



HAL
open science

Near-optimal estimation of the unseen under regularly varying tail populations

Stefano Favaro, Zacharie Naulet

► **To cite this version:**

Stefano Favaro, Zacharie Naulet. Near-optimal estimation of the unseen under regularly varying tail populations. *Bernoulli*, 2023, 29 (4), 10.3150/23-BEJ1589 . hal-04535795

HAL Id: hal-04535795

<https://hal.science/hal-04535795>

Submitted on 7 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Near-optimal estimation of the unseen under regularly varying tail populations

STEFANO FAVARO^{1,*} and ZACHARIE NAULET²

¹*Department of Economics and Statistics, University of Torino and Collegio Carlo Alberto, Torino, 10134, Italy.
E-mail: stefano.favaro@unito.it*

²*Laboratoire de mathématiques d'Orsay, CNRS, Université Paris-Saclay, 91405 Orsay, France.
E-mail: zacharie.naulet@universite-paris-saclay.fr*

Given n samples from a population of individuals belonging to different species, what is the number U of hitherto unseen species that would be observed if λn new samples were collected? This is the celebrated unseen-species problem, of interest in many disciplines, and it has been the subject of recent breakthrough studies introducing non-parametric estimators of U that are minimax near-optimal and consistent all the way up to $\lambda \asymp \log n$. These works do not rely on any assumption on the underlying unknown distribution p of the population, and therefore, while providing a theory in its greatest generality, worst-case distributions may hamper the estimation of U in concrete applications. In this paper, we strengthen the non-parametric framework for estimating U , making use of suitable assumptions on the underlying distribution p . In particular, inspired by the estimation of rare probabilities in extreme value theory, and motivated by the ubiquitous power-law type distributions in many natural and social phenomena, we make use of a semi-parametric assumption of regular variation of index $\alpha \in (0, 1)$ for the tail behaviour of p . Under this assumption, we introduce an estimator of U that is simple, linear in the sampling information, computationally efficient, and scalable to massive datasets. Then, uniformly over our class of regularly varying tail distributions, we show that the proposed estimator has provable guarantees: i) it is minimax near-optimal, up to a power of $\log n$ factor; ii) it is consistent all the way up to $\log \lambda \asymp n^{\alpha/2}/\sqrt{\log n}$, and this range is the best possible. This work presents the first study on the estimation of the unseen under regularly varying tail distributions. Our results rely on a novel approach, of independent interest, which combines the renowned method of the two fuzzy hypotheses for minimax estimation of discrete functionals, with Bayesian arguments under Poisson-Kingman priors for p . A numerical illustration of our methodology is presented for synthetic and real data.

MSC2020 subject classifications: Primary 62G05 ; Secondary 62C20

Keywords: Multinomial model; optimal minimax estimation; Poisson-Kingman prior; power-law data; regularly varying tails; tail-index; unseen-species problem

1. Introduction

Estimating the number of unseen species is an important problem in many disciplines. It first appeared in ecology (Fisher et al. [22], Good and Toulmin [27], Chao and Lee [12], Bunge and Fitzpatrick [9]), and its importance has grown in recent years driven by applications in biological sciences (Kroes et al. [40], Gao et al. [24], Ionita-Laza et al. [34], Daley and Smith [15]). Consider a population of individuals, with each individual being endowed with a value from a set of species labels $\mathbb{S} = \{s_j\}_{j \geq 1}$, and for $n, m \geq 1$ consider $(n + m)$ samples from the population to be modeled as a random sample $\mathbf{X}_{n+m} = (X_1, \dots, X_n, X_{n+1}, \dots, X_{n+m})$ from an unknown distribution $p = \sum_{j \geq 1} p_j \delta_{s_j}$ on \mathbb{S} . Assuming that only the first n elements of \mathbf{X}_{n+m} are observable, the unseen-species problem calls for

*Also affiliated to IMATI-CNR "Enrico Magenes" (Milan, Italy).

estimating

$$U_{n,m} := U_{n,m}(\mathbf{X}_{n+m}) = |\{X_{n+1}, \dots, X_{n+m}\} \setminus \{X_1, \dots, X_n\}|,$$

namely the number of hitherto unseen species that would be observed in m additional (unobservable) sample. For $\lambda = m/n > 0$, the Good-Toulmin estimator (Good and Toulmin [27]) and the smoothed Efron-Thisted estimator (Efron and Thisted [19]) are arguably the most popular estimators of $U_{n,\lambda n}$. Motivated by the increasing interest in the range $\lambda > 1$, especially in biological applications, the smoothed Efron-Thisted estimator has been the subject of recent breakthrough studies (Orlitsky et al. [46], Wu and Yang [56, 57], Polyanskiy and Wu [49]). In particular, under the Multinomial model for the observable samples, Orlitsky et al. [46] showed that a smoothed Efron-Thisted estimator is such that: i) its worst-case normalized mean-square error (NMSE) is minimax near-optimal for all $\lambda > 1$; ii) it consistently estimates $U_{n,m}$ all of the way up to $\lambda \asymp \log n$, and this range is the best possible. These results do not rely on any assumption on the underlying distribution p , and they hold uniformly over the class \mathcal{P} of distributions on \mathbb{S} , thus providing a grounded theory in its greatest generality. Under this non-parametric framework, however, worst-case distributions may hamper both the estimation of $U_{n,\lambda n}$ and the study of theoretical guarantees, leading to unreliable results in numerous concrete applications.

