C_h(X_{m+1}) = c_1] = J \binom{m}{c_1} I_1(c_1; 1);$$

ii) for $s = 2$

$$\Pr[C_h(X_{m+1}) = c_1, C_h(X_{m+2}) = c_2] \tag{60} \\ = J \mathbb{1}\{c_1 = c_2\} \binom{m}{c_1} I_1(c_1; 2) + J(J-1) \binom{m}{c_1, c_2} I_2(c_1, c_2; 1, 1).$$

We conclude by studying the numerator in (57). This expression is determined by the complete knowledge of the joint distribution of (X_1, \dots, X_{m+s}) . As above, we can start by writing

$$\Pr[\mathbf{f}_{\mathbf{X}_s} = \mathbf{l}_s, \mathbf{C}_{h(\mathbf{X}_s)} = \mathbf{c}]$$

$$\begin{aligned}
 &= \sum_{k=1}^s \sum_{(\pi_1, \dots, \pi_k) \in \Pi(s, k)} \Delta(\pi_1, \dots, \pi_k; l_1, \dots, l_s) \binom{n}{l_1^*, \dots, l_k^*} \\
 &\quad \times \Pr \left[B(m; l_1^*, \dots, l_k^*; \pi_1, \dots, \pi_k) \cap \left\{ \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_1) = c_1, \dots, \sum_{i=1}^m \mathbb{1}_{h(X_i)}(j_s) = c_s \right\} \right]
 \end{aligned}$$

where the event $B(m; l_1^*, \dots, l_k^*)$ is characterized by the relations among random variables X_{m+r} 's

$$\begin{aligned}
 &X_1 = \dots = X_{l_1^*} = X_{m+r} \quad \text{for all } r \in \pi_1 \\
 &X_{l_1^*+1} = \dots = X_{l_1^*+l_2^*} = X_{m+r} \quad \text{for all } r \in \pi_2 \\
 &\quad \dots \dots \dots \\
 &X_{l_1^*+\dots+l_{k-1}^*+1} = \dots = X_{l_1^*+\dots+l_k^*} = X_{m+r} \quad \text{for all } r \in \pi_k \\
 &\quad X_{m+r_1} \neq X_{m+r_2} \quad \text{for all } r_1 \in \pi_a, r_2 \in \pi_b \quad \text{for all } a \neq b \\
 &\{X_{l_1^*+\dots+l_k^*+1}, \dots, X_m\} \cap \{X_{m+1}, \dots, X_{m+s}\} = \emptyset.
 \end{aligned}$$

The numerator of (57) can be treated as the denominator of (57), namely by exploiting the double partition structure induced by the above relations on the random variables X_i 's and $h(X_i)$'s. We observe that the combination of this two partition structures proves particularly cumbersome to be written for general $s \geq 1$. For this reason, further manipulations of the posterior distribution (57) will be deferred to the proof the next theorem, where we assume $s = 2$.

Theorem 6 *For $m \geq 1$, let $x_{1:m}$ be a stream of tokens that are modeled as a random sample $X_{1:m}$ from $P \sim DP(\theta; \nu)$, and let (X_{m+1}, X_{m+2}) be a pair of additional random samples from P . Moreover, let h_n be a random hash function distributed as the strong universal family \mathcal{H} , and let h_n be independent of $X_{1:m}$ for any $m \geq 1$, that is h_n is independent of P . Then*

$$\begin{aligned}
 &\Pr[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2 \mid C_{n, h_n(X_{m+1})} = c_{n,1}, C_{n, h_n(X_{m+2})} = c_{n,2}] \\
 &= \frac{\text{Num}(l_1, l_2, c_{n,1}, c_{n,2})}{\text{Den}(c_{n,1}, c_{n,2})} \quad l_1, l_2 \geq 0
 \end{aligned}$$

with

i)

$$\begin{aligned}
 \text{Den}(c_{n,1}, c_{n,2}) &= J \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \frac{\binom{\theta}{J}_{(c+2)} \left(\theta - \frac{\theta}{J}\right)_{(m-c)}}{c!(m-c)!} \\
 &\quad + J(J-1) \frac{\binom{\theta}{J}_{(c_{n,1}+1)} \binom{\theta}{J}_{(c_{n,2}+1)} \left(\theta - \frac{2\theta}{J}\right)_{(m-c_{n,1}-c_{n,2})}}{c_{n,1}! c_{n,2}! (m-c_{n,1}-c_{n,2})!};
 \end{aligned}$$

ii)

$$\text{Num}(l_1, l_2, c_{n,1}, c_{n,2}) = \mathbb{1}\{l_1 = l_2 =: l, c_{n,1} = c_{n,2} = c\} \frac{\theta(l+1) \binom{\theta}{J}_{(c-l)} \left(\theta - \frac{\theta}{J}\right)_{(m-c)}}{(c-l)!(m-c)!}$$

