

Bernstein-von Mises theorem for the Pitman-Yor process of nonnegative type

Franssen, S. E.M.P.; van der Vaart, A. W.

DOI

[10.1214/22-EJS2077](https://doi.org/10.1214/22-EJS2077)

Publication date

2022

Document Version

Final published version

Published in

Electronic Journal of Statistics

Citation (APA)

Franssen, S. E. M. P., & van der Vaart, A. W. (2022). Bernstein-von Mises theorem for the Pitman-Yor process of nonnegative type. *Electronic Journal of Statistics*, 16(2), 5779-5811. <https://doi.org/10.1214/22-EJS2077>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Bernstein-von Mises theorem for the Pitman-Yor process of nonnegative type*

S. E. M. P. Franssen

TU Delft

e-mail: s.e.m.p.franssen@tudelft.nl

A. W. van der Vaart

TU Delft

e-mail: a.w.vandervaart@tudelft.nl

Abstract: The Pitman-Yor process is a random probability distribution, that can be used as a prior distribution in a nonparametric Bayesian analysis. The process is of species sampling type and generates discrete distributions, which yield of the order n^σ different values (“species”) in a random sample of size n , if the type σ is positive. Thus this type parameter can be set to target true distributions of various levels of discreteness, making the Pitman-Yor process an interesting prior in this case. It was previously shown that the resulting posterior distribution is consistent if and only if the true distribution of the data is discrete. In this paper we derive the distributional limit of the posterior distribution, in the form of a (corrected) Bernstein-von Mises theorem, which previously was known only in the continuous, inconsistent case. It turns out that the Pitman-Yor posterior distribution has good behaviour if the true distribution of the data is discrete with atoms that decrease not too slowly. Credible sets derived from the posterior distribution provide valid frequentist confidence sets in this case. For a general discrete distribution, the posterior distribution, although consistent, may contain a bias which does not converge to zero at the \sqrt{n} rate and invalidates posterior inference. We propose a bias correction that solves this problem. We also consider the effect of estimating the type parameter from the data, both by empirical Bayes and full Bayes methods. In a small simulation study we illustrate that without bias correction the coverage of credible sets can be arbitrarily low, also for some discrete distributions.

MSC2020 subject classifications: Primary 62G20; secondary 62G15.

Keywords and phrases: Pitman-Yor process, Bernstein-von Mises theorem, weak convergence, credible set, empirical Bayes, species sampling.

Received December 2021.

1. Introduction

The Pitman-Yor process [24, 20] is a random probability distribution, which can be used as a prior distribution in a nonparametric Bayesian analysis. It is

*The research leading to these results is partly financed by a Spinoza prize awarded by the Netherlands Organisation for Scientific Research (NWO).

characterised by a *type* parameter σ , which in this paper we take to be positive. The Pitman-Yor process of type $\sigma = 0$ is the Dirichlet process [12], which is well understood, while negative types correspond to finitely discrete distributions and were considered in [7]. The Pitman-Yor process is also known as the two-parameter Poisson-Dirichlet Process, is an example of a Poisson-Kingman process [23], and a species sampling process of Gibbs type [8].

The easiest definition is through *stick-breaking* ([20, 15]), as follows. The family of nonnegative Pitman-Yor processes is given by three parameters: a number $\sigma \in [0, 1)$, a number $M > -\sigma$ and an atomless probability distribution G on some measurable space $(\mathcal{X}, \mathcal{A})$. We say that a random probability measure P on $(\mathcal{X}, \mathcal{A})$ is a Pitman-Yor process (of nonnegative type), denoted $P \sim \text{PY}(\sigma, M, G)$, if P can be represented as

$$P = \sum_{i=1}^{\infty} W_i \delta_{\theta_i},$$

where $W_i = V_i \prod_{j=1}^{i-1} (1 - V_j)$ for $V_i \stackrel{\text{iid}}{\sim} \text{B}(1 - \sigma, M + i\sigma)$, independent of $\theta_i \stackrel{\text{iid}}{\sim} G$, and B the beta distribution.

It is clear from this definition that the realisations of P are discrete probability measures, with countably many atoms at random locations, with random weights. If one first draws $P \sim \text{PY}(\sigma, M, G)$, and next given P a random sample X_1, \dots, X_n from P , then ties among the latter observations are possible, or even likely. It is known ([23]) that the number K_n of different values among X_1, \dots, X_n is almost surely of the order n^σ if $\sigma > 0$, whereas it is logarithmic in n if $\sigma = 0$. This suggests that the Pitman-Yor process is a reasonable prior distribution for a dataset in which similar patterns are expected (or observed). In particular, when a large number of clusters is expected, a Pitman-Yor process of positive type could be preferred over the standard Dirichlet prior, which corresponds to $\sigma = 0$. Applications in genetics or topic modelling can be found in [30, 26, 14, 1]. The Pitman-Yor process has also been proposed as a prior for estimating the probability that a next observation is a new species [11], with applications in e.g. forensic statistics [5, 6]. The papers [3, 4] study hierarchical versions of Pitman-Yor processes, which are useful to discover structure in data beyond clustering.

In this paper we consider the properties of the Pitman-Yor posterior distribution to estimate the distribution of a random sample of observations. By definition this posterior distribution is the conditional distribution of P given X_1, \dots, X_n in the Bayesian hierarchical model $P \sim \text{PY}(\sigma, M, G)$ and $X_1, \dots, X_n | P \stackrel{\text{iid}}{\sim} P$. We assume that in reality the observations X_1, \dots, X_n are an i.i.d. (i.e. independent and identically distributed) sample from a distribution P_0 and investigate the use of the posterior distribution for inference on P_0 . It was shown in [16, 8] that in this setting, as $n \rightarrow \infty$,

$$P | X_1, \dots, X_n \rightsquigarrow \delta_{(1-\lambda)P_0^d + \lambda(1-\sigma)P_0^c + \sigma\lambda G}, \quad (1)$$

where \rightsquigarrow denotes weak convergence of measures, δ_Q denotes the Dirac measure at the probability distribution Q , and $P_0 = (1-\lambda)P_0^d + \lambda P_0^c$ is the decomposition

of P_0 in its discrete component $(1 - \lambda)P_0^d$ and the remaining (atomless) part λP_0^c . In the case that P_0 is discrete, we have $\lambda = 0$ and the measure $(1 - \lambda)P_0^d + \lambda(1 - \sigma)P_0^c + \sigma\lambda G$ in the right side reduces to $P_0^d = P_0$, and hence (1) expresses that the posterior distribution collapses asymptotically to the Dirac measure at P_0 . The posterior distribution is said to be consistent in this case. However, if P_0 is not discrete, then the posterior distribution recovers P_0 asymptotically only if $\sigma = 0$ (the case of the Dirichlet prior) or if $G = P_0^c$. The last case will typically fail and hence in the case that $\sigma > 0$ the posterior distribution will typically be consistent if and only if P_0 is discrete. This reveals the Pitman-Yor prior of positive type as a reasonable prior only for discrete distributions.

Besides for recovery, a posterior distribution is used to express remaining uncertainty, for instance in the form of a credible (or Bayesian confidence) set. To justify such a procedure from a non-Bayesian point of view, the posterior consistency must be refined to a distributional result of Bernstein-von Mises type. Such a result was obtained by [16] in the case that the true distribution P_0 is atomless, the case that the posterior distribution is inconsistent and the Pitman-Yor prior is better avoided. In the present paper we study the case of general distributions P_0 , including the case of most interest that P_0 is discrete. It turns out that discreteness per se is not enough for valid inference, but it is also needed that the weights of the atoms in P_0 decrease fast enough. In the other case, ordinary Bayesian credible sets are not valid confidence sets. For the latter case our result suggests a bias correction.

Since the type parameter σ determines the number of distinct values in a sample from the prior, it might be interpreted as influencing the discreteness of the prior, smaller σ favouring fewer distinct values and hence a more discrete prior. In the asymptotic result the type parameter plays only a secondary role. At first thought counterintuitively, a larger σ , which gives a less discrete prior, increases the bias in the posterior distribution that arises when the atoms in P_0 decrease too slowly.

In practice one may prefer to estimate the type parameter from the data. The empirical Bayes method maximizes the marginal likelihood of X_1, \dots, X_n in the Bayesian setup over σ . We show that in the consistent case, substitution of this estimator in the posterior distribution for given type parameter does not change the asymptotics of the Pitman-Yor posterior. Alternatively, we may equip σ itself with a prior, resulting in a mixture of Pitman-Yor processes as a prior for P . We show that this too results in the same posterior behaviour. Thus estimating the type parameter does not solve the inconsistency problem.

We can conclude that the Pitman-Yor process is an appropriate prior for estimating a distribution only if the sizes of the atoms of this distribution decrease sufficiently rapidly. Our results show that the speed of decay depends on the aspect of interest, for instance different for the distribution function than for the mean.

Our results depend heavily on the characterisation of the posterior distribution given in [22] (see Section 4).

2. Main result

The nonparametric maximum likelihood estimator of the distribution P of a sample of observations X_1, \dots, X_n is the empirical measure $\mathbb{P}_n = n^{-1} \sum_{i=1}^n \delta_{X_i}$, the discrete uniform measure on the observations. Therefore in analogy with the case of classical parametric models (e.g. Theorem 10.1 and page 144 in [27]), in this setting a Bernstein-von Mises theorem would give the approximation of the posterior distribution of $\sqrt{n}(P - \mathbb{P}_n)$ given X_1, \dots, X_n by the normal distribution obtained as the limit of $\sqrt{n}(\mathbb{P}_n - P_0)$. To give a precise meaning to such a distributional statement, we may evaluate all the measures involved on a collection of sets, and interpret $\sqrt{n}(P - \mathbb{P}_n)$ and $\sqrt{n}(\mathbb{P}_n - P_0)$ as stochastic processes indexed by sets. For instance, in the case that the sample space is the real line, we could use the sets $(-\infty, t]$, for $t \in \mathbb{R}$, corresponding to the distribution functions of the measures P , \mathbb{P}_n and P_0 .

More generally, we may evaluate these measures on measurable functions $f : \mathcal{X} \rightarrow \mathbb{R}$, as

$$Pf = \int f dP, \quad \mathbb{P}_n f = \int f d\mathbb{P}_n = \frac{1}{n} \sum_{i=1}^n f(X_i), \quad P_0 f = \int f dP_0.$$

Given a collection \mathcal{F} of such functions, the Bernstein-von Mises can then address the distributions of the stochastic processes $\{\sqrt{n}(Pf - \mathbb{P}_n f) : f \in \mathcal{F}\}$ and $\{\sqrt{n}(\mathbb{P}_n f - P_0 f) : f \in \mathcal{F}\}$, the first one conditionally given the observations X_1, \dots, X_n . For instance, in the case that $\mathcal{X} = \mathbb{R}$, we might choose the collection \mathcal{F} to consist of all indicator functions $x \mapsto 1_{x \leq t}$, for t ranging over \mathbb{R} , but we can also add the identity function $f(x) = x$, yielding the means of the measures.

For a set \mathcal{F} of finitely many functions, these processes are just vectors in Euclidean space and their distributions can be evaluated as usual. Furthermore, the limit law of $\{\sqrt{n}(\mathbb{P}_n f - P_0 f) : f \in \mathcal{F}\}$ is a multivariate normal distribution, in view of the multivariate central limit theorem (provided $P_0 f^2 < \infty$, for every $f \in \mathcal{F}$). It is convenient to write the latter as the distribution of a Gaussian process $\{\mathbb{G}_{P_0} f : f \in \mathcal{F}\}$, determined by its mean and covariance function

$$\mathbb{E}\mathbb{G}_{P_0} f = 0 \quad \mathbb{E}\mathbb{G}_{P_0} f \mathbb{G}_{P_0} g = P_0(f - P_0 f)(g - P_0 g).$$

The process \mathbb{G}_{P_0} is known as a P_0 -Brownian bridge (see e.g. [25, 28, 27]).

An appropriate generalisation (and strengthening) of the central limit theorem to sets \mathcal{F} of infinitely many functions is Donsker's theorem (e.g. [27], Chapter 19). The Bernstein-von Mises theorem can be strengthened in a similar fashion. For the case of indicator functions on the real line, Donsker's theorem was derived by [10], and the corresponding Bernstein-von Mises theorem for the Dirichlet process by [18, 19]. A precise formulation (in the general case, which is not more involved than the real case) is as follows.

A class of functions \mathcal{F} is called P_0 -Donsker if the sequence $\sqrt{n}(\mathbb{P}_n - P_0)$ converges in distribution to a tight, Borel measurable element in the metric space $\ell^\infty(\mathcal{F})$ of bounded functions $z : \mathcal{F} \rightarrow \mathbb{R}$, equipped with the uniform

norm $\|z\|_{\mathcal{F}} = \sup_{f \in \mathcal{F}} |z(f)|$. The limit is then a version of the Gaussian process \mathbb{G}_{P_0} . The Bernstein-von Mises theorem involves conditional convergence in distribution given the observations X_1, \dots, X_n , which is best expressed using a metric. The bounded Lipschitz metric (see for example [28], Chapter 1.12) is convenient, and leads to defining conditional convergence in distribution of the sequence $\sqrt{n}(P - \mathbb{P}_n)$ in $\ell^\infty(\mathcal{F})$ given X_1, \dots, X_n to \mathbb{G}_{P_0} as

$$\sup_{h \in \text{BL}_1} \left| \mathbb{E} \left(h(\sqrt{n}(P - \mathbb{P}_n)) \mid X_1, \dots, X_n \right) - \mathbb{E} h(\mathbb{G}_{P_0}) \right| \rightarrow 0,$$

where the convergence refers to the i.i.d. sample X_1, X_2, \dots from P_0 , and can be in (outer) probability or almost surely. The supremum is taken over the set BL_1 of all functions $h : \ell^\infty(\mathcal{F}) \rightarrow [0, 1]$ such that $|h(z_1) - h(z_2)| \leq \|z_1 - z_2\|_{\mathcal{F}}$, for all $z_1, z_2 \in \ell^\infty(\mathcal{F})$. For simplicity of notation and easy interpretation, we write the preceding display as

$$\sqrt{n}(P - \mathbb{P}_n) \mid X_1, \dots, X_n \rightsquigarrow \mathbb{G}_{P_0}.$$

Conditional convergence in distribution of other processes is defined and denoted similarly. For finite sets \mathcal{F} , the complicated definition using the bounded Lipschitz metric reduces to ordinary weak convergence of random vectors. Also, a finite set \mathcal{F} is P_0 -Donsker if and only if $P_0 f^2 < \infty$, for every $f \in \mathcal{F}$. There are many examples of infinite Donsker classes (see e.g. [28]), with the set of indicators of cells $(-\infty, t]$ as the classical example.