1.1. Our contributions

In this paper, we strengthen the non-parametric framework for estimating $U_{n,\lambda n}$, making use of suitable assumptions on the underlying distribution p . For $p \in \mathcal{P}$, we define the function $\bar{F}_p : (0, 1) \rightarrow \mathbb{Z}_+$ such that

$$\bar{F}_p(x) := \sum_{j \geq 1} \mathbf{1}_{\{p_j > x\}}, \quad x \in (0, 1),$$

which characterizes the distribution of $U_{n,\lambda n}$, and we argue that reasonable assumptions on the distribution p can be made by prescribing the behaviour of $\bar{F}_p(x)$ near $x \approx 0$. We denote by K_n the number of distinct species in a random sample (X_1, \dots, X_n) from p , such that $U_{n,m} = K_{n+m} - K_n$. Since $\mathbb{E}_p[K_{n+m}] = (n+m) \int_{(0,1)} \bar{F}_p(1-x) x^{n+m-1} dx$, that is $\mathbb{E}_p[K_{n+m}]/(n+m)$ is the Mellin transform of $y \mapsto \bar{F}_p(1-y)$, the estimation of K_{n+m} is dual to the estimation of $\bar{F}_p(x)$ for small values of x . Therefore, the problems of estimating $U_{n,m}$, K_{n+m} and the Mellin transform of $y \mapsto \bar{F}_p(1-y)$ at large arguments, or characterizing $\bar{F}_p(z)$ for small z , are essentially equivalent problems. This observation allows to draw a parallel with extreme value theory (EVT) (Bingham et al. [7], De Haan and Ferreira [17]), where interest is in estimating probabilities of rare events; that is, given a real-valued random sample (Z_1, \dots, Z_n) , interest is in the estimation of the survival function $z \mapsto \mathbb{P}(Z_1 > z)$ for large values of z , especially for z larger than $\max\{Z_1, \dots, Z_n\}$. We see that \bar{F}_p in the unseen-species problem plays an analogous role to that of the survival function in EVT, and therefore it is not surprising that the estimation of $U_{n,m}$ suffers from the same issues encountered in the estimation of probabilities of rare events. In both cases the interest is in estimating \bar{F}_p outside the range allowed by available data, which, unless λ is small or a weak loss functions is considered, is impossible because of the arbitrary behaviour of \bar{F}_p nearly zero. This precludes generic non-parametric assumptions in many realistic scenarios.

Inspired by the well-known problem of estimating probabilities of rare events in the context of EVT, we consider a semi-parametric assumption on \bar{F}_p prescribing the behaviour of $\bar{F}_p(x)$ near $x \approx 0$. In particular, for constants $L, L' > 0$ and $\alpha \in (0, 1)$, we consider the class of regularly varying

distributions

$$\mathcal{P}(\alpha, L, L') := \left\{ p \in \mathcal{P} : \sup_{x \in (0,1)} \frac{|x^\alpha \bar{F}_p(x) - L|}{\sqrt{x^\alpha \log(e/x)}} \leq L' \right\},$$

i.e. distributions with regularly varying tails of index α (Feller [20], Bingham et al. [7]). We note that not all distributions with regularly varying tails of index α permit a reliable estimation of the unseen for large λ . This is the reason to impose, in the definition of $\mathcal{P}(\alpha, L, L')$, that the error rate of approximating $\bar{F}_p(x)$ by $Lx^{-\alpha}$ must decay rapidly enough when $x \rightarrow 0$. Such a constraint is common in EVT, for instance to quantify the rate of convergence of Hill's estimator (Hill [32], Hall and Welsh [28, 29], Drees [18]). A distribution p satisfying $\lim_{x \rightarrow 0} x^\alpha \bar{F}_p(x) \sim L$ for some $(\alpha, L) \in (0, 1) \times (0, \infty)$ is referred to as a power-law distributions, and data from such a distribution are typically referred to as power-law data. As in EVT, for power-law type distributions the value of $\mathbb{E}_p[K_{n+m}]$ can be deduced from α via Tauberian arguments (Feller [20]). Gnedin et al. [26] show that $\lim_{n \rightarrow +\infty} n^{-\alpha} K_n = L\Gamma(1 - \alpha)$ almost surely, with $\Gamma(\cdot)$ being the Gamma function. The power-law behaviour is a common assumption in EVT, because it entails that the law of $\max\{Z_1, \dots, Z_n\}$ properly centered and rescaled converge to a Fréchet distribution. In the unseen-species problem, interest in the power-law behaviour is two-fold: first it enables for tractable analysis and accurate estimation of $U_{n, \lambda n}$, and second, power-law data have readily identifiable signatures which permits to assess the plausibility of the assumption on real data. See Appendix F.2 for details on checking the power-law assumption.

Remark 1. *One could argue that since the behaviour of $U_{n,m}$ for large m is determined by $\bar{F}_p(x)$ for small x , the class $\mathcal{P}(\alpha, L, L')$ is too restrictive, and we may be tempted to use instead the weaker class*

$$\tilde{\mathcal{P}}(\alpha, L, L', x_0) := \left\{ p \in \mathcal{P} : \sup_{x \in (0, x_0)} \frac{|x^\alpha \bar{F}_p(x) - L|}{\sqrt{x^\alpha \log(e/x)}} \leq L' \right\}.$$

For any $x_0 \in (0, 1)$ we have $\mathcal{P}(\alpha, L, L') \subset \tilde{\mathcal{P}}(\alpha, L, L', x_0)$ and for any $p \in \mathcal{P}$, we also have for $x \geq x_0$ that $\bar{F}_p(x) \leq \bar{F}_p(x_0) = \sum_{j \geq 1} \mathbf{1}_{p_j > x_0} \leq x_0^{-1} \sum_{j \geq 1} p_j = x_0^{-1} < \infty$. This shows that for any $p \in \mathcal{P}$ and any (α, L) , it must be that $\sup_{x \in (x_0, 1)} |x^\alpha \bar{F}_p(x) - L| (x^\alpha \log(e/x))^{-1/2} < \infty$. That is, for all (α, L, L', x_0) there is a $L'' > 0$ such that $\tilde{\mathcal{P}}(\alpha, L, L', x_0) \subset \mathcal{P}(\alpha, L, L'')$. Using the weaker class $\tilde{\mathcal{P}}(\alpha, L, L', x_0)$ in place of the class $\mathcal{P}(\alpha, L, L')$ is interesting only if the dependency in (L', x_0) can be made explicit in the rates of convergence of the estimators. Unfortunately, this is known to be a tough problem, and then we make the choice to stick with $\mathcal{P}(\alpha, L, L')$ in order to facilitate the analysis.