$$\begin{aligned}
 & + \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \frac{\theta^2 (\frac{\theta}{J})_{(c-l_1-l_2)} (\theta - \frac{\theta}{J})_{(m-c)}}{J(c-l_1-l_2)! (m-c)!} \\
 & + \left(\frac{J-1}{J} \right) \frac{\theta^2 (\frac{\theta}{J})_{(c_{n,1}-l_1)} (\frac{\theta}{J})_{(c_{n,2}-l_2)} (\theta - \frac{2\theta}{J})_{(m-c_{n,1}-c_{n,2})}}{(c_{n,1}-l_1)! (c_{n,2}-l_2)! (m-c_{n,1}-c_{n,2})!}.
 \end{aligned}$$

Proof Following the ‘‘Bayesian’’ proof for $s \geq 1$, we start by expressing the posterior distribution of $(f_{X_{m+1}}, f_{X_{m+2}})$ given $C_{n,h_n(X_{m+1})}$ and $C_{n,h_n(X_{m+2})}$ as a ratio of two probabilities, and then we deal with the numerator and denominator. That is, we write the following expression

$$\begin{aligned}
 & \Pr[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2 \mid C_{n,h_n(X_{m+1})} = c_{n,1}, C_{n,h_n(X_{m+2})} = c_{n,2}] \tag{61} \\
 & = \frac{\Pr[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2}]}{\Pr[C_{n,h_n(X_{m+1})} = c_{n,1}, C_{n,h_n(X_{m+2})} = c_{n,2}]}
 \end{aligned}$$

Observe that the denominator of the posterior distribution (61) reduces to (60). Then, by using the finite-dimensional projective property of the DP, we can write the following expressions

$$\begin{aligned}
 & J \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \binom{m}{c} I_1(c; 2) \\
 & = J \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \binom{m}{c} \\
 & \quad \times \int_0^1 p^{c+2} (1-p)^{m-c} \frac{\Gamma(\theta)}{\Gamma(\theta/J) \Gamma(\theta(1-1/J))} p^{\theta/J-1} (1-p)^{\theta(1-1/J)-1} dp \\
 & = J \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \binom{m}{c} \frac{\Gamma(\theta)}{\Gamma(\theta/J) \Gamma(\theta(1-1/J))} \frac{\Gamma(\theta/J + c + 2) \Gamma(\theta(1-1/J) + m - c)}{\Gamma(\theta + m + 2)}
 \end{aligned}$$

and

$$\begin{aligned}
 & J(J-1) \binom{m}{c_{n,1}, c_{n,2}} I_2(c_{n,2}, c_{n,2}; 1, 1) \\
 & = J(J-1) \binom{m}{c_{n,1}, c_{n,2}} \\
 & \quad \times \int_{\Delta_2} p_1^{c_{n,1}+1} p_2^{c_{n,2}+1} (1-p_1-p_2)^{m-c_{n,1}-c_{n,2}} \\
 & \quad \times \frac{\Gamma(\theta)}{\Gamma(\theta/J) \Gamma(\theta/J) \Gamma(\theta(1-2/J))} p_1^{\theta/J-1} p_2^{\theta/J-1} (1-p_1-p_2)^{\theta(1-2/J)-1} dp_1 dp_2 \\
 & = J(J-1) \binom{m}{c_{n,1}, c_{n,2}} \frac{\Gamma(\theta)}{(\Gamma(\theta/J))^2 \Gamma(\theta(1-2/J))} \\
 & \quad \times \frac{\Gamma(\theta/J + c_{n,1} + 1) \Gamma(\theta/J + c_{n,2} + 1) \Gamma(\theta(1-2/J) + m - c_{n,1} - c_{n,2})}{\Gamma(\theta + m + 2)}.
 \end{aligned}$$

Then,

$$\Pr[C_{n,h_n(X_{m+1})} = c_{n,1}, C_{n,h_n(X_{m+2})} = c_{n,2}] \tag{62}$$

$$\begin{aligned}
 &= J \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \binom{m}{c_{n,1}} \frac{\Gamma(\theta)}{\Gamma(\theta/J)\Gamma(\theta(1-1/J))} \\
 &\quad \times \frac{\Gamma(\theta/J + c + 1)\Gamma(\theta(1-1/J) + m - c)}{\Gamma(\theta + m + 2)} \\
 &+ J(J-1) \binom{m}{c_{n,1}, c_{n,2}} \frac{\Gamma(\theta)}{(\Gamma(\theta/J))^2\Gamma(\theta(1-2/J))} \\
 &\quad \times \frac{\Gamma(\theta/J + c_{n,1} + 1)\Gamma(\theta/J + c_{n,2} + 1)\Gamma(\theta(1-2/J) + m - c_{n,1} - c_{n,2})}{\Gamma(\theta + m + 2)}.
 \end{aligned}$$