We are ready to formulate the main result of the paper. Let $\tilde{X}_1, \tilde{X}_2, \dots$ be the distinct values in X_1, X_2, \dots in the order of appearance, let K_n be the number of distinct elements among X_1, \dots, X_n , and set

$$\tilde{\mathbb{P}}_n = \frac{1}{K_n} \sum_{i=1}^{K_n} \delta_{\tilde{X}_i}. \tag{2}$$

All limit results refer to a sample X_1, X_2, \dots, X_n drawn from a measure P_0 . This can always be written as $P_0 = (1 - \lambda)P_0^d + \lambda P_0^c$, where P_0^d is a discrete and P_0^c an atomless distribution and $\lambda \in [0, 1]$ is the weight of the discrete part in P_0 . The decomposition is unique unless $\lambda = 0$ or $\lambda = 1$, when P_0^c or P_0^d is arbitrary.

Theorem 1. *Let $P_0 = (1 - \lambda)P_0^d + \lambda P_0^c$ where P_0^d is a discrete and P_0^c an atomless probability distribution. The posterior distribution of P in the model $P \sim \text{PY}(\sigma, M, G)$ and $X_1, \dots, X_n \mid P \stackrel{iid}{\sim} P$ satisfies for every finite collection \mathcal{F} of functions with $(P_0 + G)f^2 < \infty$, for every $f \in \mathcal{F}$, almost surely under P_0^∞ ,*

$$\begin{aligned} & \sqrt{n} \left(P - \mathbb{P}_n - \frac{\sigma K_n}{n} (G - \tilde{\mathbb{P}}_n) \right) \mid X_1, \dots, X_n \\ & \rightsquigarrow \sqrt{1 - \lambda} \mathbb{G}_{P_0^d} + \sqrt{(1 - \sigma)\lambda} \mathbb{G}_{P_0^c} + \sqrt{\sigma(1 - \sigma)\lambda} \mathbb{G}_G \\ & \quad + \sqrt{(1 - \sigma)\sigma\lambda} \left(\frac{(1 - \lambda)P_0^d + (1 - \sigma)\lambda P_0^c}{1 - \sigma\lambda} - G \right) Z_1 \end{aligned}$$

$$+ \frac{\sqrt{(1-\sigma)\lambda(1-\lambda)}}{\sqrt{1-\sigma\lambda}} (P_0^c - P_0^d) Z_2.$$

Here $\mathbb{G}_{P_0^d}$, $\mathbb{G}_{P_0^c}$ and \mathbb{G}_G are independent Brownian bridge processes, independent of the independent standard normal variables Z_1 and Z_2 . More generally this is true, with convergence in $\ell^\infty(\mathcal{F})$ in probability, for every P_0 -Donsker class of functions \mathcal{F} for which the PY (σ, σ, G) process satisfies the central limit theorem in $\ell^\infty(\mathcal{F})$. If in addition $P_0^* \|f - P_0 f\|_{\mathcal{F}}^2 < \infty$, then the convergence is also P_0^∞ -almost surely.

The proof of the theorem is deferred to Section 4.1. The condition that the PY (σ, σ, G) process satisfies the central limit theorem in $\ell^\infty(\mathcal{F})$, is satisfied, for instance, for all classes \mathcal{F} that are suitably measurable with finite uniform entropy integral and for all classes \mathcal{F} with finite G -bracketing integral. This follows from Theorems 2.11.9 or 2.11.1 in [28].

The limit process in the theorem is Gaussian, but it is the P_0 -Brownian bridge \mathbb{G}_{P_0} only if $\lambda = 0$, i.e. if $P_0 = P_0^d$ is discrete. In addition, the behaviour of the Pitman-Yor posterior deviates from the “desired” behaviour by the presence on the left side of the term

$$\sqrt{n} B_n(f) := \frac{\sigma K_n}{\sqrt{n}} (G - \tilde{\mathbb{P}}_n). \quad (3)$$

Given the observations X_1, \dots, X_n this term is deterministic, and we can only expect it to disappear if K_n/\sqrt{n} tends to zero. While $K_n/n \rightarrow 0$ almost surely for any discrete distribution P_0 , the more stringent convergence to zero of K_n/\sqrt{n} is valid only if the sizes of the atoms of P_0 decrease fast enough. This relationship was made precise in [17] (also see the corollary below) in terms of the function

$$\alpha_0(u) = \#\{x : 1/P_0\{x\} \leq u\}. \quad (4)$$

If α_0 is regularly varying at $u = \infty$ (in the sense of Karamata, see e.g. [2] or the appendix to [9]) with exponent $\gamma_0 \in (0, 1)$, then $K_n/\alpha_0(n) \rightarrow \Gamma(1 - \gamma_0)$, almost surely, and $\alpha_0(n)$ is n^{γ_0} up to a slowly varying factor. In this case, for K_n/\sqrt{n} to tend to zero, it is necessary that the exponent be smaller than $1/2$ and sufficient that it is strictly smaller than $1/2$. For instance, if the ordered atoms $P_0\{x_j\}$ of P_0 decrease proportionally to $1/j^\alpha$, then $K_n/\sqrt{n} \rightarrow 0$ in probability if and only if $\alpha > 2$.

For bounded functions f , the convergence $K_n/\sqrt{n} \rightarrow 0$ is also enough to drive the additional term (3) to zero, as the terms $(G - \tilde{\mathbb{P}}_n)f$ will remain bounded in that case. For unbounded functions f , a still more stringent condition on P_0 is needed to make the term $(K_n/\sqrt{n})\tilde{\mathbb{P}}_n f$ go away. For instance, for the posterior mean of a distribution on \mathbb{N} with atoms $P_0\{j\}$ of the order $1/j^\alpha$, the next corollary implies that $\alpha > 4$ is needed.

We conclude that for a large class of discrete distributions P_0 , but not all, the Bernstein-von Mises theorem takes its standard form, and this also depends on which aspect of the posterior distribution we are interested in.

Corollary 2. *Under the conditions of Theorem 1, if P_0 is a discrete probability distribution, then $\sqrt{n}(P - \mathbb{P}_n - (\sigma K_n/n)(G - \tilde{\mathbb{P}}_n)|X_1, \dots, X_n \rightsquigarrow \mathbb{G}_{P_0}$ in $\ell^\infty(\mathcal{F})$, in probability or almost surely.*

- (i) *If the class of functions \mathcal{F} is uniformly bounded and the atoms $\{x_j\}$ of P_0 satisfy $P_0\{x_j\} \leq Cj^{-\alpha}$, for some constants C and $\alpha > 2$, then also $\sqrt{n}(P - \mathbb{P}_n)|X_1, \dots, X_n \rightsquigarrow \mathbb{G}_{P_0}$, in probability. If the class of functions \mathcal{F} is uniformly bounded and the function $u \mapsto \alpha_0(u) = \#\{x : 1/P_0\{x\} \leq u\}$ is regularly varying at $u = \infty$ of exponent strictly smaller than $1/2$, then this is also true almost surely.*
- (ii) *If the atoms $\{x_j\}$ of P_0 and the function f satisfy $P_0\{x_j\} \leq Cj^{-\alpha}$ and $f(x_j) \asymp j^p$, for some $p > 0$, then $\sqrt{n}(P - \mathbb{P}_n)f|X_1, \dots, X_n \rightsquigarrow \mathbb{G}_{P_0}f$, in probability if $\alpha > 2p + 2$.*

Proof. The first assertion merely specializes the limit in Theorem 1 to the case of a discrete distribution, by setting $\lambda = 0$. Assertions (i) and (ii) follow from this if the term (3) tends to zero, in probability or almost surely.

For bounded functions f , as assumed in (i), the term (3) tends to zero provided K_n/\sqrt{n} tends to zero. The almost sure convergence is immediate from [17], Theorems 9 and 1', which show that $K_n/\alpha_0(n) \rightarrow \Gamma(1 - \gamma_0)$, almost surely, for γ_0 the exponent of regular variation. For the convergence in probability, we note that $K_n = \sum_{j=1}^\infty 1_{j \in \{X_1, \dots, X_n\}}$, whence $\mathbb{E}K_n = \sum_{j=1}^\infty (1 - (1 - P_0\{x_j\})^n)$. By the inequality $(1 - p)^n \geq 1 - np$, for $p \geq 0$, we find that $\mathbb{E}K_n \leq \sum_{j=1}^\infty (nCj^{-\alpha} \wedge 1)$, which can be seen to be $o(\sqrt{n})$ if $\alpha > 2$.

The assertion in (ii) follows provided $K_n/\sqrt{n} \rightarrow 0$ and $(K_n/\sqrt{n})\tilde{\mathbb{P}}_n f \rightarrow 0$, in probability. Reasoning as before, we find

$$\mathbb{E}\left(\frac{K_n}{\sqrt{n}}\tilde{\mathbb{P}}_n f\right) = \frac{1}{\sqrt{n}} \sum_{j=1}^\infty f(x_j)(1 - (1 - P_0\{x_j\})^n).$$

Under the given assumptions on f and the atoms, this is bounded above by

$$\frac{1}{\sqrt{n}} \int_{C^{1/\alpha}}^\infty u^p \left(1 - \left(1 - \frac{C}{u^\alpha}\right)^n\right) du = \frac{n^{(p+1)/\alpha}}{\sqrt{n}} \int_{C/n}^\infty v^{(p+1)/\alpha-1} \left(1 - \left(1 - \frac{C}{nv}\right)^n\right) dv.$$

The integrand is bounded above by $Cv^{(p+1)/\alpha-1}$ and hence the integral converges near 0. By again the inequality $(1 - p)^n \geq 1 - np$, for $p \geq 0$, the integrand is also bounded above by $v^{(p+1)/\alpha-2}$ and hence the integral converges near infinity if $(p + 1)/\alpha < 1$. The middle part of the integral always gives a non-vanishing contribution and hence the full expression can tend to zero only if the leading factor tends to zero. This is true under the more stringent condition that $(p + 1)/\alpha < 1/2$. □

For $\lambda = 1$ and $P_0 = P_0^c$, Theorem 1 was obtained by [16]. In this case all observations are distinct and the left side of the theorem reduces to $\sqrt{n}(P - (1 - \sigma)\mathbb{P}_n - \sigma G)$, since $K_n = n$. As noted in the introduction, the posterior distribution is not even consistent, i.e. the asymptotic limit is “wrong” even without the \sqrt{n} multiplier.

The Bernstein-von Mises theorem is important for the validity of credible sets. A credible interval for Pf , for a given function f , could be formed as the interval between two quantiles of the marginal posterior distribution of Pf given X_1, \dots, X_n . For instance, for f equal to the indicator of a given set $A \in \mathcal{A}$, this gives a credible interval for a probability $P(A)$, and for $f(x) = x$, we obtain a credible interval for the mean. Simultaneous credible sets, for instance a credible band for a distribution function can be obtained similarly.

By the inconsistency of the posterior distribution in the case that the true distribution possesses a continuous component ($\lambda > 0$), there is no hope that in this case such an interval for Pf will cover a true value P_0f with the desired probability. However, also in the case of a discrete distribution P_0 , the coverage may not tend to the nominal value, due to the presence of the bias term (3). We need at least that K_n/\sqrt{n} tends to zero, and more for unbounded functions f .

Because the bias $B_n(f) = (\sigma K_n/n)(Gf - \tilde{\mathbb{P}}_n f)$ is observed (and σ and the center measure G are fixed by our prior choices), it is possible to correct a credible interval by shifting it by minus this amount. Thus for $Q_{n,\alpha}(f)$ the α -quantile of the posterior distribution of Pf given X_1, \dots, X_n , we consider both the credible intervals $[Q_{n,\alpha}(f), Q_{n,\beta}(f)]$ and corrected intervals $[Q_{n,\alpha}(f) - B_n(f), Q_{n,\beta}(f) - B_n(f)]$, for given $\alpha < \beta$.

Corollary 3. *Under the conditions of Theorem 1, if P_0 is a discrete probability distribution, then $\Pr_{P_0}(Q_{n,\alpha}(f) - B_n(f) \leq P_0f \leq Q_{n,\beta}(f) - B_n(f)) \rightarrow \beta - \alpha$, for every f with $(P_0 + G)f^2 < \infty$. If $\sqrt{n}B_n(f) \rightarrow 0$, in probability, then also $\Pr_{P_0}(Q_{n,\alpha}(f) \leq P_0f \leq Q_{n,\beta}(f)) \rightarrow \beta - \alpha$, for every such f . For bounded functions f , the latter is true if the atoms of P_0 satisfy $P_0\{x_j\} \leq Cj^{-\alpha}$, for some constants C and $\alpha > 2$. For $f(x) = x$, this is true for $\alpha > 4$.*

Proof. The α -quantile $Q_{n,\alpha}(f)$ of the posterior distribution of Pf is equal to $n^{-1/2}\bar{Q}_{n,\alpha}(f) + \mathbb{P}_n f + B_n(f)$, for $\bar{Q}_{n,\alpha}(f)$ the α -quantile of the posterior distribution of $\sqrt{n}(Pf - \mathbb{P}_n f - B_n(f))$. By Theorem 1, the latter posterior distribution tends to a normal distribution with mean zero and variance $\tau^2(f) = \text{Var } \mathbb{G}_{P_0} f$. It follows that

$$Q_{n,\alpha}(f) = \mathbb{P}_n f + B_n(f) + \frac{\tau(f)}{\sqrt{n}}\xi_\alpha + o_P\left(\frac{1}{\sqrt{n}}\right),$$

where ξ_α is the α -quantile of the standard normal distribution. Thus the event $P_0f \geq Q_{n,\alpha}(f) - B_n(f)$ can be rewritten as $P_0f \geq \mathbb{P}_n f + \tau(f)/\sqrt{n}\xi_\alpha + o_P(n^{-1/2})$. The probability of the latter event tends to $1 - \alpha$, by the central limit theorem applied to $\sqrt{n}(\mathbb{P}_n f - P_0f)$.

If $\sqrt{n}B_n(f)$ tends to zero in probability, then in the preceding display $B_n(f)$ can be incorporated into the $o_P(n^{-1/2})$ remainder term, and the remaining argument works for the uncorrected interval as well.

The final assertions follow from Corollary 2. □

Example. *The following explicit counterexample illustrates that the coverage can fail. Let G be the normal distribution with both mean and variance 1, let*

$P_0 = \sum_{j=1}^{\infty} p_j \delta_j$, for $p_j = 6/(\pi j)^2$, and consider the function $f = 1_{(1,\infty)} - 1_{(-\infty,1]}$. Since $Gf = 0$, we get $(\sigma K_n/\sqrt{n})(G - \tilde{\mathbb{P}}_n)f = (\sigma/\sqrt{n}) \sum_{i=1}^{K_n} f(\tilde{X}_i)$. Eventually the atom $\{1\}$ will be among the observations. Since $f(1) = -1$ and $f(j) = 1$ for all atoms $j \geq 2$, $(\sigma K_n/\sqrt{n})(G - \tilde{\mathbb{P}}_n)f = (\sigma/\sqrt{n})(-1 + (K_n - 1)) \rightarrow \sigma\sqrt{6/\pi}$, almost surely, by [17], Theorem 8, Theorem 1' and Example 4. The coverage of the uncorrected interval $[Q_{n,\alpha}(f), Q_{n,\beta}(f)]$ will tend to $\Phi(-\xi_\alpha - \sigma\sqrt{6/\pi}) - \Phi(-\xi_\beta - \sigma\sqrt{6/\pi})$.