The estimation of $U_{n, \lambda n}$ over $\mathcal{P}(\alpha, L, L')$ is directly connected to the problem of estimating the tail-index $\alpha \in (0, 1)$, and hence we construct estimators for both α and $U_{n, \lambda n}$. We assume the Multinomial model for the observable samples and, as a practical tool, we rely on Poisson-Kingman priors (Pitman [47, 48]) for the underlying unknown $p \in \mathcal{P}(\alpha, L, L')$ to obtain estimators of α and $U_{n, \lambda n}$ and to establish minimax lower bounds under suitable loss functions. The proposed estimators $\hat{\alpha}_n$ and $\hat{U}_{n, \lambda n}(\hat{\alpha}_n)$ of α and $U_{n, \lambda n}$, respectively, are simple, linear in the sampling information, computationally efficient and scalable to massive datasets. In particular, uniformly over $\mathcal{P}(\alpha, L, L')$, we show that: i) under a quadratic loss, $\hat{\alpha}_n$ consistently estimates α at rate $n^{-\alpha/2} \sqrt{\log n}$; ii) under a quadratic loss normalized by $(n\lambda)^\alpha$, the estimator $\hat{U}_{n, \lambda n}(\hat{\alpha}_n)$ consistently estimates $U_{n, \lambda n}$ all of the way up to $\log \lambda \asymp n^{\alpha/2} / \sqrt{\log n}$. We also establish lower bounds for the minimax risks of estimating α and $U_{n, \lambda n}$. In particular, uniformly over $\mathcal{P}(\alpha, L, L')$, we show that: i) the estimators $\hat{\alpha}_n$ and $\hat{U}_{n, \lambda n}(\hat{\alpha}_n)$

are near-optimal, in the sense of matching minimax lower bounds, up to a power of $\log n$ factor; ii) the range $\log \lambda \asymp n^{\alpha/2}/\sqrt{\log n}$ is the best possible for consistently estimating $U_{n,\lambda n}$. This work presents the first study on the estimation of the unseen under regularly varying tail distributions. Our results rely on a novel approach, of independent interest, that combines the renowned method of the two fuzzy hypotheses for minimax estimation of discrete functionals (Tsybakov [53], Wu and Yang [56, 57]), with Bayesian arguments under Poisson-Kingman priors for the underlying $p \in \mathcal{P}(\alpha, L, L')$; interestingly, our approach is critical both for deriving an estimator of the tail-index α and for establishing minimax lower bounds. An illustration of our methodology is presented for synthetic data, text data arising from Gutenberg books and Wikipedia pages, and data from humans' electronic activities in email communications and Twitter posts. All these data have power-law behaviours, and our empirical analysis emphasizes the substantial gain in estimation, over smoothed Efron-Thisted estimators, allowed by leveraging such a power-law behaviour. This is of course expected, and we make clear that such a results do not mean that our estimator is preferable in general. The good empirical performance of our estimator on these datasets is due to their power-law behaviour, on which our estimator is designed to work at best.

Power-law data occur in many situations of scientific interest, and nowadays they have significant consequences for the understanding of natural and social phenomena. See Clauset et al. [14], and references therein, for a detailed account on power-law data. Besides the context of text data in natural languages (Zipf [58], Cancho and Solé [11], Harald [31]), power-law phenomena have emerged for data arising from humans' electronic activities, e.g. patterns of website visits, email messages, relations and interactions on social networks, password innovation, tags in annotation systems and edits of webpages (Huberman and Adamic [33], Barabási [5], Rybski [50], Muchnik et al. [44]). In these contexts, the estimation of $U_{n,m}$ is critical in decisions concerning with, e.g., collective attention monitoring, resources managements, novelty and popularity triggering, and language learning. For power-law data, our study shows a remarkable gain in the estimation accuracy of $U_{n,\lambda n}$ in the large λ regime if we leverage this prior information on the structure of \bar{F}_p . Of course, as in the context of EVT, estimates have to be considered with precautions, since they rely on a modeling hypothesis that can be easily rejected if not satisfied, but it is impossible to accept with absolute certainty on the sole basis of the observed data (at best we can be satisfied with plausible evidences). Thus, validating the result of the estimation shall always be made based upon external information about the data generating mechanism, such as biological or physical insights, that can confirm it has an actual power-law behaviour (Clauset et al. [14]).