Now, we focus on the numerator of the posterior distribution (61), which is rewritten as follows

$$\begin{aligned}
 &\Pr \left[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2} \right] \quad (63) \\
 &= \Pr \left[f_{X_{n+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2}, X_{m+1} = X_{m+2} \right] \\
 &\quad + \Pr \left[f_{X_{n+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2}, X_{m+1} \neq X_{m+2} \right].
 \end{aligned}$$

First, we consider the first term on the right-hand side of the probability (63). In particular, we write

$$\begin{aligned}
 &\Pr \left[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2}, X_{m+1} = X_{m+2} \right] \\
 &= \mathbb{1}\{l_1 = l_2 =: l, c_{n,1} = c_{n,2} = c\} \binom{m}{l} \\
 &\quad \times \Pr \left[X_1 = \dots, X_l = X_{m+1} = X_{m+2}, \{X_{l+1}, \dots, X_m\} \cap \{X_{m+1}\} = \emptyset, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c \right] \\
 &= \mathbb{1}\{l_1 = l_2 =: l, c_{n,1} = c_{n,2} = c\} \binom{m}{l} \\
 &\quad \times \Pr \left[X_1 = \dots, X_l = X_{m+1} = X_{m+2}, \{X_{l+1}, \dots, X_m\} \cap \{X_{m+1}\} = \emptyset, \sum_{i=l+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c - l \right]
 \end{aligned}$$

which is determined by the distribution of (X_1, \dots, X_{m+2}) . In view of Sangalli (2006, Equation 3.5)

$$\Pr[X_1 \in C_1, \dots, X_{m+2} \in C_{m+2}] = \sum_{k=1}^{m+2} \frac{\theta^k}{(\theta)^{(m+2)}} \sum_{(\pi_1, \dots, \pi_k) \in \Pi(m+2, k)} \prod_{i=1}^k (|\pi_i| - 1)! \nu(\cap_{r \in \pi_i} C_r).$$

We set $D(m, l) := \{X_1 = \dots, X_l = X_{m+1} = X_{m+2}, \{X_{l+1}, \dots, X_m\} \cap \{X_{m+1}\} = \emptyset\}$, and we define $\mu_{(\pi_1, \dots, \pi_k)}$ as the probability measure on $(\mathcal{V}^{m+2}, \mathcal{Y}^{m+2})$ generated by the following

identity

$$\nu_{\pi_1, \dots, \pi_k}(C_1 \times \dots \times C_{m+2}) := \prod_{i=1}^k \nu(\cap_{r \in \pi_i} C_r),$$

It is clear that such measures attach to $D(m, l)$ a probability value that is either 0 or 1. In particular, $\nu_{\pi_1, \dots, \pi_k}(D(m, l)) = 1$ if and only if one of the π 's (e.g. π_k , being these partitions given up to the order) is exactly equal to the set $\{1, \dots, l, m+1, m+2\}$. Accordingly, we write

$$\begin{aligned} & \Pr \left[D(m, l), \sum_{i=l+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c - l \right] \\ &= \sum_{k=2}^{m-l+1} \frac{\theta^k}{(\theta)_{(m+2)}} \sum_{(\pi_1, \dots, \pi_{k-1}) \in \Pi(m-l, k-1)} (l+1)! \prod_{i=1}^{k-1} (|\pi_i| - 1)! \nu_{\pi_1, \dots, \pi_k} \left(\sum_{i=l+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c - l \right) \\ &= \frac{\theta(\theta)_{(m-l)}}{(\theta)_{(m+2)}} (l+1)! \sum_{r=1}^{m-l} \frac{\theta^r}{(\theta)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (|\pi_i| - 1)! \times \\ & \quad \times \left\{ \sum_{j=1}^J \nu(\{j\}) \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=l+1}^m \mathbb{1}_{h_n(X_i)}(j) = c - l \right) \right\} \\ &= \frac{\theta(\theta)_{(m-l)}}{J(\theta)_{(m+2)}} (l+1)! \sum_{r=1}^{m-l} \frac{\theta^r}{(\theta)_{(m-l)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l, r)} \prod_{i=1}^r (|\pi_i| - 1)! \times \\ & \quad \times \left\{ \sum_{j=1}^J \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=l+1}^m \mathbb{1}_{h_n(X_i)}(j) = c - l \right) \right\}. \end{aligned}$$