The joint convergence in collections of functions f allows to study simultaneous credible sets and credible bands, besides univariate intervals. For instance, in the case that the sample space is the real line, we can take \mathcal{F} equal to the set of all indicators of cells $(-\infty, t]$, and obtain a credible band for the distribution function $F_0(t) = P_0(-\infty, t]$, as follows. Let $F(t) = P(-\infty, t]$ be the distribution function of the posterior process, and for $m_n(t)$ and $s_n(t)$ two functions dependent on X_1, \dots, X_n , let $\xi_{n,\alpha}$ be the α -quantile of the posterior distribution of $\sup_{t \in \mathbb{R}} |(F(t) - m_n(t))/s_n(t)|$. Consider the credible band of functions

$$C_n(\alpha) := \{F : m_n(t) - \xi_{n,1-\alpha}s_n(t) \leq F(t) \leq m_n(t) + \xi_{n,1-\alpha}s_n(t), \forall t\}.$$

Possible choices for the functions m_n and s_n are the pointwise posterior mean $m_n(t) = \mathbb{E}(F(t) | X_1, \dots, X_n)$ and the pointwise posterior standard deviation $s_n(t) = \text{sd}(F(t) | X_1, \dots, X_n)$. The quantiles $\xi_{n,\alpha}$ will typically be computed approximately from an MCMC sample from the posterior distribution, or approximated using tables for the limiting Brownian bridge process.

Corollary 4. *If P_0 is a discrete probability distribution with atoms such that $P_0\{x_j\} \leq Cj^{-2-\epsilon}$, for some constants C and $\epsilon > 0$, then $\Pr_{P_0}(F_0 \in C_n(\alpha)) \rightarrow 1 - 2\alpha$.*

Proof. Because the class of indicator functions is Donsker, both the classical empirical process $\{\sqrt{n}(\mathbb{F}_n - F_0)(t) : t \in \mathbb{R}\}$ and the posterior empirical process $\{\sqrt{n}(F - \mathbb{F}_n)(t) : t \in \mathbb{R}\} | X_1, \dots, X_n$ tend to the process $\mathbb{G}_U \circ F_0$, for \mathbb{G}_U a standard (classical) Brownian bridge process, by Theorem 1. The result follows from this along the same lines as the proof of Corollary 3. \square

The bias term (3) vanishes as $\sigma \downarrow 0$, which is in agreement with the fact that in this case the Pitman-Yor prior approaches the Dirichlet prior, which is well known to give asymptotically correct inference for any distribution P_0 . The bias term increases with σ , which is counterintuitive, as the bias appears only for heavy-tailed P_0 (having many large atoms), while large σ gives more different atoms in the prior.

One might hope that a data-dependent choice of σ could solve this bias problem. The empirical Bayes method is to estimate σ by the maximum likelihood estimator based on observing X_1, \dots, X_n in the Bayesian model, i.e. the maximiser of the marginal likelihood, and plug this into the posterior distribution of P for known σ . The hierarchical Bayes method is to put a prior on σ , and given σ , put the Pitman-Yor prior on P . Disappointingly, these methods do not change the limit behaviour of the posterior distribution of P . This is explained

by the fact that these methods yield a reasonable estimator of a value of σ connected to the discreteness of the true distribution P_0 , and we already noted the counterintuitive fact that a better match of discreteness does not solve the bias problem, but even makes it worse.

A sample X_1, \dots, X_n from a realisation of the Pitman-Yor process induces a (random) partition of the set $\{1, 2, \dots, n\}$ through the equivalence relation $i \equiv j$ if and only if $X_i = X_j$. An alternative way to generate the sample is to generate first the partition and next attach to each set in the partition a value generated independently from the center measure G (see e.g. [13], Lemma 14.11 for a precise statement), duplicating this as many times as there are indices in the set, in order to form the observations X_1, \dots, X_n . Because the parameter σ enters only in creating the partition, the partition is a sufficient statistic for σ . Because of exchangeability, the vector $(N_{n,1}, \dots, N_{n,K_n})$ of cardinalities of the partitioning sets is already sufficient for σ and hence the empirical Bayes estimator and posterior distribution of σ based on observations (X_1, \dots, X_n) or on observations $(K_n, N_{n,1}, \dots, N_{n,K_n})$ are the same.

The likelihood function for σ is therefore equal to the probability of a particular partition, called the exchangeable partition probability function (EPPF). For the Pitman-Yor process this is given by (see [22], or [13, page 465])

$$p_\sigma(N_{n,1}, \dots, N_{n,K_n}) = \frac{\prod_{i=1}^{K_n-1} (M + i\sigma)}{(M + 1)^{[n-1]}} \prod_{j=1}^{K_n} (1 - \sigma)^{[N_{n,j}-1]}. \quad (5)$$

Here $a^{[n]} = a(a+1)\cdots(a+n-1)$ is the ascending factorial, with $a^{[0]} = 1$ by convention. For the case that $M = 0$, it is shown in [11], that provided the partition is nontrivial ($1 < K_n < n$), the maximiser $\hat{\sigma}_n$ of this likelihood exists. Moreover, if the true distribution P_0 is discrete, with atoms satisfying, for $\alpha_0(u) = \#\{x : 1/P_0\{x\} \leq u\}$ and some $\sigma_0 \in (0, 1)$,

$$\sup_{u>1} \frac{|\alpha_0(u) - Lu^{\sigma_0}|}{\sqrt{u^{\sigma_0} \log(eu)}} < \infty,$$

then [11] shows that the maximum likelihood estimator satisfies $\hat{\sigma}_n = \sigma_0 + O_P(n^{-\sigma_0/2} \sqrt{\log n})$. Thus the coefficient of regular variation σ_0 may be viewed as a true value of σ , identified by the maximum likelihood estimator.

For the following theorem we need only the consistency of $\hat{\sigma}_n$, which we prove in Section 4.3 for general M , under the condition that α_0 is regularly varying. We also consider the full Bayes approach, and show that the posterior distribution of σ concentrates asymptotically around the empirical likelihood estimator, and hence contracts to σ_0 , under the same condition.

Theorem 5. *Let $P_0 = (1 - \lambda)P_0^d + \lambda P_0^c$ where P_0^d is a discrete and P_0^c an atomless probability distribution. If $\hat{\sigma}_n$ are estimators based on X_1, \dots, X_n such that $\hat{\sigma}_n \rightarrow \sigma_0$ in probability, and P is the posterior Pitman-Yor process of Theorem 1, then the process*

$$\sqrt{n} \left(P - \mathbb{P}_n - \frac{\hat{\sigma}_n K_n}{n} (G - \tilde{\mathbb{P}}_n) \right) \Big| X_1, \dots, X_n$$

tends to the same limit process as in Theorem 1 with σ replaced by σ_0 , in probability. If P_0 is discrete with atoms such that α_0 given in (4) is regularly varying of exponent $\sigma_0 \in (0, 1)$, then this is true for the maximum likelihood estimator $\hat{\sigma}_n$. Furthermore, in this case for Π_σ a prior distribution on σ with continuous positive density on $[0, 1]$, the posterior distribution of P in the model $\sigma \sim \Pi_\sigma$, $P|\sigma \sim \text{PY}(\sigma, M, G)$ and $X_1, \dots, X_n|P, \sigma \sim P$ satisfies the assertion of Theorem 1, with σ in the left side also interpreted as a random posterior variable and σ in the right side replaced by σ_0 . Finally if P_0 possesses a nontrivial atomless component (i.e. $\lambda > 0$), then $\hat{\sigma}_n \rightarrow \sigma_0 := 1$.

The proof of the theorem is deferred to Section 4.2. The final assertion of the theorem underlines again the deficiency of the Pitman-Yor process for distributions with a continuous component, which is not solved by estimating the type parameter. The type estimate tends to type 1 instead of the desired type 0 corresponding to the Dirichlet prior.

Besides the type parameter, the prior precision parameter M could be replaced by a data-dependent version. However, unlike the type parameter, this prior precision does not appear in the asymptotics of the posterior distribution of $\sqrt{n}(P - \mathbb{P}_n)$. Moreover, inspection of the proof of Theorems 1 and 5 shows that the convergence in these theorems is uniform in $M \ll \sqrt{n}$. Thus data-dependent M will not lead to new insights.

In the case of a discrete distribution P_0 for which the atoms decrease too slowly to ensure that K_n/\sqrt{n} tends to zero, the bias term (3) could still tend to zero if $\tilde{\mathbb{P}}_n \rightarrow G$. However, we show below that $\tilde{\mathbb{P}}_n(A) \rightarrow 0$, for any set A that contains only finitely many atoms of P_0 , and hence such convergence is false in any reasonable sense. Furthermore, the (in)consistency result (1) shows that in the case that P_0 possesses a continuous component, the center measure $G = P_0^c$ is the only choice for which the posterior distribution is even consistent. A data-dependent center measure might achieve this, but in the present context would come down to the original problem of estimating P_0 . Hierarchical choices (and hence random) of the center measure are considered in [3, 4], but with the different aim of finding hierarchical structures in the data.

Lemma 6. *If P_0 is discrete with infinitely many support points, then $\tilde{\mathbb{P}}_n f \rightarrow 0$ in probability for any bounded function f with finite support. Furthermore, $\tilde{\mathbb{P}}_n f \rightarrow f_\infty$ in probability, for any bounded function f for which there exists a number f_∞ such that $\sup_{x: P_0\{x\} < \delta} |f(x) - f_\infty| \rightarrow 0$, as $\delta \downarrow 0$.*

Proof. Let x_j be the atoms of P_0 , ordered by decreasing size $p_j := P_0\{x_j\}$ and set $f_j = f(x_j)$. Arguing as in the proof of Corollary 2, we can obtain (for the variance also see [17], formulas (39)–(40))

$$\begin{aligned} \mathbb{E}(K_n \tilde{\mathbb{P}}_n f) &= \sum_{j=1}^{\infty} f_j (1 - (1 - p_j)^n), \\ \text{Var}(K_n \tilde{\mathbb{P}}_n f) &= \sum_{j=1}^{\infty} f_j^2 [(1 - p_j)^n - (1 - p_j)^{2n}] \end{aligned}$$

$$+ \sum_{i \neq j} f_i f_j [(1 - p_i - p_j)^n - (1 - p_i)^n (1 - p_j)^n].$$

For $f = 1$ these expressions reduce to $\mathbb{E}K_n$ and $\text{Var} K_n$. As all terms of the series in $\mathbb{E}K_n$ tend to 1, it can be seen that $\mathbb{E}K_n \rightarrow \infty$. Furthermore, it can be seen that $\text{Var} K_n \leq \mathbb{E}K_n$, as the terms in the second series in $\text{Var} K_n$ are negative and the terms of the first series are bounded above by the terms in $\mathbb{E}K_n$. Because the terms of the series tend to zero as $n \rightarrow \infty$, for fixed i, j , and $f_j \rightarrow f_\infty$, as $j \rightarrow \infty$, for general f as in the second assertion of the lemma, the expressions are asymptotically equivalent to $f_\infty \mathbb{E}K_n + o(1)$ and $f_\infty^2 \text{Var} K_n + o(1)$, as $n \rightarrow \infty$. It follows that

$$\text{Var} \left(\frac{K_n \tilde{\mathbb{P}}_n f}{\mathbb{E}(K_n \tilde{\mathbb{P}}_n f)} - 1 \right) = \frac{\text{Var}(K_n \tilde{\mathbb{P}}_n f)}{(\mathbb{E}(K_n \tilde{\mathbb{P}}_n f))^2} = \frac{f_\infty^2 \text{Var} K_n + o(1)}{(f_\infty \mathbb{E}K_n + o(1))^2}.$$

Since $\text{Var} K_n \leq \mathbb{E}K_n \rightarrow \infty$, the right side tends to zero if $f_\infty \neq 0$. Then $K_n \tilde{\mathbb{P}}_n f / \mathbb{E}(K_n \tilde{\mathbb{P}}_n f) \rightarrow 1$, in probability. Taking $f = 1$, we see that $K_n / \mathbb{E}K_n \rightarrow 1$, in probability. Combining the preceding, we conclude that $\tilde{\mathbb{P}}_n f / f_\infty \rightarrow 1$, in probability.

If $f_\infty = 0$, then it follows that $\mathbb{E}(K_n \tilde{\mathbb{P}}_n f) \rightarrow 0$. Combination with the fact that $K_n \rightarrow \infty$, almost surely, gives that $\tilde{\mathbb{P}}_n f \rightarrow 0$, in probability.

If f has finite support, then the condition on f in the second part holds with $f_\infty = 0$ and hence $\tilde{\mathbb{P}}_n f \rightarrow 0$, in probability. \square

3. Numerical illustration

To illustrate that credible sets can be off, we carried out three simulation experiments, involving three discrete true probability distributions P_1, P_2, P_3 on \mathbb{N} . We focused on a credible interval for the probability of the set $[2, \infty)$. The measure P_1 is finitely discrete and given in Table 1, while P_2 and P_3 are given by the formulas

$$P_2\{k\} \propto \frac{1}{k^2}, \quad P_3\{k\} \propto \frac{1}{k^{1.5}}.$$

By the results of [17], as $n \rightarrow \infty$ the number K_n of distinct observations in a sample of size n from these distributions are asymptotically equal to 6, and proportional to \sqrt{n} and to $n^{2/3}$, respectively, for P_1, P_2 and P_3 . Thus K_n / \sqrt{n} tends to 0, a positive constant and ∞ , respectively, and a bias is expected for P_2 and P_3 , but not for P_1 , where P_2 is a boundary case.

TABLE 1
Probability distribution P_1

k	1	2	3	4	5	6
$P_1(X = k)$	0.1	0.1	0.2	0.2	0.3	0.1

As prior parameters we used $\sigma = 1/2$ and $M = 1$ and G the normal distribution with mean and variance 1. The choice $M = 1$ means that the prior is not biased exceedingly against the true distribution.

The Pitman-Yor posterior distribution can be simulated using the explicit representation given by [21] (see Section 4). Following Algorithm 1 from [1], we truncated the infinite series in the representation at a finite value, ensuring that the total weight of the tail is smaller than $n^{-1/2}$ so that the approximation is accurate within our context. We simulated 10000 samples from each of P_1 , P_2 and P_3 and for five different sample sizes: $n = 10, 10^2, 10^3, 10^4, 10^5$. For each sample we computed a 95% credible interval for $P[2, \infty)$ from its marginal posterior distribution, constructed using the 0.025 and the 0.975 posterior quantiles. We next computed coverage as the proportion of the 10000 replications that the true value, $P_1[2, \infty)$, $P_2[2, \infty)$ or $P_3[2, \infty)$, belonged to the interval. We did the same with the credible interval shifted by $(\sigma K_n / \sqrt{n})(G[2, \infty) - \tilde{\mathbb{P}}_n[2, \infty))$, derived from (3).