Remark 2. *In general, the problem of estimating the unseen in the large λ regime do require to make modeling assumptions on the underlying p . In this paper, we claim that assuming $\bar{F}_p(x)$ behaves as a polynomial near zero is a convenient assumption, often realistic in many problems of practical interest, and thus is worth investigating. Other forms of modeling assumptions can be made, such as constraining the shape of the distribution. In this direction, Chee and Wang [13], Giguelay and Huet [25] and Balabdaoui and Kulagina [3] investigate the same problem under the modeling assumption that p is k -monotone. Although not directly related to the unseen-species problem, the approaches developed in Anevski et al. [1] and Jana et al. [35] may be adapted to construct estimators for the unseen.*

1.2. Organization of the paper

The paper is structured as follows. In Section 2 we recall the Multinomial and the Poisson-Kingman partition models, and we highlight the major challenges in deriving lower bounds for the minimax risk

of estimating $U_{n,\lambda n}$. Section 3 contains our results: i) we introduce an estimator of α and a plug-in estimator of $U_{n,\lambda n}$, and we establish their convergence rates; ii) we study optimality of the proposed estimators by establishing lower bounds for the minimax risks of estimating α and $U_{n,\lambda n}$. In Section 4 we present a numerical illustrations of our method on synthetic and real data, and Section 5 contains a discussion of our work, including related estimation problems, and open challenges. Proofs are deferred to appendices.

2. Partition models for $U_{n,\lambda n}$ and preliminaries

The Multinomial model for the observable samples is the most common model for estimating $U_{n,\lambda n}$ (Good and Toulmin [27], Efron and Thisted [19], Orlicsky et al. [46]). Let $p = \sum_{j \geq 1} p_j \delta_{s_j}$ be an unknown distribution on the set of species labels \mathbb{S} , that is $p \in \mathcal{P}$. Under the Multinomial model it is assumed that $n \geq 1$ observable samples are modeled as a random sample $\mathbf{X}_n = (X_1, \dots, X_n)$ from p , i.e.,

$$X_i \stackrel{\text{iid}}{\sim} p \quad i = 1, \dots, n.$$

That is, the random variable X_i takes value $s_j \in \mathbb{S}$ with unknown probability p_j , for $j \geq 1$. Species labels s_j 's are not relevant in our context, and therefore they are assumed to be fixed. Because of the discreteness of p , the random sample \mathbf{X}_n from p induces a random partition Π_n of the set $\mathbb{N}_n = \{1, \dots, n\}$ whose blocks are the (equivalence) classes induced by the equivalence relations $i \sim j \iff X_i = X_j$ almost surely. Similarly, for $m \geq 1$ we denote by Π_{n+m} the random partition of the set $\mathbb{N}_{n+m} = \{1, \dots, n, n+1, \dots, n+m\}$ induced by the random sample $\mathbf{X}_{n+m} = (X_1, \dots, X_n, X_{n+1}, \dots, X_{n+m})$ from p . According to its definition, $U_{n,\lambda n}$ is uniquely determined by the random partition Π_{n+m} . We assume that the random sample \mathbf{X}_n can be ideally extended to a sequence $\mathbf{X} = (X_i)_{i \geq 1}$, of which the first n elements \mathbf{X}_n are observable, and we denote by $\Pi = (\Pi_n)_{n \geq 1}$ the random partition of \mathbb{N} induced by the sequence \mathbf{X} . The conditional distribution of Π given \mathbf{X}_n coincides with the conditional distribution of Π given Π_n , and hence it is sufficient to consider a model for Π_n rather than for \mathbf{X}_n . The next definition summarizes the model for the random partition Π_n under \mathcal{P} .

DEFINITION 1 (Multinomial partition model). *Let \mathbf{X}_n be a random sample from $p \in \mathcal{P}$, and let \mathbb{P}_p^n be the distribution of the random partition Π_n of \mathbb{N}_n induced by \mathbf{X}_n . Then the Multinomial partition model for Π_n is $\{\mathbb{P}_p^n : p \in \mathcal{P}\}$, and we denote by \mathbb{E}_p^n the expectation under \mathbb{P}_p^n . Similarly, we denote by $\bar{\mathbb{P}}_p$ and $\bar{\mathbb{E}}_p$ the distribution of the random partition Π and the corresponding expectation under $\bar{\mathbb{P}}_p$, respectively.*

Under the model of Definition 1, the derivation of sharp lower bounds for the minimax risk of an estimator of $U_{n,\lambda n}$ is challenging, the main difficulty being that \mathbf{X}_n contains only few information of $U_{n,\lambda n}$ (Orlicsky et al. [46], Wu and Yang [57]); this is a common feature of discrete functionals of Π_{n+m} , of which the unseen $U_{n,\lambda n}$ is an instance. Although the estimation of $U_{n,\lambda n}$ will be carried under sub-models of the Multinomial partition model, i.e. over $p \in \mathcal{P}(\alpha, L, L')$, we will also make use of a partition model where the distribution p is itself random. Such a model will serve as a practical tool to obtain an estimator of $U_{n,\lambda n}$, and then to establish a lower bound for its minimax risk. In particular, we assume that p is a discrete (almost surely) random probability measure in the broad class of Poisson-Kingman models (Pitman [47]), namely discrete random probability measures obtained by a suitable normalization of Poisson point processes. To define a Poisson-Kingman model, let $(J_j)_{j \geq 1}$

be the decreasing ordered jumps of an inhomogeneous Poisson point process on \mathbb{R}_+ with Lévy measure ρ such that

$$\int_{\mathbb{R}_+} \min(1, x) \rho(dx) < +\infty. \quad (1)$$

Under the assumption (1) it holds that $T := \sum_{j \geq 1} J_j < +\infty$, i.e. finite total mass, almost surely (Kingman [39]). A Poisson-Kingman distribution with parameter ρ is defined as the law of the random probability masses $(J_j/T)_{j \geq 1}$, or equivalently the law of the discrete (almost surely) random probability measure

$$P := \sum_{j \geq 1} \frac{J_j}{T} \delta_{s_j}.$$

We denote by \mathcal{L} the class of Lévy measures satisfying the assumption (1), and we shall write $P \sim \text{PK}(\rho)$ for abbreviating the Poisson-Kingman distribution with Lévy measure $\rho \in \mathcal{L}$ (Pitman [48, Chapter 3])