Hence,

$$\begin{aligned} & \Pr \left[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2}, X_{m+1} = X_{m+2} \right] \\ & \quad \quad \quad (64) \\ &= \mathbb{1}\{l_1 = l_2 =: l, c_{n,1} = c_{n,2} = c\} \frac{m!}{(c-l)!(m-c)!} \frac{\theta(l+1)}{\Gamma(\theta+m+2)} \times \\ & \quad \times \frac{\Gamma(\theta)}{\Gamma(\theta/J)\Gamma(\theta(1-1/J))} \Gamma(\theta/J+c-l) \Gamma(\theta(1-1/J)+m-c). \end{aligned}$$

Now, we consider the second term on the right-hand side of the probability (63). In particular, we write

$$\begin{aligned} & \Pr \left[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2}, X_{m+1} \neq X_{m+2} \right] \\ &= \binom{m}{l_1, l_2} \Pr \left[X_1 = \dots, X_{l_1} = X_{m+1}, X_{l_1+1} = \dots, X_{l_1+l_2} = X_{m+2}, X_{m+1} \neq X_{m+2}, \right] \end{aligned}$$

$$\begin{aligned}
 & \left. \{X_{l_1+l_2+1}, \dots, X_m\} \cap \{X_{m+1}, X_{m+2}\} = \emptyset, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2} \right] \\
 = & \binom{m}{l_1, l_2} \Pr \left[X_1 = \dots, X_{l_1} = X_{m+1}, X_{l_1+1} = \dots, X_{l_1+l_2} = X_{m+2}, X_{m+1} \neq X_{m+2}, \right. \\
 & \left. \{X_{l_1+l_2+1}, \dots, X_m\} \cap \{X_{m+1}, X_{m+2}\} = \emptyset, \right. \\
 & \left. l_2 \mathbb{1}_{h_n(X_{l_1+1})}(X_{m+1}) + \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1} - l_1, \right. \\
 & \left. l_1 \mathbb{1}_{h_n(X_1)}(X_{m+2}) + \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2} - l_2 \right].
 \end{aligned}$$

Setting

$$E(m, l_1, l_2) := \left\{ X_1 = \dots, X_{l_1} = X_{m+1}, X_{l_1+1} = \dots, X_{l_1+l_2} = X_{m+2}, X_{m+1} \neq X_{m+2}, \right. \\
 \left. \{X_{l_1+l_2+1}, \dots, X_m\} \cap \{X_{m+1}, X_{m+2}\} = \emptyset \right\},$$

we have that $\nu_{\pi_1, \dots, \pi_k}(E(m, l_1, l_2)) = 1$ if and only if two of the π 's (e.g. π_{k-1} and π_k , being these partitions given up to the order) are exactly equal to the sets $\{1, \dots, l_1, m+1\}$ and $\{l_1+1, \dots, l_1+l_2, m+2\}$, respectively. Therefore, from above, we write the following probability

$$\begin{aligned}
 & \Pr \left[E(m, l_1, l_2), l_2 \mathbb{1}_{h_n(X_{l_1+1})}(X_{m+1}) + \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1} - l_1, \right. \\
 & \left. l_1 \mathbb{1}_{h_n(X_1)}(X_{m+2}) + \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2} - l_2 \right] \\
 = & \sum_{k=3}^{n-l_1-l_2+2} \frac{\theta^k}{(\theta)_{(m+2)}} \sum_{(\pi_1, \dots, \pi_{k-2}) \in \Pi(m-l_1-l_2, k-2)} l_1! l_2! \prod_{i=1}^{k-2} (|\pi_i| - 1)! \times \\
 & \times \nu_{\pi_1, \dots, \pi_k} \left(l_2 \mathbb{1}_{h_n(X_{l_1+1})}(X_{m+1}) + \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1} - l_1, \right. \\
 & \left. l_1 \mathbb{1}_{h_n(X_1)}(X_{m+2}) + \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2} - l_2 \right) \\
 = & \frac{\theta^2 (\theta)_{(m-l_1-l_2)}}{(\theta)_{(m+2)}} l_1! l_2! \sum_{r=1}^{m-l_1-l_2} \frac{\theta^r}{(\theta)_{(m-l_1-l_2)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l_1-l_2, r)} \prod_{i=1}^r (|\pi_i| - 1)! \\
 & \times \left[\sum_{(j_1, j_2) \in [J]^2} \nu(\{j_1\}) \nu(\{j_2\}) \right]
 \end{aligned}$$

$$\times \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(j_1) = c_{n,1} - l_1 - l_2 \mathbb{1}\{j_1 = j_2\}, \right. \\ \left. \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(j_2) = c_{n,2} - l_2 - l_1 \mathbb{1}\{j_1 = j_2\} \right).$$