Tables 2 and 3 summarise the results. For P_1 both the corrected uncorrected intervals perform satisfactorily, whereas for P_2 and P_3 the uncorrected intervals undercover, severely so for P_3 , while the corrected intervals perform reasonably well, although not perfectly. The simulation results thus confirm the theoretical findings.

TABLE 2
Coverage of uncorrected posterior 95% credible intervals

n	10	100	1000	10000	100000
P_1	0.660	0.940	0.957	0.940	0.947
P_2	0.707	0.772	0.790	0.845	0.838
P_3	0.559	0.231	0.035	0.0	0.0

TABLE 3
Coverage of corrected posterior 95% credible intervals

n	10	100	1000	10000	100000
P_1	0.990	0.967	0.958	0.942	0.945
P_2	0.814	0.941	0.958	0.971	0.971
P_3	0.884	0.956	0.959	0.985	0.949

To illustrate the asymptotic normality of the posterior distribution, Figure 1 shows density plots of the marginal posterior distribution of $P[2, \infty)$, given samples of various sizes from P_1 . The plots were computed from the 100000 replicates, using the R “density” function. The normal approximation is satisfactory for $n = 1000$, but the posterior is visibly skewed for $n = 100$.

4. Proofs

Let $\tilde{X}_1, \dots, \tilde{X}_{K_n}$ be the distinct values in X_1, \dots, X_n , and let $N_{1,n}, \dots, N_{K_n,n}$ be their multiplicities. By Corollary 20 in [22] (or see [13, Theorem 14.37]), the posterior distribution of the Pitman-Yor process can be characterised as the distribution of

$$PY_n = R_n S_n + (1 - R_n) Q_n, \tag{6}$$




FIG 1. Density of the marginal posterior distribution of $P[2, \infty)$ based on n observations from P_1 , for $n = 10, 100, 1000$ (top row) and $n = 10^4, 10^5$. The true value of the parameter is $P_1[2, \infty) = 0.9$.

for

$$S_n = \sum_{i=1}^{K_n} W_{n,i} \delta_{\tilde{X}_i}, \quad (7)$$

and independent variables R_n, W_n, Q_n with, conditionally on X_1, \dots, X_n , distributed according to:

- $R_n \sim \text{B}(n - \sigma K_n, M + \sigma K_n)$,
- $Q_n \sim \text{PY}(\sigma, M + \sigma K_n, G)$,
- $W_n = (W_{n,1}, \dots, W_{n,K_n}) \sim \text{Dir}(K_n; N_{n,1} - \sigma, \dots, N_{n,K_n} - \sigma)$.

Here B and Dir refer to the beta and Dirichlet distributions, respectively. The number K_n will tend almost surely to the total number of atoms of P_0 in the case that P_0 is finitely discrete, and it will tend to infinity otherwise. In the latter case the rate of growth can have any order n^γ , for $0 < \gamma \leq 1$. (See Theorem 8 of [17], where it is shown that $K_n/\mathbb{E}K_n \rightarrow 1$, almost surely, where any rate can occur for $\mathbb{E}K_n$.) The proofs below use that K_n/n tends to the mass λ of the continuous part of P_0 , and the limit of the related sequence $n^{-1}K_n\tilde{\mathbb{P}}_n f = n^{-1}\sum_{i=1}^{K_n} f(\tilde{X}_i)$.

Lemma 7. *The number K_n of distinct values among $X_1, \dots, X_n \stackrel{iid}{\sim} \lambda P_0^c + (1 - \lambda)P_0^d$ satisfies $K_n/n \rightarrow \lambda$, almost surely. The number K_n^d of those values that belong to the set S of atoms of P_0^d satisfies $K_n^d/n \rightarrow 0$, almost surely.*

Proof. The number of distinct values not in S is $K_n^c := n\mathbb{P}_n(S^c)$ and hence $K_n^c/n \rightarrow P_0(S^c) = \lambda$, almost surely. If $S = \{x_1, x_2, \dots\}$, then the number of distinct values in S is bounded above by $m + n\mathbb{P}_n\{x_{m+1}, x_{m+2}, \dots\}$, for any m , and hence $K_n^d/n \leq m/n + \mathbb{P}_n\{x_{m+1}, x_{m+2}, \dots\} \rightarrow P_0\{x_{m+1}, x_{m+2}, \dots\}$, almost surely, for every m . \square

A class \mathcal{F} of measurable functions $f : \mathcal{X} \rightarrow \mathbb{R}$ is P_0 -Glivenko-Cantelli if the uniform law of large numbers holds: $\sup_{f \in \mathcal{F}} |\mathbb{P}_n f - P_0 f| \rightarrow 0$, outer almost surely (see e.g. [28], Chapter 2.4; we write “outer” because the supremum may not be measurable; for standard examples this is superfluous). An *envelope function* of \mathcal{F} is a measurable function $F : \mathcal{X} \rightarrow \mathbb{R}$ such that $|f| \leq F$, for every $f \in \mathcal{F}$.

Lemma 8. *Suppose \mathcal{F} has an envelope function with $P_0 F < \infty$. If S is the set of atoms of P_0^d , then $\sup_{f \in \mathcal{F}} |n^{-1} \sum_{i=1}^{K_n} f(\tilde{X}_i) 1_{\tilde{X}_i \in S}| \rightarrow 0$, outer almost surely. Furthermore, if $\sigma_n \rightarrow \sigma \in [0, 1]$, and \mathcal{F} is a P_0 -Glivenko-Cantelli class, then uniformly in $f \in \mathcal{F}$, outer almost surely,*

$$\mathbb{P}_n f + \frac{\sigma_n K_n}{n} \tilde{\mathbb{P}}_n f \rightarrow T f := (1 - \lambda) P_0^d f + (1 - \sigma) \lambda P_0^c f.$$

Proof. For any M , the supremum is bounded above by $n^{-1} K_n^d M + \mathbb{P}_n F 1_{F > M} \rightarrow P_0 F 1_{F > M}$, almost surely. The first term tends to zero by Lemma 7, for any M . The second term can be made arbitrarily small by choosing M large.

For the convergence in the display we write $K_n \tilde{\mathbb{P}}_n f = \sum_{i=1}^{K_n} f(X_i) 1_{X_i \notin S} + \sum_{i=1}^{K_n} f(\tilde{X}_i) 1_{\tilde{X}_i \in S}$. By the first assertion, the second sum divided by n tends to zero, uniformly in f . The first sum divided by n tends to $\lambda P_0^c f$, where the convergence is uniform in $f \in \mathcal{F}$ if \mathcal{F} is a Glivenko-Cantelli class (which implies that the set of functions $x \mapsto f(x) 1_{S^c}(x)$ is a Glivenko-Cantelli class, in view of [29]). Thus the left side of the display tends to $P_0 f + \sigma \lambda P_0^c f$, which is equal to $T f$. \square

4.1. Proof of Theorem 1

The left side $\sqrt{n}(\text{PY}_n - \mathbb{P}_n + (\sigma K_n/n)(\tilde{\mathbb{P}}_n - G))$ of the theorem can be decomposed as

$$\begin{aligned} \sqrt{n} \left(R_n - 1 + \frac{\sigma K_n}{n} \right) (S_n - Q_n) + \sqrt{n} \left(S_n \left(1 - \frac{\sigma K_n}{n} \right) - \mathbb{P}_n + \frac{\sigma K_n}{n} \tilde{\mathbb{P}}_n \right) \\ + \sigma \sqrt{K_n} (Q_n - G) \sqrt{\frac{K_n}{n}}. \end{aligned} \tag{8}$$

We derive the limit distributions of these three terms in Lemmas 9–11 below. For later use it will be helpful to allow $\sigma \in (0, 1)$ to depend on n . For this reason we give precise proofs of the first two lemmas, although they are very similar to results obtained in [16, 13]. The main novelty is in the third lemma. For simplicity we assume that $\sigma_n \in (0, 1)$ converges to a limit, which we allow to be 0 or 1.

Lemma 9. *If $\sigma_n \rightarrow \sigma \in [0, 1]$, then*

$$\sqrt{n} \left(R_n - 1 + \frac{\sigma_n K_n}{n} \right) | X_1, \dots, X_n \rightsquigarrow N(0, (1 - \sigma \lambda) \sigma \lambda), \quad a.s. \tag{9}$$

Proof. We can represent the beta variable R_n as the quotient $R_n = U_n/(U_n + V_n)$, for independent gamma variables $U_n \sim \Gamma(u_n, 1)$ and $V_n \sim \Gamma(v_n, 1)$, for $u_n = n - \sigma_n K_n$ and $v_n = M + \sigma_n K_n$ the means, and also variances, of the latter variables. We can decompose

$$(U_n + V_n) \left(R_n - \frac{u_n}{u_n + v_n} \right) = \frac{v_n}{u_n + v_n} (U_n - u_n) - \frac{u_n}{u_n + v_n} (V_n - v_n).$$

Since $\sigma_n K_n/n \rightarrow \sigma\lambda \in [0, 1]$, we have $v_n/(u_n + v_n) \rightarrow \sigma\lambda$ and $u_n/(u_n + v_n) \rightarrow 1 - \sigma\lambda$. Furthermore, $(U_n + V_n)/n \rightarrow 1$, almost surely, by the law of large numbers.

If $\sigma\lambda < 1$, then $n - \sigma_n K_n \rightarrow \infty$ and hence $(U_n - u_n)/\sqrt{u_n} \rightsquigarrow Z_1 \sim N(0, 1)$, by the central limit theorem. It follows that $(U_n - u_n)/\sqrt{n} \rightsquigarrow Z_1 \sqrt{1 - \sigma\lambda}$. If $\sigma\lambda = 1$, then $\text{Var}(U_n/\sqrt{n}) = u_n/n \rightarrow 0$ and hence $(U_n - u_n)/\sqrt{n} \rightsquigarrow 0$, where the limit 0 is identical to $Z_1 \sqrt{1 - \sigma\lambda}$ in this case. Thus in all cases $(U_n - u_n)/\sqrt{n} \rightsquigarrow Z_1 \sqrt{1 - \sigma\lambda}$.

If $\sigma\lambda > 0$, then $\sigma_n K_n \rightarrow \infty$ and hence $(V_n - v_n)/\sqrt{v_n} \rightsquigarrow Z_2 \sim N(0, 1)$, by the central limit theorem. It follows that $(V_n - v_n)/\sqrt{n} \rightsquigarrow Z_2 \sqrt{\sigma\lambda}$. If $\sigma\lambda = 1$, then $\text{Var}(V_n/\sqrt{n}) = v_n/n \rightarrow 0$ and hence $(V_n - v_n)/\sqrt{n} \rightsquigarrow 0$, where the limit 0 is identical to $Z_2 \sqrt{\sigma\lambda}$ in this case. Thus in all cases $(V_n - v_n)/\sqrt{n} \rightsquigarrow Z_2 \sqrt{\sigma\lambda}$.

Combining the preceding, we see that the sequence $\sqrt{n}(R_n - u_n/(u_n + v_n))$ converges weakly to $\sigma\lambda Z_1 \sqrt{1 - \sigma\lambda} + (1 - \sigma\lambda) Z_2 \sqrt{\sigma\lambda}$. As the limit variable has variance $(1 - \sigma\lambda)\sigma\lambda$ and $u_n/(u_n + v_n) = (1 - \sigma_n K_n/n)(1 + O(1/n))$, this concludes the proof. \square

Lemma 10. *If $\sigma_n \rightarrow \sigma \in [0, 1]$ and $K_n \rightarrow \infty$ and \mathcal{F} is a class of finitely many G -square-integrable functions, then in $\mathbb{R}^{\mathcal{F}}$,*

$$\sigma_n \sqrt{K_n} (Q_n - G) | X_1, \dots, X_n \rightsquigarrow \sqrt{\sigma(1 - \sigma)} \mathbb{G}_G. \quad \text{a.s.} \quad (10)$$

The convergence is also true in $\ell^\infty(\mathcal{F})$ if \mathcal{F} possesses a G -square integrable envelope function and the Pitman-Yor process $\text{PY}(\sigma, \sigma, G)$ satisfies the central limit theorem in this space.

Proof. The process $Q_n \sim \text{PY}(\sigma_n, M + \sigma_n K_n, G)$ centered at mean zero can be represented as

$$Q_n - G \sim \sum_{i=0}^{K_n} W_{n,i} (P_i - G),$$

where $(W_{n,0}, \dots, W_{n,K_n}) \sim \text{Dir}(K_n + 1; M, \sigma_n, \dots, \sigma_n)$ is independent of the independent processes $P_0 \sim \text{PY}(\sigma_n, M, G)$ and $P_i \stackrel{\text{iid}}{\rightsquigarrow} \text{PY}(\sigma_n, \sigma_n, G)$, for $i = 1, \dots, K_n$ (see e.g. Proposition 14.35 in [13]). The variable $W_{n,0}$ is $B(M, K_n \sigma_n)$ -distributed, whence

$$\sigma_n \sqrt{K_n} \mathbb{E} |W_{n,0} (P_0 - G) f| = \frac{\sigma_n \sqrt{K_n} M}{M + K_n \sigma_n} \mathbb{E} |P_0 - G| f \leq \frac{M}{\sqrt{K_n}} \sqrt{G f^2},$$

where the moment of $(P_0 - G)f$ can be obtained from Proposition 14.34 in [13]. Next by the gamma representation of the Dirichlet distribution (e.g. Propositions G.2 and G.3 in [13]), we can represent

$$\sigma_n \sqrt{K_n} \sum_{i=1}^{K_n} W_{n,i}(P_i - G) \sim (1 - W_{n,0}) \frac{K_n^{-1/2} \sum_{i=1}^{K_n} V_{n,i}(P_i - G)}{K_n^{-1} \sum_{i=1}^{K_n} V_{n,i}/\sigma_n},$$

where the variables $V_{i,n} \stackrel{\text{iid}}{\sim} \Gamma(\sigma_n, 1)$ are independent of $W_{n,0}$ and the P_i . The triangular array of variables $V_{n,i}(P_i - G)$ are i.i.d. for every n with

$$\begin{aligned} \mathbb{E}V_{n,1}^2((P_1 - G)f)^2 &= \sigma_n(1 + \sigma_n)G(f - Gf)^2 \frac{1 - \sigma_n}{1 + \sigma_n}, \\ \mathbb{E}V_{n,1}^2((P_1 - G)f)^2 1_{|V_{n,1}(P_1 - G)f| \geq M_n} &\rightarrow 0, \end{aligned}$$

for any $M_n \rightarrow \infty$. The second claim is implied by the uniform integrability of the set of variables $W_\sigma := V_\sigma^2((P_\sigma - G)f)^2$, for $\sigma \in [0, 1]$, where $V_\sigma \sim \Gamma(\sigma, 1)$ is independent of $P_\sigma \sim \text{PY}(\sigma, \sigma, G)$, and W_0 and W_1 are defined to be degenerate at 0, in agreement with the first line of the preceding display. This itself is a consequence of the continuity of the map $\sigma \mapsto W_\sigma$ from $[0, 1]$ to $L_2(\Omega)$ and the Dunford-Pettis theorem. The continuity follows from the norm continuity, $\mathbb{E}W_{\sigma_n}^2 \rightarrow \mathbb{E}W_\sigma^2$, if $\sigma_n \rightarrow \sigma$, by the first assertion in the display, combined with the continuity in distribution of $\sigma \mapsto W_\sigma$. Therefore, the sequence $K_n^{-1/2} \sum_{i=1}^{K_n} V_{n,i}(P_i - G)$ tends to a normal distribution with mean zero and variance $\sigma(1 - \sigma)G(f - Gf)^2$, by the Lindeberg central limit theorem. The linearity of the process in f shows that as a process it tends marginally in distribution to the process $\sqrt{\sigma(1 - \sigma)}\mathbb{G}_G$. Because $\text{Var}(K_n^{-1} \sum_{i=1}^{K_n} V_{n,i}/\sigma_n) = 1/(K_n\sigma_n)$, we have $K_n^{-1} \sum_{i=1}^{K_n} V_{n,i}/\sigma_n \rightarrow 1$, in probability, if $K_n\sigma_n \rightarrow \infty$. Since also $1 - W_{n,0} \rightarrow 1$, the proof is complete in the case that $K_n\sigma_n \rightarrow \infty$.