According to the celebrated de Finetti's representation theorem for exchangeable random variables, a random sample from a Poisson-Kingman model P is part of an exchangeable sequence $\mathbf{X} = (X_i)_{i \geq 1}$ with directing (de Finetti) measure $\text{PK}(\rho)$, that is $\lim_{n \rightarrow +\infty} n^{-1} \sum_{1 \leq i \leq n} \delta_{X_i} = P$ almost surely (Pitman [48, Chapter 2 and Chapter 4]). For $n \geq 1$, let \mathbf{X}_n be a random sample from P , that is we write

$$\begin{aligned} X_i | P &\stackrel{\text{iid}}{\sim} P \quad i = 1, \dots, n, \\ P &\sim \text{PK}(\rho). \end{aligned} \quad (2)$$

Under the sampling model (2), the distribution $\text{PK}(\rho)$ takes on the natural interpretation as a (non-parametric) prior distribution for the unknown $P \in \mathcal{P}$. Hence, the distribution $\text{PK}(\rho)$ is also referred to as Poisson-Kingman prior (Lijoi and Prünster [41]). Because of the (almost sure) discreteness of P , the random sample \mathbf{X}_n from P induces a random partition Π_n of the set \mathbb{N}_n whose blocks are the (equivalence) classes induced by the equivalence relations $i \sim j \iff X_i = X_j$ almost surely. The random partition Π_n is exchangeable, namely the distribution of Π_n is a symmetric function of the sizes of its blocks, and it is referred to as Poisson-Kingman random partition with parameter ρ . Moreover, the sequence of exchangeable random partitions $\Pi = (\Pi_n)_{n \geq 1}$ is consistent in the sense that Π_r is the restriction of Π_n to the first r elements, for all $r < n$ (Pitman [48, Chapter 2 and Chapter 3]). The assumption of consistency of Π implies that Π defines an exchangeable random partition of \mathbb{N} , where exchangeability of Π means that the distribution of Π is invariant under finite permutations of its elements.

DEFINITION 2 (Poisson-Kingman partition model). *Let \mathbf{X}_n be a random sample from P , where $P \sim \text{PK}(\rho)$, and let \mathbb{P}_ρ^n be the distribution of the random partition Π_n of \mathbb{N}_n induced by \mathbf{X}_n . Then the Poisson-Kingman partition model for Π_n is $\{\mathbb{P}_\rho^n : \rho \in \mathcal{L}\}$, and we denote by \mathbb{E}_ρ^n the expectation under \mathbb{P}_ρ^n .*

3. Near-optimal estimation of $U_{n, \lambda n}$

Under the Multinomial model for the observable samples, with $p \in \mathcal{P}(\alpha, L, L')$, we consider the problem of estimating the unseen $U_{n, \lambda n}$ for $\lambda > 1$. We start by introducing an estimator of the tail-index

α and a plug-in estimator of $U_{n,\lambda n}$, and by establishing their convergence rates under a suitable class of loss functions. We make use of the Poisson-Kingman partition model of Definition 2, for a suitable specification of $\rho \in \mathcal{L}$. In particular, we consider the class of Lévy measures $\mathcal{L}' = \{\rho_\alpha : \alpha \in (0, 1)\}$ where

$$\rho_\alpha(dx) = \frac{\alpha}{\Gamma(1-\alpha)} x^{-\alpha-1} dx.$$

\mathcal{L}' is referred to as the class of α -stable Lévy measures. Assuming that $P \sim \text{PK}(\rho_\alpha)$ (Kingman [38]), then

$$\lim_{x \rightarrow 0} x^\alpha \bar{F}_P(x) = \frac{T^{-\alpha}}{\Gamma(1-\alpha)} \quad \text{almost surely.} \quad (3)$$

See Kingman [38] and Gnedin et al. [26] for a proof of (3). According to (3), the parameter $\alpha \in (0, 1)$ of the α -stable Poisson-Kingman model is precisely the tail-index of the random function \bar{F}_P almost-surely. This observation suggests to make use of the Poisson-Kingman partition model $\{\mathbb{P}_{\rho_\alpha}^n : \alpha \in (0, 1)\}$ to obtain an estimator of the tail-index α . We remark that the model $\{\mathbb{P}_{\rho_\alpha}^n : \alpha \in (0, 1)\}$ is used uniquely to construct an estimator of α , whereas the study of such an estimator will be carried under the Multinomial partition model. The next theorem characterizes the maximum likelihood estimator $\hat{\alpha}_n$ of the parameter α under $\{\mathbb{P}_{\rho_\alpha}^n : \alpha \in (0, 1)\}$, and it guarantees its existence and uniqueness under weak conditions.

THEOREM 1. *Let Π_n be a partition of \mathbb{N}_n , and denote by $C_{n,\ell}$ the number of blocks of Π_n with at least ℓ elements, for $\ell = 1, \dots, n$. Then, under the model $\{\mathbb{P}_{\rho_\alpha}^n : \alpha \in (0, 1)\}$ the maximum likelihood estimator $\hat{\alpha}_n$ of the parameter α based on the observation of Π_n satisfies $\varphi(\hat{\alpha}_n) = 0$, where φ is such that*

$$\varphi_n(\alpha) := \sum_{k=1}^{n-1} \frac{\alpha}{k-\alpha} \cdot C_{n,k+1} - (C_{n,1} - 1).$$

The equation $\varphi_n(\alpha)x = 0$ has a unique solution in $(0, 1)$ whenever $C_{n,1} \neq n$ and $C_{n,1} \neq 1$.