We observe that the expression within the brackets in the last term, as a sum over $[J]^2$, can be split into the sum of two terms, according on whether $j_1 = j_2$ or not. Therefore, we write

$$\sum_{r=1}^{m-l_1-l_2} \frac{\theta^r}{(\theta)_{(m-l_1-l_2)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l_1-l_2, r)} \prod_{i=1}^r (|\pi_i| - 1)! \\ \times \left[\sum_{j_1=j_2 \in [J]} \nu(\{j_1\}) \nu(\{j_2\}) \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(j_1) = c_{n,1} - l_1 - l_2, \right. \right. \\ \left. \left. \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(j_2) = c_{n,2} - l_2 - l_1 \right) \right] \\ = \frac{1}{J} \mathbb{1}\{c_{n,1} = c_{n,2} =: c\} \binom{m-l_1-l_2}{c-l_1-l_2} \frac{\Gamma(\theta)}{\Gamma(\theta/J) \Gamma(\theta(1-1/J))} \\ \times \frac{\Gamma(\theta/J + c - l_1 - l_2) \Gamma(\theta(1-1/J) + m - c)}{\Gamma(\theta + m - l_1 - l_2)}.$$

On the other hand, assuming $J \geq 3$

$$\sum_{r=1}^{m-l_1-l_2} \frac{\theta^r}{(\theta)_{(m-l_1-l_2)}} \sum_{(\pi_1, \dots, \pi_r) \in \Pi(m-l_1-l_2, r)} \prod_{i=1}^r (|\pi_i| - 1)! \\ \times \left[\sum_{\substack{(j_1, j_2) \in [J]^2 \\ j_1 \neq j_2}} \nu(\{j_1\}) \nu(\{j_2\}) \nu_{\pi_1, \dots, \pi_r} \left(\sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(j_1) = c_{n,1} - l_1, \right. \right. \\ \left. \left. \sum_{i=l_1+l_2+1}^m \mathbb{1}_{h_n(X_i)}(j_2) = c_{n,2} - l_2 \right) \right] \\ = \frac{J-1}{J} \binom{m-l_1-l_2}{c_{n,1}-l_1, c_{n,2}-l_2} \frac{\Gamma(\theta)}{[\Gamma(\theta/J)]^2 \Gamma(\theta(1-2/J))} \\ \times \frac{\Gamma(\theta/J + c_{n,1} - l_1) \Gamma(\theta/J + c_{n,2} - l_2) \Gamma(\theta(1-2/J) + m - c_{n,1} - c_{n,2})}{\Gamma(\theta + m - l_1 - l_2)}.$$

Then,

$$\Pr \left[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2}, X_{m+1} \neq X_{m+2} \right] \quad (65)$$

$$\begin{aligned}
 &= \binom{m}{l_1, l_2} \frac{\theta^2(\theta)_{(m-l_1-l_2)}}{(\theta)_{(m+2)}} l_1! l_2! \\
 &\quad \times \left[\frac{1}{J} \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \binom{m-l_1-l_2}{c-l_1-l_2} \frac{\Gamma(\theta)}{\Gamma(\theta/J)\Gamma(\theta(1-1/J))} \frac{\Gamma(\theta/J+c-l_1-l_2)\Gamma(\theta(1-1/J)+m-c)}{\Gamma(\theta+m-l_1-l_2)} \right. \\
 &\quad + \frac{J-1}{J} \binom{m-l_1-l_2}{c_{n,1}-l_1, c_{n,2}-l_2} \frac{\Gamma(\theta)}{[\Gamma(\theta/J)]^2\Gamma(\theta(1-2/J))} \\
 &\quad \left. \times \frac{\Gamma(\theta/J+c_{n,1}-l_1)\Gamma(\theta/J+c_{n,2}-l_2)\Gamma(\theta(1-2/J)+m-c_{n,1}-c_{n,2})}{\Gamma(\theta+m-l_1-l_2)} \right].
 \end{aligned}$$

Then, by combining the probability (64) and the probability (65) we write the following expression