If $K_n\sigma_n$ remains bounded, then necessarily $\sigma_n \rightarrow 0$, as $K_n \rightarrow \infty$, by assumption. Then

$$\begin{aligned} \sigma_n^2 K_n \mathbb{E} \left(\sum_{i=0}^{K_n} W_{n,i}(P_i - G)f \right)^2 &= \sigma_n^2 K_n \sum_{i=0}^{K_n} \sum_{j=0}^{K_n} \mathbb{E}W_{n,i}W_{n,j}(P_i - G)f(P_j - G)f \\ &\leq \sigma_n^2 K_n \mathbb{E} \left(\sum_{i=0}^{K_n} W_{n,i} \right)^2 Gf^2 \leq \sigma_n^2 K_n Gf^2. \end{aligned}$$

Since this tends to zero, the lemma holds also in this case, with a limit process equal to 0, which is equal to $\sqrt{\sigma(1 - \sigma)}\mathbb{G}_G$.

For the final assertion we note that the preceding argument gives the convergence of $\sup_{f \in \mathcal{F}} \sigma_n \sqrt{K_n} W_{n,0}(P_0 - G)f$ to zero for any class \mathcal{F} with square-integrable envelope function. The convergence of $K_n^{-1/2} \sum_{i=1}^{K_n} V_{n,i}(P_i - G)$ in $\infty(\mathcal{F})$ follows from the convergence of $K_n^{-1/2} \sum_{i=1}^{K_n} (P_i - G)$ by the multiplier central limit theorem (e.g. Lemma 2.9.1 and Theorem 2.9.2 in [28]). \square

Lemma 11. *If $\sigma_n \rightarrow \sigma \in [0, 1]$, where $\sigma\lambda < 1$, then for any P_0 -Donsker class with square-integrable envelope function*

$$\sqrt{n} \left(S_n \left(1 - \frac{\sigma_n K_n}{n} \right) - \mathbb{P}_n + \frac{\sigma K_n \tilde{\mathbb{P}}_n}{n} \right) \rightsquigarrow \mathbb{W} - \frac{1}{1 - \sigma\lambda} \mathbb{W}1T, \quad \text{a.s.}, \quad (11)$$

in $\ell^\infty(\mathcal{F})$, where $\mathbb{W} = \sqrt{\lambda(1 - \sigma)}\mathbb{G}_{P_0^c} + \sqrt{1 - \lambda}\mathbb{G}_{P_0^d}$, for independent Brownian bridge processes $\mathbb{G}_{P_0^c}$ and $\mathbb{G}_{P_0^d}$, and T is the (deterministic) process defined in Lemma 8. The convergence is true in probability for any P_0 -Donsker class. If $\sigma_n \rightarrow \sigma \in [0, 1]$, where $\sigma\lambda = 1$, then the sequence of processes tends to the zero proces.

Proof. A gamma representation for the multinomial vector W_n in the definition of S_n is

$$W_{n,i} = \frac{U_{i,0} + \sum_{j=1}^{N_{n,i}-1} U_{i,j}}{\sum_{i=1}^{K_n} \left(U_{i,0} + \sum_{j=1}^{N_{n,i}-1} U_{i,j} \right)},$$

for all $U_{i,j}$ independent, $U_{i,0} \sim \Gamma(1 - \sigma, 1)$ and $U_{i,j} \sim \Gamma(1, 1)$, for $j \geq 1$. Relabel the n variables $U_{i,j}$ as $\xi_{n,1}, \dots, \xi_{n,n}$, as follows. Let S be the set of all atoms of P_0 . An observation X_i that is not contained in S appears exactly once in the set $\{X_1, \dots, X_n\}$ of observations; set the variable $\xi_{n,i}$ with the corresponding i equal to $U_{i,0}$. Every X_i that is contained in S appears $N_{n,i} \geq 1$ times among X_1, \dots, X_n ; set the $\xi_{n,j}$ with indices corresponding to these appearances equal to $U_{i,0}, U_{i,1}, \dots, U_{i,N_{n,i}-1}$. Then

$$S_n = \sum_{i=1}^{K_n} W_{n,i} f(\tilde{X}_i) = \frac{n^{-1} \sum_{i=1}^n \xi_{n,i} f(X_i)}{n^{-1} \sum_{i=1}^n \xi_{n,i}} =: \frac{\bar{S}_n f}{\bar{S}_n 1}, \quad (12)$$

and the left side of the lemma can be decomposed as

$$\begin{aligned} S_n f \sqrt{n} \left(1 - \frac{\sigma_n K_n}{n} - \bar{S}_n 1 \right) + \sqrt{n} \left(\bar{S}_n f - \mathbb{P}_n f + \frac{\sigma_n K_n \tilde{\mathbb{P}}_n f}{n} \right) \\ = -S_n f \sqrt{n} (\bar{S}_n 1 - T_n 1) + \sqrt{n} (\bar{S}_n f - T_n f), \end{aligned}$$

where $T_n f = \mathbb{P}_n f - (\sigma_n K_n/n)\tilde{\mathbb{P}}_n f$ tends to Tf , by Lemma 8. We shall show that $\sqrt{n}(\bar{S}_n - T_n) | X_1, \dots, X_n \rightsquigarrow \mathbb{W}$. Then $S_n f \rightarrow Tf/T1 = Tf/(1 - \sigma\lambda)$, and the result follows in the case that $\sigma\lambda < 1$.

The variables $\xi_{n,1}, \dots, \xi_{n,n}$ are independent. The K_n variables corresponding to the distinct values are $\Gamma(1 - \sigma, 1)$ -distributed; the others are $\Gamma(1, 1)$ -distributed. Thus the conditional mean and variance of $\bar{S}_n f$ are given by

$$\begin{aligned} \sum_{i=1}^n (\mathbb{E} \xi_{n,i}) f(X_i) &= \sum_{i=1}^n f(X_i) - \sigma \sum_{i=1}^{K_n} f(\tilde{X}_i) = T_n f, \\ \frac{1}{n} \sum_{i=1}^n (\text{Var } \xi_{n,i}) f^2(X_i) &= \frac{1}{n} \sum_{i=1}^n f^2(X_i) - \frac{\sigma}{n} \sum_{i=1}^{K_n} f^2(\tilde{X}_i) \rightarrow Tf^2, \quad \text{a.s.}, \end{aligned}$$

by Lemma 8. The limit variance is equal to $\text{Var } \mathbb{W}f$. To complete the proof of the convergence $\sqrt{n}(\bar{S}_n - T_n)f | X_1, \dots, X_n \rightsquigarrow \mathbb{W}$, it suffices to verify the Lindeberg-Feller condition. We have, for $\xi_n \sim \Gamma(1 - \sigma_n, 1)$ and $\bar{\xi}_n \sim \Gamma(1, 1)$,

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left(\xi_{n,i}^2 f^2(X_i) 1_{|\xi_{n,i} f(X_i)| > \epsilon \sqrt{n}} | X_1, \dots, X_n \right) \\ & \leq \max \left(\mathbb{E} \xi_n^2 1_{|\xi_n| \max_{1 \leq i \leq n} |f(X_i)| > \epsilon \sqrt{n}}, \mathbb{E} \bar{\xi}_n^2 1_{|\bar{\xi}_n| \max_{1 \leq i \leq n} |f(X_i)| > \epsilon \sqrt{n}} \right) \mathbb{P}_n f^2. \end{aligned}$$

This tends to zero for every sequence X_1, X_2, \dots such that both $\mathbb{P}_n f^2 = O(1)$ and $\max_{1 \leq i \leq n} |f(X_i)|/\sqrt{n} \rightarrow 0$, which is almost every sequence if $P_0 f^2 < \infty$.

By the Cramér-Wold device and linearity in f , the convergence is then implied for finite sets of f .

For convergence as processes in $\ell^\infty(\mathcal{F})$ for a general Donsker class, it suffices to prove asymptotic tightness (see e.g. Theorem 1.5.4 in [28]). The processes $n^{-1/2} \sum_{i=1}^n (\xi_{n,i} - \mathbb{E} \xi_{n,i}) f(X_i)$ are multiplier processes with mean zero, independent multipliers. Because the multipliers are not i.i.d., a direct application of the conditional multiplier central limit theorem (see Theorem 2.9.7 in [28]) is not possible. However, the multipliers have two forms $\Gamma(1 - \sigma, 1)$ and $\Gamma(1, 1)$. By Jensen's inequality, for any collection \mathcal{G} of functions,

$$\mathbb{E}_\xi \left\| \sum_{i=1}^n (\xi_{n,i} - \mathbb{E} \xi_{n,i}) f(X_i) \right\|_{\mathcal{G}}^* \leq \mathbb{E}_{\xi, \xi'} \left\| \sum_{i=1}^n (\xi_{n,i} - \mathbb{E} \xi_{n,i} + \xi'_{n,i} - \mathbb{E} \xi'_{n,i}) f(X_i) \right\|_{\mathcal{G}}^*,$$

for any random variables $\xi'_{n,i}$ independent of the $\xi_{n,i}$. We can choose these variables so that all $\xi_{n,i} + \xi_{n,i} \stackrel{\text{iid}}{\sim} \Gamma(1, 1)$. The process in the right side then does have i.i.d. multipliers, and the asymptotic tightness follows from the i.i.d. case (as in [28]), Theorems 3.6.13, 2.9.6 and 2.9.7; also see Corollary 2.9.9; we apply the preceding inequality with \mathcal{G} equal to the set of differences $f - g$ of functions $f, g \in \mathcal{F}$ with $L_2(P_0)$ -norm of $f - P_0 f - g + P_0 g$ smaller than δ .

Finally if $\sigma\lambda = 1$, then both $\sigma = 1$ and $\lambda = 1$. The second implies that $P_0 = P_0^c$, $K_n = n$ and $\tilde{\mathbb{P}}_n = \mathbb{P}_n$. Thus in this case $S_n(1 - \sigma_n K_n/n) = \sum_{i=1}^n W_{n,i} f(X_i)(1 - \sigma_n)$, for $(W_{n,1}, \dots, W_{n,n}) \sim \text{Dir}(n, 1 - \sigma_n, \dots, 1 - \sigma_n)$, and $T_n f = \mathbb{P}_n f(1 - \sigma_n)$. We can now compute

$$\begin{aligned} \mathbb{E} \left(\sum_{i=1}^n W_{n,i} f(X_i) | X_1, \dots, X_n \right) &= \mathbb{P}_n f, \\ \text{Var} \left(\sum_{i=1}^n W_{n,i} f(X_i) | X_1, \dots, X_n \right) &= \sum_{i=1}^n \sum_{j=1}^n \text{Cov}(W_{n,i}, W_{n,j}) f(X_i) f(X_j) \\ &\leq \sum_{i=1}^n \frac{(n-1)f^2(X_i)}{n^2(n(1-\sigma_n)+1)} \leq \frac{\mathbb{P}_n f^2}{n(1-\sigma_n)}, \end{aligned}$$

as the covariances between the $W_{n,i}$ are negative. This implies that the conditional mean and variance of $\sqrt{n}(S_n(1 - \sigma_n) - T_n f)$ tend to zero, as $\sigma_n \rightarrow 1$. \square

We are ready to complete the proof of Theorem 1. If $K_n \rightarrow \infty$, then Lemmas 11–10 together with the convergence $K_n/n \rightarrow \lambda$ immediately give the convergence of the second and third terms in the decomposition (8). Furthermore, these lemmas give that $S_n - Q_n \rightarrow Tf/(1 - \sigma\lambda)$, which combined with Lemma 9 gives the convergence of the first term in (8).

If K_n remains bounded, then Lemma 10 does not apply. However, since the process Q_n will run through finitely many different Pitman-Yor processes, we have $Q_n - G = O_P(1)$ and hence the third term in (8) is $O_P(1/\sqrt{n})$, still under the assumption that K_n is bounded. Lemma 11 is still valid, and hence the second term in (8) converges to a Gaussian process as before. We can divide this term by $1 - \sigma K_n/n \rightarrow 1$, to see that $S_n \rightarrow T$, in view of Lemma 8. The sequence K_n can remain bounded only if $\lambda = 0$ and then the normal limit in Lemma 9 is degenerate, whence $\sqrt{n}(R_n - 1) = -\sigma K_n/\sqrt{n} + o_P(1) = o_P(1)$, almost surely, again under the assumption that K_n is bounded. Combined this shows that the first term in (8) tends to zero.