Theorem 1 shows that the maximum likelihood estimator of the parameter $\alpha \in (0, 1)$ exists uniquely whenever Π_n is not the partition consisting on n singletons or the partition consisting of a single block. Therefore, without loss of generality, we assume that the estimator $\hat{\alpha}_n$ always exists; alternatively, we may assume the convention that $\hat{\alpha}_n = 0$ if Π_n is the partition that consists of a single block and $\alpha = 1$ if Π_n is the partition that consists of n blocks. Under the Multinomial partition model with $p \in \mathcal{P}(\alpha, L, L')$, the next theorem shows that $\hat{\alpha}_n$ is a consistent estimator of the tail-index α , as $n \rightarrow +\infty$, with respect to a quadratic loss function. Moreover, it holds that the convergence rate of $\hat{\alpha}_n$ is $n^{-\alpha/2} \sqrt{\log n}$.

THEOREM 2. *For $\bar{L} > 0$ and $0 < \underline{\alpha} < \bar{\alpha} < 1$ let $\mathbb{A} = \{(\alpha, L, L') : L, L' \in (0, \bar{L}), \alpha \in (\underline{\alpha}, \bar{\alpha})\}$. For every $\bar{L} > 0$, $0 < \underline{\alpha} < \bar{\alpha} < 1$, and $\varepsilon > 0$ there exists a constant $D > 0$ such that*

$$\limsup_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p \left(|\hat{\alpha}_n - \alpha|^2 > \frac{D \log(n)}{n^\alpha} \right) \leq \varepsilon.$$

According to Theorem 1 and Theorem 2, we make use of $\hat{\alpha}_n$ to introduce an estimator of the unseen $U_{n,\lambda n}$ for $\lambda > 1$. Because $\lambda > 1$ and $p \in \mathcal{P}(\alpha, L, L')$, it is well-known that $U_{n,\lambda n} \asymp (\lambda n)^\alpha$ almost-surely (Gnedin et al. [26]). This suggests to measure the performance of an estimator of $U_{n,\lambda n}$ through

the loss function

$$\ell_\alpha(u, v) = \frac{(u - v)^2}{(\lambda n)^{2\alpha}}. \quad (4)$$

Observe that, under the loss functions (4), the null estimator $\hat{U} = 0$ is such that $\ell_\alpha(\hat{U}, U_{n,\lambda n}) = O_p(1)$. Thus, we look for estimators of $U_{n,\lambda n}$ such that the loss function diminishes, in probability, as $n \rightarrow +\infty$. It is known that the number K_n of distinct species in a random sample \mathbf{X}_n from $p \in \mathcal{P}(\alpha, L, L')$ is such that $\lim_{n \rightarrow +\infty} n^{-\alpha} K_n = L\Gamma(1 - \alpha)$ almost-surely (Gnedin et al. [26]). Accordingly, by assuming that the tail-index α is known, it is natural to consider as an estimator of $U_{n,\lambda n}$ the following quantity

$$K_n((1 + \lambda)^\alpha - 1). \quad (5)$$

In general, the tail-index α is unknown, and hence it must be estimated. This leads to introduce an estimator of $U_{n,\lambda n}$ by combining (5) with $\hat{\alpha}_n$. Here we consider the plug-in (threshold) estimator of the form

$$\hat{U}_{n,\lambda n} := K_n((1 + \lambda)^{\hat{\alpha}_n} - 1) \mathbf{1}_{\{\log(\lambda) \leq C\sqrt{n^{\hat{\alpha}_n}/\log(n)}\}},$$

where $C > 0$ is any constant. Under the Multinomial partition model of Definition 1, with $p \in \mathcal{P}(\alpha, L, L')$, and with respect to the class of loss functions (4), the next theorem shows that $\hat{U}_{n,\lambda n}$ consistently estimates $U_{n,\lambda n}$ all of the way up to $\log(\lambda) \asymp \sqrt{n^\alpha/\log n}$. As a direct consequence of this result, we obtain a range or threshold for $\lambda > 1$ in order to consistently estimate $U_{n,\lambda n}$ over the class $\mathcal{P}(\alpha, L, L')$.

THEOREM 3. *For $\bar{L} > 0$ and $0 < \underline{\alpha} < \bar{\alpha} < 1$ let $\mathbb{A} = \{(\alpha, L, L') : L, L' \in (0, \bar{L}), \alpha \in (\underline{\alpha}, \bar{\alpha})\}$. Then for every $\bar{L} > 0, 0 < \underline{\alpha} < \bar{\alpha} < 1$, and $\varepsilon > 0$ there exists a constant $D > 0$ such that*

$$\limsup_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \mathbb{P}_p \left(\ell_\alpha(\hat{U}_{n,\lambda n}, U_{n,\lambda n}) > D \min \left\{ 1, \frac{\log(n) \log(\lambda)^2}{n^\alpha} \right\} \right) \leq \varepsilon.$$

3.1. Optimality of the estimators $\hat{\alpha}_n$ and $\hat{U}_{n,\lambda n}$

We establish lower bounds for the minimax risks of estimating the tail-index α and the unseen $U_{n,\lambda n}$. Under the Multinomial model for the observable samples, with $p \in \mathcal{P}(\alpha, L, L')$, we start by determining a lower bound for the minimax risks of estimating α , with respect to a quadratic loss, and then we determine a lower bound for the minimax risks of estimating $U_{n,\lambda n}$, with respect to the loss function (4). These results, in combination with Theorem 2 and Theorem 5, allows for an assessment on the optimality of the estimators $\hat{\alpha}_n$ and $\hat{U}_{n,\lambda n}$, where optimality is in the sense of matching minimax lower bounds. To construct minimax lower bounds we make use of the Poisson-Kingman partition model of Definition 2. Precisely, minimax lower bounds are obtained by means of a suitable variation of the method the two fuzzy hypotheses (Tsybakov [53, Chapter 2]), which exploits the Poisson-Kingman distribution as a fuzzy prior. Our problem requires a minor adaptation over Tsybakov [53, Theorem 2.14], mostly to deal with the problem that we aim at deriving minimax lower bounds over a class $\mathcal{P}' \subseteq \mathcal{P}$ but our priors may have a larger support. The adaption of the method is straightforward as long as the prior charges enough \mathcal{P}' . For the sake of completeness, the next theorem gives the desired extension of Tsybakov [53, Theorem 2.14], and it will be critical for the derivation of our minimax lower bounds.