$$\begin{aligned}
 &\Pr \left[f_{X_{m+1}} = l_1, f_{X_{m+2}} = l_2, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+1}) = c_{n,1}, \sum_{i=1}^m \mathbb{1}_{h_n(X_i)}(X_{m+2}) = c_{n,2} \right] \\
 &= \frac{m!}{\Gamma(\theta+m+2)} \left\{ \mathbb{1}\{l_1 = l_2 =: l, c_{n,1} = c_{n,2} = c\} \frac{\theta(l+1)}{(c-l)!(m-c)!} \beta_1(\theta, J) \times \right. \\
 &\quad \times \Gamma(\theta/J+c-l)\Gamma(\theta(1-1/J)+m-c) \\
 &\quad + \mathbb{1}\{c_{n,1} = c_{n,2} = c\} \frac{\theta^2}{J(c-l_1-l_2)!(m-c)!} \beta_1(\theta, J) \Gamma(\theta/J+c-l_1-l_2)\Gamma(\theta(1-1/J)+m-c) \\
 &\quad + \left(\frac{J-1}{J} \right) \frac{\theta^2}{(c_{n,1}-l_1)!(c_{n,2}-l_2)!(m-c_{n,1}-c_{n,2})!} \beta_2(\theta, J) \times \\
 &\quad \left. \times \Gamma(\theta/J+c_{n,1}-l_1)\Gamma(\theta/J+c_{n,2}-l_2)\Gamma(\theta(1-2/J)+m-c_{n,1}-c_{n,2}) \right\} \quad (66)
 \end{aligned}$$

with

$$\beta_1(\theta, J) := \frac{\Gamma(\theta)}{\Gamma(\theta/J)\Gamma(\theta(1-1/J))}$$

and

$$\beta_2(\theta, J) := \frac{\Gamma(\theta)}{[\Gamma(\theta/J)]^2\Gamma(\theta(1-2/J))}.$$

The proof is completed by combing the posterior distribution (61) with probabilities (62) and (66). \blacksquare

Theorem 6 extends Theorem 1 to the more general problem of estimating 2-range queries. In particular, for the collection of hash functions h_1, \dots, h_N , the posterior distribution of $(f_{X_{m+1}}, f_{X_{m+2}})$ given $\{(C_{n,h_n(X_{m+1})}, C_{n,h_n(X_{m+2})})\}_{n \in [N]}$ follows from Theorem 6 by the assumption that the h_n 's are i.i.d. according to the strong universal family \mathcal{H} , and Bayes theorem. CMS-DP estimates of the 2-range query $\bar{f}_2 = f_{x_{m+1}} + f_{x_{m+2}}$ are then obtained as functionals of the posterior distribution of \bar{f}_2 , e.g. posterior mode, posterior mean and posterior median. To conclude, it remains to estimate the prior's parameter $\theta > 0$ based on hashed frequencies; this is obtained following the empirical Bayes procedure described in Section 2.

Acknowledgement

The authors are grateful to the Action Editor (Professor Erik Sudderth) and four anonymous Referees for their comments and corrections that allow to improve remarkably the paper. Stefano Favaro wishes to thank Graham Cormode, Matteo Sesia and Luca Trevisan for stimulating discussions on sketches and generalizations thereof. Emanuele Dolera and Stefano Favaro received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 817257. Emanuele Dolera and Stefano Favaro are thankful for the financial support of Italian Ministry of Education, University and Research (MIUR), “Dipartimenti di Eccellenza” grant 2018-2022. Stefano Favaro is also affiliated to IMATI-CNR “Enrico Magenes” (Milan, Italy).

References

- AAMAND, A., INDYK, P. AND VAKILIAN, A. (2019). Frequency estimation algorithms under Zipfian distribution. *Preprint arXiv:1908.05198*.
- AGGARWAL, C. AND YU, P. (2010). On classification of high-cardinality data streams. In *Proceedings of the 2010 SIAM International Conference on Data Mining*.
- BACALLADO, S., BATTISTON, M., FAVARO, S. AND TRIPPA, L. (2017). Sufficientness postulates for Gibbs-type priors and hierarchical generalizations. *Statistical Science* **32**, 487–500.
- BALDI, P. (2017) *Stochastic calculus*. Springer.
- BARABÁSI, A.L. (2005) The origin of bursts and heavy tails in human dynamics. *Nature* **435**, 227.
- BERNTON, E., JACOB, P.E., GERBER, M. AND ROBERT, C.P. (2019). On parameter estimation with the Wasserstein distance. *Information and Inference* **8**, 657–676.
- CAI, D., MITZENMACHER, M. AND ADAMS, R.P. (2018). A Bayesian nonparametric view on count–min sketch. In *Advances in Neural Information Processing Systems*.
- CANCHO, R.F. AND SOLÉ, R.V. (2003). Least effort and the origins of scaling in human language. *Proceeding of the National Academy of Sciences of USA* **100**, 788–791.
- CHARALAMBIDES, C. (2005) *Combinatorial methods in discrete distributions*. Wiley.
- CLAUSET, A., SHALIZI, C.R. AND NEWMAN, M.E.J. (2009). Power-law distributions in empirical data. *SIAM Review* **51**, 661–703.
- CORMODE, G., GAROFALAKIS, M. AND HAAS, P.J. (2012). *Synopses for massive data: samples, histograms, wavelets, sketches*. Foundations and Trends in Databases.
- CORMODE, G. AND MUTHUKRISHNAN, S. (2005). An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms* **55**, 58–75.