4.2. Proof of Theorem 5

Make the dependence on σ of the Pitman-Yor posterior process and its limit explicit by writing $\text{PY}_n(\sigma)$ and $\mathbb{G}(\sigma)$ for the process PY_n in (6) and the right side in Theorem 1, and set

$$\text{CPY}_n(\sigma) = \sqrt{n}(\text{PY}_n(\sigma) - \mathbb{P}_n - \frac{\sigma K_n}{n}(G - \tilde{\mathbb{P}}_n)).$$

Lemmas 9–11 give

$$\sup_{\sigma \in (0,1)} \sup_{h \in \text{BL}_1} \left| \mathbb{E} \left(h(\text{CPY}_n(\sigma)) \mid X_1, \dots, X_n \right) - \mathbb{E} h(\mathbb{G}(\sigma)) \right| \rightarrow 0, \quad (13)$$

in probability. This immediately gives that for every data-dependent $\hat{\sigma}_n$ that take their values in the interval $(0, 1)$,

$$\sup_{h \in \text{BL}_1} \left| \mathbb{E} \left(h(\text{CPY}_n(\hat{\sigma}_n)) \mid X_1, \dots, X_n \right) - \mathbb{E} h(\mathbb{G}(\hat{\sigma}_n)) \right| \rightarrow 0,$$

in probability, where the second expectation is on the limit process $\mathbb{G}(\hat{\sigma}_n)$ for given, fixed $\hat{\sigma}_n$. The continuity of the limit process in σ shows that, for $\hat{\sigma}_n \rightarrow \sigma_0$ in probability,

$$\sup_{h \in \text{BL}_1} \left| \mathbb{E} h(\mathbb{G}(\hat{\sigma}_n)) - \mathbb{E} h(\mathbb{G}(\sigma_0)) \right| \rightarrow 0,$$

in probability. Combined the two preceding displays give the first assertion of Theorem 5.

For discrete P_0 with regularly varying atoms, the convergence of the maximum likelihood estimator $\hat{\sigma}_n$ to its coefficient of regular variation $\sigma_0 \in (0, 1)$ is shown in Theorem 12, and hence the preceding argument applies.

In a hierarchical Bayesian setup with a prior on σ and given σ the Pitman-Yor prior on P , the posterior distribution of P can be decomposed as

$$\begin{aligned} &\mathbb{E}(h(\text{CPY}_n(\sigma) | X_1, \dots, X_n)) \\ &= \int \mathbb{E}\left(h(\text{CPY}_n(\sigma) | \sigma, X_1, \dots, X_n)\right) \Pi_n(d\sigma | X_1, \dots, X_n), \end{aligned}$$

where $\Pi_n(d\sigma | X_1, \dots, X_n)$ refers to the posterior distribution of σ given the observations X_1, \dots, X_n , and $\text{CPY}_n(\sigma) | \sigma, X_1, \dots, X_n$ is the standardised Pitman-Yor posterior distribution for given σ , considered in Theorem 1. The uniformity (13) shows that the expectation in the integral on the right side can be replaced asymptotically by $\mathbb{E}h(\mathbb{G}(\sigma))$, uniformly in $h \in \text{BL}_1$, whenever the posterior distribution of σ concentrates with probability tending to one on the interval $(0, 1)$. In particular, this is true if the posterior distribution of σ is consistent for some value $\sigma_0 \in (0, 1)$, i.e. if it concentrates asymptotically within the interval $(\sigma_0 - \epsilon, \sigma_0 + \epsilon)$, for every $\epsilon > 0$. This consistency is shown in the proposition below. Given posterior consistency, by the continuity of the limit process in σ , the expectation $\mathbb{E}h(\mathbb{G}(\sigma))$ can in turn in the limit be replaced by $\mathbb{E}h(\mathbb{G}(\sigma_0))$, uniformly in $h \in \text{BL}_1$. This gives the second assertion of Theorem 5.

4.3. Estimating the type parameter

A measurable function $\alpha : [1, \infty) \rightarrow \mathbb{R}_+$ is said to be *regularly varying* (at ∞) of order γ if, for all $u > 0$, as $n \rightarrow \infty$,

$$\frac{\alpha(nu)}{\alpha(n)} \rightarrow u^\gamma. \tag{14}$$

It is known (see e.g. [2] or the appendix to [9]) that if the limit of the sequence of quotients on the left exists for every u , then it necessarily has the form u^γ , for some γ , as in (14). If we write $\alpha(u) = u^\gamma L(u)$, then L will be *slowly varying*: a function that is regularly varying of order 0. Then $\alpha(n) = n^\gamma L(n)$, and it can be shown that $n^{\gamma-\delta} \ll \alpha(n) \ll \alpha^{n+\delta}$, for every $\delta > 0$, so that the rate of growth of α is n^γ to “first order”. (See Potter’s theorem, [2], Theorem 1.5.6, or [9], Proposition B.1.9-5).

Example. For the probability distribution $(p_j)_{j \in \mathbb{N}}$ with $p_j = C/j^\alpha$, for some $\alpha > 1$, the function $\alpha(u) := \#\{j : 1/p_j \leq u\} = \lfloor (Cu)^{1/\alpha} \rfloor$ is regularly varying of order $\gamma = 1/\alpha$.

We consider the empirical Bayes estimator $\hat{\sigma}_n$, the maximum likelihood estimator in the model $P | \sigma \sim \text{PY}(\sigma, M, G)$ and $X_1, \dots, X_n | P, \sigma \sim P$ given observations X_1, \dots, X_n . We also consider the posterior distribution of σ given X_1, \dots, X_n in the model $\sigma \sim \Pi_\sigma$, $P | \sigma \sim \text{PY}(\sigma, M, G)$ and $X_1, \dots, X_n | P, \sigma \sim P$, for a given prior distribution Π_σ on $(0, 1)$. In both cases the likelihood for observing X_1, \dots, X_n is proportional to (5). Hence $\hat{\sigma}_n$ is the point of maximum of this function and, by Bayes theorem, the posterior distribution has density relative to Π_σ proportional to (5).

In the following theorem we consider these objects under the assumption that X_1, \dots, X_n are an i.i.d. sample from a distribution P_0 . Consistency of $\hat{\sigma}_n$ for σ_0 means that $\hat{\sigma}_n \rightarrow \sigma_0$ in probability. Contraction of the posterior distribution to σ_0 means that $\Pi_n(\sigma : |\sigma - \sigma_0| > \epsilon | X_1, \dots, X_n)$ tends to zero in probability, for every $\epsilon > 0$.

Theorem 12. *If P_0 is discrete with atoms such that $\alpha_0(u) := \#\{x : 1/P_0\{x\} \leq u\}$ is regularly varying of exponent $\sigma_0 \in (0, 1)$, then the empirical Bayes estimator $\hat{\sigma}_n$ is consistent for σ_0 . Furthermore, for a prior distribution Π_σ on σ with a density that is bounded away from zero and infinity, the posterior distribution of σ contracts to σ_0 .*

Proof. Up to an additive constant the log likelihood can be written

$$\begin{aligned} \Lambda_n(\sigma) &= \sum_{l=1}^{K_n-1} \log(M + l\sigma) + \sum_{j=1: N_{n,j} \geq 2}^{K_n} \sum_{l=0}^{N_{n,j}-2} \log(1 - \sigma + l) \\ &= \sum_{l=1}^{K_n-1} \log(M + l\sigma) + \sum_{l=1}^{n-1} \log(l - \sigma) Z_{n,l+1}, \end{aligned}$$

where $Z_{n,l} = \#\{1 \leq j \leq K_n : N_{n,j} \geq l\}$ is the number of distinct values of multiplicity at least l in the sample X_1, \dots, X_n . (In the case that all observations are distinct and hence $N_{n,j} = 1$ for every j , the second term of the likelihood is equal to 0.) The concavity of the logarithm shows that the log likelihood is a strictly concave function of σ . For $\sigma \downarrow 0$, it tends to a finite value, while for $\sigma \uparrow 1$ it tends to $-\infty$ if the term with $l = 1$ is present in the second sum, i.e. if there is at least one tied observation. This happens with probability tending to 1 as $n \rightarrow \infty$. The derivative of the log likelihood is equal to

$$\Lambda'_n(\sigma) = \sum_{l=1}^{K_n-1} \frac{l}{M + l\sigma} - \sum_{l=1}^{n-1} \frac{1}{l - \sigma} Z_{n,l+1}. \quad (15)$$

The left limit at $\sigma = 0$ is $\Lambda'_n(0) = \frac{1}{2}K_n(K_n - 1) - \sum_{l=1}^{n-1} l^{-1} Z_{n,l+1}$. Since $Z_{n,l} \leq Z_{n,1} = K_n$, a crude bound on the sum is $K_n \log n$, which shows that the derivative at $\sigma = 0$ tends to infinity if $K_n \gg \log n$. In that case the unique maximum of the log likelihood in $[0, 1]$ is taken in the interior of the interval, and hence $\hat{\sigma}_n$ satisfies $\Lambda'_n(\hat{\sigma}_n) = 0$.

Under the condition that α_0 is regularly varying of exponent $\sigma_0 \in (0, 1)$, the sequence $\alpha_n := \alpha_0(n)$ is of the order n^{σ_0} up to slowly varying terms. By Theorems 9 and 1' of [17], the sequence K_n/α_n tends almost surely to $\Gamma(1 - \sigma_0)$ and hence in particular $K_n \gg \log n$.

We show below that $\Lambda'_n(\sigma)/\alpha_n \rightarrow \lambda(\sigma)$ in probability, for every σ , and a strictly decreasing function λ with $\lambda(\sigma_0) = 0$. It follows that $\Lambda'_n(\sigma_0 - \epsilon) > 0$ and $\Lambda'_n(\sigma_0 + \epsilon) < 0$ with probability tending to one, for every fixed $\epsilon > 0$. Then $\sigma_0 - \epsilon < \hat{\sigma}_n < \sigma_0 + \epsilon$ with probability tending to one, by the monotonicity of $\sigma \mapsto \Lambda'_n(\sigma)$, and hence the consistency of $\hat{\sigma}_n$ follows.

The monotonicity of Λ'_n and the fact that $\Lambda'_n(\hat{\sigma}_n) = 0$, give that on the event $\sigma_0 + \epsilon > \hat{\sigma}_n$,

$$\begin{aligned} \Lambda_n(\sigma) &\geq \Lambda_n(\sigma_0 + \epsilon), && \text{if } \hat{\sigma}_n < \sigma < \sigma_0 + \epsilon, \\ \Lambda_n(\sigma) &\leq \Lambda_n(\sigma_0 + \epsilon) + \Lambda'_n(\sigma_0 + \epsilon)(\sigma - \sigma_0 - \epsilon), && \text{if } \sigma > \sigma_0 + \epsilon. \end{aligned}$$

It follows that on the event $\sigma_0 + \epsilon > \hat{\sigma}_n$,

$$\begin{aligned} \Pi_n(\sigma > \sigma_0 + \epsilon | X_1, \dots, X_n) &= \frac{\int_{\sigma_0 + \epsilon}^1 e^{\Lambda_n(\sigma)} d\Pi_\sigma(\sigma)}{\int_0^1 e^{\Lambda_n(\sigma)} d\Pi_\sigma(\sigma)} \\ &\leq \frac{\int_{\sigma_0 + \epsilon}^1 e^{\Lambda_n(\sigma_0 + \epsilon) + \Lambda'_n(\sigma_0 + \epsilon)(\sigma - \sigma_0 - \epsilon)} d\Pi_\sigma(\sigma)}{\int_{\hat{\sigma}_n}^{\sigma_0 + \epsilon} e^{\Lambda_n(\sigma_0 + \epsilon)} d\Pi_\sigma(\sigma)} \\ &\lesssim \frac{\int_0^\infty e^{\Lambda'_n(\sigma_0 + \epsilon)u} du}{\sigma_0 + \epsilon - \hat{\sigma}_n} = \frac{1}{-\Lambda'_n(\sigma_0 + \epsilon)(\sigma_0 + \epsilon - \hat{\sigma}_n)}, \end{aligned}$$

where the proportionality constant depends on the density of Π_σ only. Since $-\Lambda'_n(\sigma_0 + \epsilon)/\alpha_n \rightarrow -\lambda(\sigma_0 + \epsilon) > 0$ and $\sigma_0 + \epsilon - \hat{\sigma}_n \rightarrow \epsilon$ in probability, the right side tends to zero in probability. Combined with a similar argument on the left tail of the posterior distribution, this shows that the posterior distribution contracts to σ_0 .

It remains to be shown that $\Lambda'_n(\sigma)/\alpha_n \rightarrow \lambda(\sigma)$, in probability for a strictly decreasing function λ with a unique zero at σ_0 . The variables $Z_{n,l}$ can be written as $Z_{n,l} = \sum_{j=1}^\infty 1_{M_{n,j} \geq l}$, for $M_{n,j}$ the number of observations equal to x_j . As $K_n = Z_{n,1}$, the function Λ'_n can be written in the form

$$\Lambda'_n(\sigma) = \sum_{l=1}^{K_n-1} \frac{l}{M + l\sigma} - \sum_{l=1}^\infty \sum_{j=1}^\infty \frac{1_{M_{n,j} \geq l+1}}{l - \sigma} = \sum_{j=1}^\infty \left[\frac{1_{M_{n,j} \geq 1}}{\sigma} - g_\sigma(M_{n,j}) \right] - \frac{h_\sigma(K_n)}{\sigma},$$

where $g_\sigma(0) = g_\sigma(1) = 0$ and $g_\sigma(m) = \sum_{l=1}^{m-1} \frac{1}{l - \sigma}$, for $m \geq 2$, and $h_\sigma(k) = 1 + \sum_{l=1}^{k-1} M/(M + l\sigma) \leq 1 + (M/\sigma) \log(1 + k\sigma/M)$. It is shown in [17] (and repeated below) that $\mathbb{E}K_n/\alpha_n \rightarrow \Gamma(1 - \sigma_0)$ and hence $\mathbb{E}h_\sigma(K_n) \leq 1 + (M/\sigma) \log(1 + \mathbb{E}K_n\sigma/M) = O(\log n) = o(\alpha_n)$, so that the term on the far right is asymptotically negligible.

It is shown in Lemma 13 that

$$\mathbb{E} \frac{1}{\alpha_n} \sum_{j=1}^\infty \left[\frac{1_{M_{n,j} \geq 1}}{\sigma} - g_\sigma(M_{n,j}) \right] \rightarrow \frac{\Gamma(1 - \sigma_0)}{\sigma} - \sum_{m=1}^\infty \frac{\Gamma(m + 1 - \sigma_0)}{m!(m - \sigma)} =: \lambda(\sigma).$$

The limit function λ is strictly decreasing. The value of the series at $\sigma = \sigma_0$ is equal to

$$\sum_{m=1}^\infty \frac{\Gamma(m - \sigma_0)}{m!} = \int_0^\infty (e^x - 1)x^{-\sigma_0-1}e^{-x} dx = \int_0^\infty (1 - e^{-x})x^{-\sigma_0-1} dx.$$

By partial integration, this can be further rewritten as $\int_0^\infty x^{-\sigma_0}/\sigma_0 e^{-x} dx = \Gamma(1 - \sigma_0)/\sigma_0$. We conclude that $\lambda(\sigma_0) = 0$.