THEOREM 4. *Let ν_0 and ν_1 be prior distributions over \mathcal{P} , let $\mathcal{P}' \subseteq \mathcal{P}$, and let $\delta(\Pi, p) \in \mathbb{R}$ be a (possibly random) functional to be estimated. Write F_0 and F_1 for the joint distributions of (Π, P) , respectively under the prior ν_0 and ν_1 . If there exist $c \in \mathbb{R}$ and $\kappa, \xi \in (0, 1)$ such that $\min\{F_0(\delta(\Pi, P) \leq c), F_1(\delta(\Pi, P) \geq c + 2\varepsilon)\} \geq 1 - \kappa$ and $\min\{\nu_0(\mathcal{P}'), \nu_1(\mathcal{P}')\} \geq 1 - \xi$, then it holds*

$$\inf_{\hat{\delta}} \sup_{p \in \mathcal{P}'} \bar{\mathbb{P}}_p(|\hat{\delta}(\Pi_n) - \delta(\Pi, p)| > \varepsilon) \geq \frac{1}{2}(1 - \|\nu_0(\mathbb{P}_P^n) - \nu_1(\mathbb{P}_P^n)\|_{\text{TV}}) - \kappa - \xi,$$

where $\nu_j(\mathbb{P}_P^n)(\cdot)$ are the mixture distributions $\int_{\mathcal{P}} \mathbb{P}_P^n(\cdot) \nu(dP)$, $j = 0, 1$, and where the infimum is taken over all measurable maps $\hat{\delta}$ that are solely function of Π_n .

The idea consists in the application of Theorem 4 to the case where ν_0 and ν_1 are Poisson-Kingman distributions with Lévy measures ρ_0 and ρ_1 , respectively. In this case, $\nu_j(\mathbb{P}_P^n)$ is equal to the distribution $\mathbb{P}_{\rho_j}^n$ of the Poisson-Kingman partition model of Definition 2. The main challenge then reduces to find an upper bound on $\|\mathbb{P}_{\rho_0}^n - \mathbb{P}_{\rho_1}^n\|_{\text{TV}}$ for general Lévy measures ρ_0 and ρ_1 . In particular, with the help of Pinsker's inequality, it is enough to upper bound $\text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n)$ which turns out to be more convenient. In the next proposition, we derive an upper bound on $\text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n)$ depending only on the moments

$$\mu_j(u, \ell) := \int_0^\infty x^\ell e^{-ux} \rho_j(dx), \quad j = 0, 1, \quad \ell = 1, \dots, n,$$

and the statistic $\mathbf{M}_n := (M_{n,1}, \dots, M_{n,n})$, where $M_{n,\ell}$ denotes the number of blocks of size ℓ in the partition Π_n .

PROPOSITION 1. *Let $G_n \sim \text{Gamma}(n, 1)$ be independent of the Poisson process $(J_j)_{j \geq 1}$ used to construct the random probability measure $P = \sum_{j \geq 1} \frac{J_j}{T} \delta_{s_j}$, and let $U_n := G_n/T$. Then,*

$$\text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n) \leq \sum_{\ell=1}^n \mathbb{E}_{\rho_0}^n \left[M_{n,\ell} \left\{ -\log \frac{\mu_1(U_n, \ell)}{\mu_0(U_n, \ell)} + \frac{\mu_1(U_n, \ell)}{\mu_0(U_n, \ell)} - 1 \right\} \right].$$

We apply Theorem 4 and Proposition 1 to determine a lower bound for the minimax risk of estimating the tail-index α . In general, both Theorem 4 and Proposition 1 are of independent interest, and they may be applied in the context of the estimation of other discrete functionals under the Multinomial model for the observable samples, with $p \in \mathcal{P}(\alpha, L, L')$. We refer to Section 5 for a discussion. Under the Multinomial partition model for the observable samples, with $p \in \mathcal{P}(\alpha, L, L')$, and with respect to a quadratic loss function, the next theorem shows that the estimator $\hat{\alpha}_n$ is near-optimal up to a power of $\log n$ factor.

THEOREM 5. *For $\bar{L} > 0$ and $0 < \underline{\alpha} < \bar{\alpha} < 1$ let $\mathbb{A} = \{(\alpha, L, L') : L, L' \in (0, \bar{L}), \alpha \in (\underline{\alpha}, \bar{\alpha})\}$. For every $0 < \underline{\alpha} < \bar{\alpha} < 1$ there exists a constant $D > 0$ such that*

$$\liminf_{n \rightarrow \infty} \inf_{\hat{\alpha}_n} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p \left(|\hat{\alpha}_n - \alpha|^2 \geq \frac{D}{n^\alpha \log(n)^2} \right) \geq \frac{1}{4},$$

where the infimum is with respect to all measurable maps $\hat{\alpha}_n$ depending only on Π_n .

Under the Multinomial partition model with $p \in \mathcal{P}(\alpha, L, L')$, and with respect to the loss function (4), the next theorem shows that the estimator $\hat{U}_{n,\lambda n}(\hat{\alpha}_n)$ is near-optimal up to a power of $\log n$ factor. The proof of the theorem consists in establishing that the estimation of $U_{n,\lambda n}$ is a harder problem than the estimation of α if λ is sufficiently large, thus reducing the problem to the estimation of the tail-index.