- CORMODE, G. AND YI, K. (2020). *Small summaries for big data*. Cambridge University Press.
- DE BLASI, P., FAVARO, S., LIJOI, A., MENA, R.H., PRÜNSTER, I. AND RUGGIERO, M. (2015). Are Gibbs-type priors the most natural generalization of the Dirichlet process? *IEEE Transactions on Pattern Analysis and Machine Intelligence* **37**, 212–229.
- DEVROYE, L. (2009). Random variate generation for exponentially and polynomially tilted stable distributions. *ACM Transactions on Modeling and Computer Simulation* **19**, 4.
- DOLERA, E. (2013). Estimates of the approximation of weighted sums of conditionally independent random variables by the normal law. *J. Inequal. Appl.* **2013**, 320.
- DOLERA, E. AND FAVARO, S. (2020). A Berry–Esseen theorem for Pitman’s α -diversity. *The Annals of Applied Probability* **30**, 847–869.
- DOLERA, E. AND FAVARO, S. (2020). Rates of convergence in de Finetti’s representation theorem, and Hausdorff moment problem. *Bernoulli* **26**, 1294–1322.
- DOLERA, E., FAVARO, S. AND PELUCHETTI, S. (2021). A Bayesian nonparametric approach to count-min sketch under power-law data stream. In *International Conference on Artificial Intelligence and Statistics*.
- DWORK, C. AND NAOR, M. AND PITASSI, T. AND ROTHBLUM, G. AND YEKHANIN, S. (2010). Pan-private streaming algorithms. In *Proceedings of the Symposium on Innovations in Computer Science*.
- FAVARO, S., LIJOI, A., R.H., MENA AND PRÜNSTER, I. (2009). Bayesian nonparametric inference for species variety with a two parameter Poisson-Dirichlet process prior. *Journal of the Royal Statistical Society Series B* **71**, 992–1008.
- FAVARO, S., LIJOI, A. AND PRÜNSTER, I. (2013). Conditional formulae for Gibbs-type exchangeable random partitions. *The Annals of Applied Probability* **23**, 1721–1754.
- FAVARO, S. AND NIPOTI, B. AND TEH, Y.W. (2015). Random variate generation for Laguerre-type exponentially tilted alpha-stable distributions. *Electronic Journal of Statistics* **9**, 1230–1242.
- FERGUSON, T.S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics* **1**, 209–230.
- GHOSAL, S. AND VAN DER VAART, A. (2017) *Fundamentals of Nonparametric Bayesian Inference*. Cambridge University Press.
- GNEDIN, A., HANSEN, B. AND PITMAN, J. (2007). Notes on the occupancy problems with infinitely many boxes: general asymptotics and power law. *Probability Surveys* **4**, 146–171.
- Gnedin, A. and Pitman J. (2006). Exchangeable Gibbs partitions and Stirling triangles. *Journal of Mathematical Sciences*, **138**, 5674–5685.

- GOYAL, A., DAUMÉ, H. AND CORMODE, G. (2012). Sketch algorithms for estimating point queries in NLP. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*.
- GOYAL, A., DAUMÉ, H. AND VENKATASUBRAMANIAN, S. (2009). Streaming for large scale NLP: language modeling. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*.
- HARALD, B.R. (2001). *Word Frequency Distributions*. Springer
- HARRISON, B.A. (2010) *Move prediction in the game of Go*. Ph.D Thesis, Harvard University.
- HSU, C., INDYK, P., KATABI, D. AND VAKILIAN, A. (2019) Learning-based frequency estimation algorithms. In *Proceedings of the International Conference on Learning Representations*.
- HUBERMAN, B.A. AND ADAMIC, L.A. (1999) Internet: growth dynamics of the World-Wide Web. *Nature* **401**, 131
- JAMES, L.F. (2002). Poisson process partition calculus with applications to exchangeable models and Bayesian nonparametrics. *Preprint arXiv:math/0205093*.
- JAMES, L.F., PRÜNSTER, I., LIJOI, A. (2009). Posterior analysis for normalized random measures with independent increments. *Scandinavian Journal of Statistics* **36**, 76–97.
- JOHNSON, N.L., KEMP, A.W. AND KOTZ, S. (2005) *Univariate discrete distributions*, Wiley Series in Probability and Statistics.
- KINGMAN, J.F.C. (1993). *Poisson processes*. Wiley Online Library.
- LEO ELWORTH, L.A., WANG, Q., KOTA, P.K., BARBERAN, C.J., COLEMAN, B., BALAJI, A., GUPTA, G., BARANIUK, R.G., SHRIVASTAVA, A. AND TREANGEN, T.J. (2020). To petabytes and beyond: recent advances in probabilistic and signal processing algorithms and their application to metagenomics. *Nucleic Acids Research* **48** 5217–5234.
- LETHAM, B., KARRER, B., OTTONI, G. AND BAKSHY, E. (2019). Constrained Bayesian optimization with noisy experiments. *Bayesian Analysis* **14**, 495–519.
- LIJOI, A., MENA, R. H., AND PRÜNSTER, I. (2005). Hierarchical mixture modeling with normalized inverse-Gaussian priors. *Journal of the American Statistical Association* **100**, 1278–1291.
- LIJOI, A., PRÜNSTER, I. AND WALKER, S.G. (2008). Bayesian nonparametric estimators derived from conditional Gibbs structures. *The Annals of Applied Probability* **18**, 1519–1547.
- MONECHI, B., RUIZ-SERRANO, A., TRIA, F., AND LORETO, V. (2017). Waves of novelties in the expansion into the adjacent possible. *PloS ONE* **12**