To complete the proof it suffices to show that the variance of the variables in the left side of the second last display tend to zero. For $i \neq j$, the conditional distribution of $M_{n,i}$ given $M_{n,j} = m$ is binomial with parameters $(n - m, p_i)$, which is stochastically smaller than the marginal binomial (n, p_i) distribution of $M_{n,i}$. It follows that $\mathbb{E}(h(M_{n,i}) | M_{n,j}) \leq \mathbb{E}h(M_{n,i})$, for every nondecreasing function h , whence $h(M_{n,i})$ and $h(M_{n,j})$ are negatively correlated for every nonnegative, nondecreasing function h . Applying this with $h(m) = 1_{m \geq 1}$ and $h = g_\sigma$, we find that

$$\begin{aligned} \text{Var} \frac{1}{\alpha_n} \sum_{j=1}^{\infty} 1_{M_{n,j} \geq 1} &\leq \frac{1}{\alpha_n^2} \sum_{j=1}^{\infty} \text{Var} 1_{M_{n,j} \geq 1} \leq \frac{1}{\alpha_n^2} \sum_{j=1}^{\infty} \mathbb{E} 1_{M_{n,j} \geq 1}, \\ \text{Var} \frac{1}{\alpha_n} \sum_{j=1}^{\infty} g_\sigma(M_{n,j}) &\leq \frac{1}{\alpha_n^2} \sum_{j=1}^{\infty} \text{Var} g_\sigma(M_{n,j}) \leq \frac{1}{\alpha_n^2} \sum_{j=1}^{\infty} \mathbb{E} g_\sigma^2(M_{n,j}). \end{aligned}$$

By Lemma 13, both right sides are of the order $O(1/\alpha_n)$. This concludes the proof that $\Lambda'_n(\sigma)/\alpha_n \rightarrow \lambda(\sigma)$, in probability. \square

Lemma 13. *Suppose that $\alpha(u) := \#\{j : 1/p_j \leq u\}$ is regularly varying at ∞ of order $\gamma \in (0, 1)$. For any $\sigma \in (0, 1)$, and $g_\sigma(m) = \sum_{l=1}^{m-1} \frac{1}{l-\sigma}$, for $m \geq 2$, and $M_{n,j} \sim \text{Binomial}(n, p_j)$,*

$$\begin{aligned} (i) \quad &\frac{1}{\alpha(n)} \sum_{j=1}^{\infty} \mathbb{E} 1_{M_{n,j} \geq 1} \rightarrow \Gamma(1 - \gamma), \\ (ii) \quad &\frac{1}{\alpha(n)} \sum_{j=1}^{\infty} \mathbb{E} g_\sigma(M_{n,j}) \rightarrow \sum_{m=1}^{\infty} \frac{\Gamma(m+1-\gamma)}{m!(m-\sigma)}, \\ (iii) \quad &\frac{1}{\alpha(n)} \sum_{j=1}^{\infty} \mathbb{E} g_\sigma^2(M_{n,j}) \rightarrow \sum_{k=1}^{\infty} \sum_{l=1}^{\infty} \frac{\Gamma(k+l+1-\gamma)}{(k-\sigma)(l-\sigma)(k+l)!}. \end{aligned}$$

Proof. Because $\Pr(M_{n,j} = 0) = (1 - p_j)^n$, the series in the left side of (i) is equal to

$$\sum_{j=1}^{\infty} \left(1 - (1 - p_j)^n\right) = \int_1^{\infty} \left(1 - \left(1 - \frac{1}{u}\right)^n\right) d\alpha(u) = n \int_0^1 \alpha\left(\frac{1}{s}\right) (1 - s)^{n-1} ds,$$

by Fubini's theorem, since $1 - (1 - 1/u)^n = \int_0^{1/u} n(1 - s)^{n-1} ds$. It follows that the left side of (i) can be written

$$\int_0^n \frac{\alpha(n/s)}{\alpha(n)} \left(1 - \frac{s}{n}\right)^{n-1} ds.$$

By regular variation of α , the integrand tends pointwise to $s^{-\gamma}e^{-s}$, as $n \rightarrow \infty$. By Potter's theorem, the quotient $\alpha(n/s)/\alpha(n)$ is bounded above by a multiple of $(1/s)^{\gamma-\delta} \vee (1/s)^{\gamma+\delta}$, for any given $\delta > 0$, while $(1 - s/n)^{n-1} \leq e^{-s(1-\delta)}$, by the inequality $1 - x \leq e^{-x}$, for $x \in \mathbb{R}$. Therefore, by the dominated convergence theorem the integral converges to $\int_0^\infty s^{-\gamma}e^{-s} ds = \Gamma(1 - \gamma)$.

The series in the left side of (ii) is equal to

$$\sum_{j=1}^{\infty} \sum_{m=2}^n g_{\sigma}(m) \binom{n}{m} p_j^m (1-p_j)^{n-m} = \sum_{m=2}^n g_{\sigma}(m) \binom{n}{m} \int_1^{\infty} \left(\frac{1}{u}\right)^m \left(1-\frac{1}{u}\right)^{n-m} d\alpha(u).$$

Writing $(1/u)^m (1-1/u)^{n-m} = \int_0^{1/u} s^{m-1} (1-s)^{n-m-1} (m-ns) ds$ (for $m \geq 1$) and applying Fubini's theorem, we can rewrite this as

$$\begin{aligned} & \sum_{m=2}^n g_{\sigma}(m) \binom{n}{m} \int_0^1 \alpha\left(\frac{1}{s}\right) s^{m-1} (1-s)^{n-m-1} (m-ns) ds \\ &= \int_0^1 \sum_{l=1}^{n-1} \frac{1}{l-\sigma} \sum_{m=l+1}^n \binom{n}{m} s^{m-1} (1-s)^{n-m-1} (m-ns) \alpha\left(\frac{1}{s}\right) ds \\ &= \int_0^1 \sum_{l=1}^{n-1} \frac{n-l}{l-\sigma} \binom{n}{l} s^l (1-s)^{n-l-1} \alpha\left(\frac{1}{s}\right) ds = \sum_{l=1}^{n-1} \frac{1}{l-\sigma} \mathbb{E} \alpha\left(\frac{1}{S_{l,n}}\right), \end{aligned}$$

for $S_{l,n} \sim \text{Beta}(l+1, n-l)$, where the second last equality follows from Lemma 14. Representing $S_{l,n}$ as $\Gamma_l / (\Gamma_l + \Gamma_{n-l})$, for independent variables $\Gamma_l \sim \Gamma(l+1, 1)$ and $\Gamma_{n-l} \sim \Gamma(n-l, 1)$, we see that the left side of (ii) is equal to

$$\sum_{l=1}^{n-1} \frac{1}{l-\sigma} \mathbb{E} \frac{\alpha(1 + \Gamma_{n-l}/\Gamma_l)}{\alpha(n)} = \sum_{l=1}^{n-1} \frac{1}{l-\sigma} \mathbb{E} \frac{\alpha((n^{-1} + n^{-1}\Gamma_{n-l}/\Gamma_l)n)}{\alpha(n)}.$$

The sequence $U_{l,n} := (n^{-1} + n^{-1}\Gamma_{n-l}/\Gamma_l)$ tends almost surely to $1/\Gamma_l$, by the law of large numbers, as $n \rightarrow \infty$, for fixed l . Since the convergence in (14) is automatically uniform in compacta contained in $(0, \infty)$ (see [9], Theorem B.1.4), it follows that $\alpha(U_{l,n}n)/\alpha(n) \rightarrow (1/\Gamma_l)^\gamma$, almost surely. Furthermore, by Potter's theorem $\alpha(U_{l,n}n)/\alpha(n) \lesssim U_{l,n}^{\gamma+\delta} \vee U_{l,n}^{\gamma-\delta}$, where $U_{l,n}^\beta \leq 1 + (n^{-1}\Gamma_{n-l})^\beta (1/\Gamma_l)^\beta$ is uniformly integrable for every $\beta < 1$, since $n^{-1}\Gamma_{n-l} \rightarrow 1$ in L_1 and $\mathbb{E}(1/\Gamma_l)^\beta < \infty$, so that $n^{-1}\Gamma_{n-l}/\Gamma_l \rightarrow 1/\Gamma_l$ in L_1 , in view of the independence of Γ_{n-l} and Γ_l . By dominated convergence we conclude that $\mathbb{E}\alpha(U_{l,n}n)/\alpha(n) \rightarrow \mathbb{E}(1/\Gamma_l)^\gamma = \Gamma(l+1-\gamma)/l!$. Since $\mathbb{E}U_{l,n}^{\gamma+\delta} \vee U_{l,n}^{\gamma-\delta} \lesssim \mathbb{E}(1/\Gamma_l)^{-\gamma+\delta} \lesssim l^{-\gamma+\delta}$, a second application of the dominated convergence theorem shows that the preceding display tends to $\sum_{l=1}^{\infty} (l-\sigma)^{-1} \Gamma(l+1-\gamma)/l!$.

For the proof of (iii) we write $g_{\sigma}^2(m) = \sum_{k=1}^{m-1} \sum_{l=1}^{m-1} (k-\sigma)^{-1} (l-\sigma)^{-1}$ and follow the same steps as in (ii) to write the left side of (iii) as

$$\sum_{k=1}^{n-1} \sum_{l=1}^{n-1} \frac{1}{k-\sigma} \frac{1}{l-\sigma} \mathbb{E} \frac{\alpha(1/S_{k \vee l, n})}{\alpha(n)}.$$

This is seen to converge to the limit as claimed by the same arguments as under (ii). \square

Lemma 14. For every $p \in [0, 1]$ and $l \in \mathbb{N} \cup \{0\}$ and $n \in \mathbb{N}$,

$$\sum_{m=l+1}^n \binom{n}{m} p^{m-1} (1-p)^{n-m-1} (m-np) = (n-l) \binom{n}{l} p^l (1-p)^{n-l-1}.$$

Proof. For X_{n-1} and X_n the numbers of successes in the first $n-1$ and n independent Bernoulli trials with success probability p , we have $\{X_n \geq l+1\} \subset \{X_{n-1} \geq l\}$ and $\{X_{n-1} \geq l\} - \{X_n \geq l+1\} = \{X_{n-1} = l, B_n = 0\}$, for B_n the outcome of the n th trial. This gives the identity $\Pr(X_{n-1} \geq l) - \Pr(X_n \geq l+1) = \Pr(X_{n-1} = l)(1-p)$. We multiply this by $n/(1-p)$ to obtain the identity given by the lemma, which we first rewrite using that $m \binom{n}{m} = n \binom{n-1}{m-1}$ and $(n-l) \binom{n}{l} = n \binom{n-1}{l}$. \square

Finally consider the situation that P_0 possesses a nontrivial continuous component. In this case the empirical Bayes estimator tends to 1.

Theorem 15. If $P_0 = (1-\lambda)P_0^d + \lambda P_0^c$ where P_0^d is a discrete and P_0^c an atomless probability distribution with $\lambda > 0$ and such that $\alpha_0(u) := \#\{x : 1/P_0\{x\} \leq u\}$ is regularly varying of exponent $\sigma_0 \in (0, 1)$, then $\hat{\sigma}_n \rightarrow 1$ in probability.

Proof. By Lemma 7 the sequence K_n/n tends to λ in probability. The second term in the derivative of the log likelihood (15) depends on tied observations only (through the variables $Z_{n,l}$ with $l \geq 2$), and the arguments from the proof of Theorem 12 show that this term retains the order $O_P(\alpha_0(n))$. Thus it follows that $\Lambda'_n(\sigma)/n \rightarrow \lambda/\sigma$ in probability, whence it is positive with probability tending to one and the likelihood increasing in σ . \square

5. Acknowledgements

The authors would like to thank Botond Szabó for his extensive feedback. This work has been presented several times and the ensuing discussions and remarks by the referees helped identify where the exposition could be improved.

Appendix A: Mean and variance of posterior distribution

In this appendix we derive explicit formulas for the mean and variance of the posterior distribution. The limit of the variances can be seen to be equal to variance of the limit variable in Theorem 1.

Lemma 16. Let $P \sim \text{PY}(\sigma, M, G)$ where $\sigma \geq 0$. Then the mean and variance of the posterior distribution of P based on observations $X_1, \dots, X_n | P \stackrel{iid}{\sim} P$ are as follows

$$\mathbb{E}[Pf | X_1, \dots, X_n] = \sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n + M} f(\tilde{X}_j) + \frac{M + \sigma K_n}{n + M} Gf,$$

$$\begin{aligned} \text{Var}(Pf|X_1, \dots, X_n) &= \left[\sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n - K_n \sigma} f(\tilde{X}_j) - Gf \right]^2 \frac{(n - \sigma K_n)(M + \sigma K_n)}{(n + M)^2(n + M + 1)} \\ &\quad - \frac{\left(\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j) \right)^2}{(n - \sigma K_n)(n + M)(n + M + 1)} + \frac{\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j)^2}{(n + M)(n + M + 1)} \\ &\quad + \frac{(1 - \sigma)(M + \sigma K_n + 1)}{(n + M)(n + M + 1)} \text{Var}_G(f). \end{aligned}$$

Lemma 17. Suppose $X_1, \dots, X_n \stackrel{iid}{\sim} P_0$, where $P_0 = (1 - \lambda)P_0^d + \lambda P_0^c$. If P follows a PY(σ, M, G) process, then the posterior distribution in the model $X_1, \dots, X_n | P \sim P, P_0$ almost surely

$$\begin{aligned} \mathbb{E}[Pf|X_1, \dots, X_n] &\rightarrow (1 - \lambda)P_0^d + (1 - \sigma)\lambda P_0^c + \lambda \sigma G \\ n \text{Var}(Pf|X_1, \dots, X_n) &\rightarrow (1 - \lambda) \text{Var}_{P_0^d}(f) + (1 - \sigma)\lambda \text{Var}_{P_0^c}(f) \\ &\quad + (1 - \sigma)\sigma \lambda \text{Var}_G(f) \\ &\quad + \frac{(1 - \sigma)\lambda(1 - \lambda)}{1 - \sigma\lambda} (P_0^d(f) - P_0^c(f))^2 \\ &\quad + (1 - \sigma\lambda)\sigma \lambda \left(\frac{(1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f)}{1 - \sigma\lambda} - Gf \right)^2. \end{aligned}$$

Proof of Lemma 16. We begin by recalling the posterior distribution from Section 4. Note that we have the following results:

- $\mathbb{E}[R_n] = \frac{n - K_n \sigma}{n + M}$ and $\text{Var}(R_n) = \frac{(n - K_n \sigma)(M + K_n \sigma)}{(n + M)^2(n + M + 1)}$.
- $\mathbb{E}[Q_n(f)] = G(f)$, $\text{Var}(Q_n(f)) = \frac{1 - \sigma}{M + \sigma K_n} \text{Var}_G(f)$.