- MUCHNIK, L., PEI, S., PARRA, L.C., REIS, S.D.S, ANDRADE, J.S., HAVLIN, S. AND MAKSE, H.A. (2013). Origins of power-law degree distribution in the heterogeneity of human activity in social networks. *Nature Scientific Reports* **3**, 1783
- PERMAN, M., PITMAN, J. AND YOR, M. (1992). Size-biased sampling of Poisson point processes and excursions. *Probability Theory and Related Fields* **92**, 21–39.
- PITEL, G. AND FOUQUIER, G. (2015). Count-min-log sketch: approximately counting with approximate counters. In *Proceedings of the International Symposium on Web Algorithm*.
- PITMAN, J. (1995). Exchangeable and partially exchangeable random partitions. *Probability Theory and Related Fields* **102**, 145–158.
- PITMAN, J. (2003). Poisson-Kingman partitions. In *Science and Statistics: A Festschrift for Terry Speed*, Goldstein, D.R. Eds. Institute of Mathematical Statistics.
- PITMAN, J. (2006). *Combinatorial stochastic processes*. Lecture Notes in Mathematics, Springer Verlag.
- PITMAN, J. AND YOR, M. (1997). The two parameter Poisson-Dirichlet distribution derived from a stable subordinator. *The Annals of Probability* **25**, 855–900.
- PRÜNSTER, I. (2002). *Random probability measures derived from increasing additive processes and their application to Bayesian statistics*. Ph.d thesis, University of Pavia.
- REGAZZINI, E. (1978). *Intorno ad alcune questioni relative alla definizione del premio secondo la teoria della credibilità*. *Giornale dell’Istituto Italiano degli Attuari* **41**, 77–89.
- REGAZZINI, E. (2001). *Foundations of Bayesian statistics and some theory of Bayesian nonparametric methods*. Lecture Notes, Stanford University.
- 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.
- RYBSKI, D., BULDYREV, S.V., HAVLIN, S., LILJEROS, F. AND MAKSE, H A. (2016). Scaling laws of human interaction activity. *Proceeding of the National Academy of Sciences of USA* **106**, 12640.
- SANGALLI, M.L. (2006). Some developments of the normalized random measures with independent increments. *Sankhya A* **68**, 461–487.
- SETHURAMAN, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* **4**, 639–650.
- SONG, H.H., CHO, T.W., DAVE, V., ZHANG, Y. AND QIU, L. (2009). Scalable proximity estimation and link prediction in online social networks. In *Proceedings of the ACM SIGCOMM Conference on Internet measurement*.

- TING, D. (2018). Count-min: optimal estimation and tight error bounds using empirical error distributions. In *International Conference on Knowledge Discovery and Data Mining*.
- TRIA, F., LORETO, V., SERVEDIO, V.D.P AND STROGATZ, S.H. (2014). The dynamics of correlated novelties. *Nature Scientific Reports* **4**, 5890.
- ZABELL, S.L. (1997). The continuum of inductive methods revisited. In *The cosmos of science: essays in exploration*, Earman, J. and Norton, J.D. Eds. Universty of Pittsburgh Press.
- ZHANG, Q., PELL, J., CANINO-KONING, R., HOWE, A.C. AND BROWN, C.T. (2014). These are not the k-mers you are looking for: efficient online k -mer counting using a probabilistic data structure. *PloS one* **9**.
- ZIPF, G.K. (1949). *Human behaviour and the principle of least effort: an introduction to human ecology*. Addison-Wesley.
- ZOLOTAREV, V.M. (1986). *One dimensional stable distributions*. American Mathematical Society.