The first two results are standard results for Beta distributed random variables, and the last two results are because Q_n is a Pitman-Yor process. Now we just need to compute the moments for the weights W_j . We use the following results from the Dirichlet distribution. If $\tilde{X} \sim \text{Dir}(K_n, \alpha_1, \dots, \alpha_{K_n})$, then

$$\begin{aligned} \mathbb{E}[\tilde{X}_i] &= \frac{\alpha_i}{\sum_{k=1}^{K_n} \alpha_k}, \\ \text{Var}(\tilde{X}_i) &= \frac{\alpha_i(\sum_{k=1}^{K_n} \alpha_k - \alpha_i)}{(\sum_{k=1}^{K_n} \alpha_k)^2(1 + \sum_{k=1}^{K_n} \alpha_k)}, \end{aligned}$$

and

$$\text{Cov}(\tilde{X}_i, \tilde{X}_j) = \frac{-\alpha_i \alpha_j}{(\sum_{k=1}^{K_n} \alpha_k)^2(1 + \sum_{k=1}^{K_n} \alpha_k)}.$$

In our case $\alpha_i = N_{i,n} - \sigma$, $K = K_n$ and $\sum_{k=1}^{K_n} \alpha_k = n - \sigma K_n$. Then a direct computation shows that

$$\mathbb{E}\left[\sum_{j=1}^{K_n} W_j f(\tilde{X}_j)\right] = \sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n - K_n \sigma} f(\tilde{X}_j).$$

For the variance we use that, for independent random variables, the variance of the sum is the sum of the covariances.

$$\begin{aligned}
& \text{Var}\left(\sum_{i=1}^{K_n} W_j f(\tilde{X}_j) \mid X_1, \dots, X_n\right) \\
&= \sum_{i \neq j} \text{Cov}(W_i, W_j) f(\tilde{X}_i) f(\tilde{X}_j) + \sum_{i=1}^{K_n} \text{Var}(W_i) f(\tilde{X}_i)^2 \\
&= \sum_{i \neq j} \frac{-(N_{i,n} - \sigma)(N_{j,n} - \sigma)}{(n - \sigma K_n)^2 (n - \sigma K_n + 1)} f(\tilde{X}_i) f(\tilde{X}_j) \\
&\quad + \sum_{i=1}^{K_n} \frac{(N_{i,n} - \sigma)(n - \sigma K_n - N_{i,n} + \sigma)}{(n - \sigma K_n)^2 (n - \sigma K_n + 1)} f(\tilde{X}_i)^2 \\
&= -\frac{\left(\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j)\right)^2}{(n - \sigma K_n)^2 (n - \sigma K_n + 1)} + \frac{\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j)^2}{(n - \sigma K_n)(n - \sigma K_n + 1)}.
\end{aligned}$$

Now we can compute the mean and variance. Using independence between R_n , W and Q_n and linearity we see that

$$\mathbb{E}[P(f) \mid X_1, \dots, X_n] = \sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n + M} f(\tilde{X}_j) + \frac{M + \sigma K_n}{n + M} G(f).$$

In order to compute the variance we apply the law of total variance. For any two random variables X, Y with finite second moment we have that

$$\text{Var}(X) = \mathbb{E}[\text{Var}(X \mid Y)] + \text{Var}(\mathbb{E}[X \mid Y]).$$

We split into conditioning on R_n and the rest, so we can use the independence between W and Q_n . We compute these piece by piece. First consider

First consider

$$\mathbb{E}\left[\text{Var}\left(R_n \sum_{j=1}^{K_n} W_j f(\tilde{X}_j) + (1 - R_n)Q_n(f) \mid R_n\right)\right].$$

Due to the independence of W and Q_n given R_n

$$= \mathbb{E}\left[R_n^2 \text{Var}\left(\sum_{j=1}^{K_n} W_j f(\tilde{X}_j)\right) + (1 - R_n)^2 \text{Var}(Q_n(f))\right].$$

Simplifying the expression yields

$$= \mathbb{E}[R_n^2] \text{Var}\left(\sum_{j=1}^{K_n} W_j f(\tilde{X}_j)\right) + \mathbb{E}[(1 - R_n)^2] \text{Var}(Q_n(f)).$$

Filling in the known moments results in

$$= \frac{(n - \sigma K_n)(n + 1 - \sigma K_n)}{(n + M)(n + M + 1)} \text{Var} \left(\sum_{j=1}^{K_n} W_j f(\tilde{X}_j) \right) + \frac{(M + \sigma K_n)(M + \sigma K_n + 1)}{(n + M)(n + M + 1)} \text{Var} (Q_n(f)).$$

Expanding the variance terms and simplifying gives

$$= - \frac{\left(\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j) \right)^2}{(n - \sigma K_n)(n + M)(n + M + 1)} + \frac{\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j)^2}{(n + M)(n + M + 1)} + \frac{(1 - \sigma)(M + \sigma K_n + 1)}{(n + M)(n + M + 1)} \text{Var}_G(f).$$

Next we deal with

$$\text{Var} \left(\mathbb{E} \left[R_n \sum_{j=1}^{K_n} W_j f(\tilde{X}_j) + (1 - R_n) Q_n(f) \mid R_n \right] \right).$$

Computing the expected value gives

$$= \text{Var} \left(R_n \sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n - K_n \sigma} f(\tilde{X}_j) + (1 - R_n) G(f) \right).$$

Reorganising terms

$$= \text{Var} \left(G(f) + R_n \left(\sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n - K_n \sigma} f(\tilde{X}_j) - G(f) \right) \right).$$

The constant term does not contribute to the variance so can be, and then taking the square of the constant in front of R_n results in

$$= \left(\sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n - K_n \sigma} f(\tilde{X}_j) - G(f) \right)^2 \text{Var} (R_n).$$

Computing the variance of R_n gives

$$= \left(\sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n - K_n \sigma} f(\tilde{X}_j) - G(f) \right)^2 \frac{(n - \sigma K_n)(M + \sigma K_n)}{(n + M)^2(n + M + 1)}.$$

Therefore by the law of total variance we find the result. □

Proof of Lemma 17. We begin with some basic results which we will apply in several places. We note the following two almost sure limits: $\frac{K_n}{n} \rightarrow \lambda$ P_0 -almost surely and

$$\frac{\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j)}{n} \rightarrow (1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f) \quad P_0\text{-a.s.}$$

For the posterior mean we know the exact formula by 16 and therefore the following limit can be computed:

$$\begin{aligned} \mathbb{E}[P(f)|X_1, \dots, X_n] &= \sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n + M} f(\tilde{X}_j) + \frac{M + \sigma K_n}{n + M} G(f) \\ &\rightarrow (1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f) + \lambda \sigma G(f) P_0\text{-a.s.} \end{aligned}$$

Recall from 16 the formula for the posterior variance. We analyse this term by term. They all follow directly from the remarks at the beginning of the the proof, and the limits hold P_0 -almost surely.

First we find that

$$\sum_{j=1}^{K_n} \frac{N_{j,n} - \sigma}{n - K_n \sigma} f(\tilde{X}_j) \rightarrow \frac{(1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f)}{1 - \sigma \lambda}.$$

Secondly,

$$n \frac{(n - \sigma K_n)(M + \sigma K_n)}{(n + M)^2(n + M + 1)} \rightarrow (1 - \sigma \lambda) \sigma \lambda.$$

Next,

$$-n \frac{\left(\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j)\right)^2}{(n - \sigma K_n)(n + M)(n + M + 1)} \rightarrow -\frac{\left((1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f)\right)^2}{1 - \sigma \lambda}.$$

Also,

$$n \frac{\sum_{j=1}^{K_n} (N_{j,n} - \sigma) f(\tilde{X}_j)^2}{(n + M)(n + M + 1)} \rightarrow (1 - \lambda)P_0^d(f^2) + (1 - \sigma)\lambda P_0^c(f^2).$$

And finally,

$$n \frac{(1 - \sigma)(M + \sigma K_n + 1)}{(n + M)(n + M + 1)} \text{Var}_G(f) \rightarrow (1 - \sigma) \sigma \lambda \text{Var}_G(f).$$

This means we now have computed the limit of the posterior variance. We will now add all the terms together, and by the continuous mapping theorem we find that,

$$\begin{aligned} n\text{Var}(P(f)|X_1, \dots, X_n) &\rightarrow (1 - \sigma\lambda)\sigma\lambda \\ &\quad \left(\frac{(1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f)}{1 - \sigma\lambda} - G(f) \right)^2 \\ &\quad - \frac{((1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f))^2}{1 - \sigma\lambda} \\ &\quad + (1 - \lambda)P_0^d(f^2) + (1 - \sigma)\lambda P_0^c(f^2) \\ &\quad + (1 - \sigma)\sigma\lambda \text{Var}_G(f) \quad \text{a.s. } P_0. \end{aligned}$$

Note that

$$\begin{aligned} & - \frac{((1 - \lambda)P_0^d(f) + (1 - \sigma)\lambda P_0^c(f))^2}{1 - \sigma\lambda} \\ & + (1 - \lambda)P_0^d(f^2) + (1 - \sigma)\lambda P_0^c(f^2) \\ & = (1 - \lambda)\text{Var}_{P_0^d}(f) + (1 - \sigma)\lambda \text{Var}_{P_0^c}(f) \\ & + \frac{(1 - \sigma)\lambda(1 - \lambda)}{1 - \sigma\lambda} (P_0^d(f) - P_0^c(f))^2. \end{aligned}$$

Combining everything yields the Lemma. □

References

- [1] ARBEL, J., DE BLASI, P., AND PRÜNSTER, I. Stochastic approximations to the Pitman-Yor process. *Bayesian Analysis* (2018). Advance publication. [MR4136558](#)
- [2] BINGHAM, N. H., GOLDIE, C. M., AND TEUGELS, J. L. *Regular variation*, vol. 27 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1989. [MR0898871](#)
- [3] CAMERLENGHI, F., DUNSON, D. B., LIJOI, A., PRÜNSTER, I., AND RODRÍGUEZ, A. Latent nested nonparametric priors (with discussion). *Bayesian Anal.* 14, 4 (2019), 1303–1356. With discussions and a rejoinder. [MR4044854](#)
- [4] CAMERLENGHI, F., LIJOI, A., ORBANZ, P., AND PRÜNSTER, I. Distribution theory for hierarchical processes. *Ann. Statist.* 47, 1 (2019), 67–92. [MR3909927](#)
- [5] CEREDA, G. *Current challenges in statistical DNA evidence evaluation*. PhD thesis, Leiden University, 2017.
- [6] CEREDA, G., AND GILL, R. D. A nonparametric bayesian approach to the rare type match problem, 2019. [MR4217515](#)

- [7] DE BLASI, P., LIJOI, A., AND PRÜNSTER, I. An asymptotic analysis of a class of discrete nonparametric priors. *Statistica Sinica* 23, 3 (2013), 1299–1321. [MR3114715](#)
- [8] DE BLASI ET AL. Are Gibbs-type priors the most natural generalization of the Dirichlet process? *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37, 2 (2015), 212–229.
- [9] DE HAAN, L., AND FERREIRA, A. *Extreme value theory*. Springer Series in Operations Research and Financial Engineering. Springer, New York, 2006. An introduction. [MR2234156](#)
- [10] DONSKER, M. D. Justification and extension of Doob’s heuristic approach to the Komogorov-Smirnov theorems. *Ann. Math. Statistics* 23 (1952), 277–281. [MR0047288](#)
- [11] FAVARO, S., AND NAULET, Z. Near-optimal estimation of the unseen under regularly varying tail populations, 2021.
- [12] FERGUSON, T. Prior distributions on spaces of probability measures. *Ann. Statist.* 2 (1974), 615–629. [MR0438568](#)
- [13] GHOSAL, S., AND VAN DER VAART, A. *Fundamentals of Nonparametric Bayesian Inference*. Cambridge University Press, 2017. [MR3587782](#)
- [14] GOLDWATER, S., GRIFFITHS, T. L., AND JOHNSON, M. Interpolating between types and tokens by estimating power-law generators. In *Advances in neural information processing systems* (2005).
- [15] ISHWARAN, H., AND JAMES, L. F. Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association* 96, 453 (2001), 161–173. [MR1952729](#)
- [16] JAMES, L. Large sample asymptotics for the two-parameter Poisson-Dirichlet process. *Pushing the limits of contemporary Statistics: Contributions in Honor of Jayanta k. Ghosh* 3 (2008). [MR2459225](#)
- [17] KARLIN, S. Central limit theorems for certain infinite urn schemes. *Journal of Mathematics and Mechanics* 17, 4 (1967). [MR0216548](#)
- [18] LO, A. Weak convergence for Dirichlet processes. *Sankhyā Ser. A* 45, 1 (1983), 105–111. [MR0749358](#)
- [19] LO, A. A remark on the limiting posterior distribution of the multiparameter Dirichlet process. *Sankhyā Ser. A* 48, 2 (1986), 247–249. [MR0905464](#)
- [20] PERMAN, M., PITMAN, J., AND YOR, M. Size-biased sampling of Poisson point processes and excursions. *Probab. Theory Related Fields* 92, 1 (1992), 21–39. [MR1156448](#)
- [21] PITMAN, J. Random discrete distributions invariant under size-biased permutation. *Adv. in Appl. Probab.* 28, 2 (1996), 525–539. [MR1387889](#)
- [22] PITMAN, J. Some developments of the Blackwell-MacQueen urn scheme. *Institute of Mathematical Statistics Lecture Notes - Monograph Series* 30 (1996), 245–267. [MR1481784](#)
- [23] PITMAN, J. Poisson-Kingman partitions. In *Statistics and Science: a Festschrift for Terry Speed*, vol. 40 of *IMS Lecture Notes Monogr. Ser.* Inst. Math. Statist., Beachwood, OH, 2003, pp. 1–34. [MR2004330](#)
- [24] PITMAN, J., AND YOR, M. The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann. Probab.* 25, 2 (1997), 855–900.

[MR1434129](#)

- [25] POLLARD, D. *Convergence of Stochastic Processes*. Springer Series in Statistics. Springer-Verlag, New York, 1984. [MR0762984](#)
- [26] TEH, Y. W. A hierarchical Bayesian language model based on Pitman-Yor processes. In *ACL-44 Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics* (2006), pp. 985–992.
- [27] VAN DER VAART, A. *Asymptotic Statistics*, vol. 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998. [MR1652247](#)
- [28] VAN DER VAART, A., AND WELLNER, J. *Weak convergence and Empirical processes*. Springer-Verlag, 1996. [MR1385671](#)
- [29] VAN DER VAART, A., AND WELLNER, J. A. Preservation theorems for Glivenko-Cantelli and uniform Glivenko-Cantelli classes. In *High dimensional probability, II (Seattle, WA, 1999)*, vol. 47 of *Progr. Probab.* Birkhäuser Boston, Boston, MA, 2000, pp. 115–133. [MR1857319](#)
- [30] WOOD, F., ARCHAMBEAU, C., GASTHAUS, J., JAMES, L., AND TEH, Y. W. A stochastic memoizer for sequence data. In *Proceedings of the 26th Annual International Conference on Machine Learning* (New York, NY, USA, 2009), ICML '09, Association for Computing Machinery, p. 1129–1136.