THEOREM 6. *For $\bar{L} > 0$ and $0 < \underline{\alpha} < \bar{\alpha} < 1$ let $\mathbb{A} = \{(\alpha, L, L') : L, L' \in (0, \bar{L}), \alpha \in (\underline{\alpha}, \bar{\alpha})\}$, and let $b > 0$ be such that it holds true $\log \lambda \geq b \log n$. Then for every $0 < \underline{\alpha} < \bar{\alpha} < 1$ there exists constant $D > 0$ such that*

$$\liminf_{n \rightarrow \infty} \inf_{\hat{U}} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p \left(\ell_{\alpha}(\hat{U}, U_{n,\lambda n}) \geq D \min \left\{ 1, \frac{\log(\lambda)^2}{n^{\alpha} \log(n)^2} \right\} \right) \geq \frac{1}{8},$$

where the infimum is with respect to all measurable maps \hat{U} depending only on Π_n .

According to Theorem 6, no estimator can estimate $U_{n,\lambda n}$ in the range $\log(\lambda) \geq n^{\alpha/2} \log(n)$; moreover, according to Theorem 3, this range or threshold (of estimability) is almost attained by the estimator $\hat{U}_{n,\lambda n}$. As a corollary of Theorem 6 and Theorem 3 it holds that, up to polylog factors, the range $\log \lambda \asymp n^{\alpha/2}$ is the best possible range for consistently estimating $U_{n,\lambda n}$, and such a range is achieved by the estimator $\hat{U}_{n,\lambda n}$. In the range $\log(\lambda) \gtrsim n^{\alpha/2} \log(n)$, the trivial estimator $\hat{U} = 0$ is a minimax estimator, with no extra $\log(n)$ factor. That is, in the range $\log(\lambda) \gtrsim n^{\alpha/2} \log(n)$ the best we can do to estimate $U_{n,\lambda n}$ is to do nothing. We observe that Theorem 6 holds true for large values of λ , typically $\lambda \gtrsim n^b$. This is because of the arguments used in the proof, which establish that the problem of estimating $U_{n,\lambda n}$ is at least as hard as estimating the tail-index α ; these two problems are essentially equivalent, provided that λ is not too small. For small values of λ , however, we do not believe that the two problems are equivalent. In particular, if $\lambda < 1$ we expect that smoothed Efron-Thisted estimator (Good and Toulmin [27], Efron and Thisted [19], Orlicsky et al. [46]) may have a stronger theoretical and experimental performance than our estimator. Furthermore, we believe the two problems may be equivalent at a much lower range of λ than that provided by Theorem 3. Understanding this range requires to obtain an exact matching for the upper and lower bounds of the risk in the estimation of α , which remains a challenging problem. This constitutes an interesting problem to investigate in future research.

4. Numerical illustrations

Power-law type distributions are arguably the most common examples of distributions in the class of regularly varying tail distributions $\mathcal{P}(\alpha, L, L')$. Power-law type distributions occur in many situations of scientific interest, and nowadays have significant consequences for the understanding of natural and social phenomena (Clauset et al. [14]). In this section, we present a numerical illustration of our methodology for synthetic data from Zipf distributions (a favorable scenario), modified Zipf distributions with two tail indices (a less favorable scenario), text data from Gutenberg books and Wikipedia pages, and data from humans' electronic activities in email communication and Twitter posts. We compare our estimator to smoothed Efron-Thisted estimators, which may be considered as the state-of-the-art estimators for the unseen-species problem (Good and Toulmin [27], Efron and Thisted [19], Orlicsky et al. [46]).

4.1. Synthetic data : favorable scenario

The most favorable scenario for our estimators is undoubtedly when the model is well-specified, i.e. Π_n has a Poisson-Kingman distribution with Lévy measure ρ_α for some $\alpha \in (0, 1)$. As we are mainly interested in data generated according to the Multinomial model, we leave the ideal case apart and consider instead a situation where the model $\{\mathbb{P}_{\rho_\alpha} : \alpha \in (0, 1)\}$ is misspecified, but in a benign fashion. A prototypical example of such benign misspecification is when (X_1, \dots, X_n) generated (i.i.d) from a Zipf distribution (Clauset et al. [14]). For $s > 1$, the Zipf(s) distribution on \mathbb{Z}_+ with parameter s has probability mass function $P(j) \propto j^{-s}$ for all $j \in \mathbb{Z}_+$. The parameter s controls the tail behaviour of Zipf distribution: the smaller s the heavier is the tail of the distribution, i.e., the smaller s the larger the fraction of species with low frequency. In particular, it follows that the Zipf(s) distribution has a tail-index $\alpha = 1/s$ and belongs to $\mathcal{P}(\alpha, L, L')$ for some $L, L' > 0$. Therefore, we expect synthetic Zipf data to be a favorable situation for our estimators. We start by an empirical validation of the maximum likelihood estimator $\hat{\alpha}_n$ of the tail-index. We compute Monte Carlo (MC) estimates of $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2]$, where the expectation is understood with respect to data generated from a Zipf($1/\alpha_0$) distribution. The expectation is approximated using 1000 MC samples. Simulations were run for $n \in \{1000, 10000, 100000, 1000000\}$ and $\alpha_0 \in \{0.2, 0.5, 0.8\}$. Estimated values of $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2]$ are reported in Table 1 and plotted in Figure 4.1. About the plots, we have adjusted a curve $\log \mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2] = -r \log(n) + C$ using least squares estimation. We found estimated values of r that are 0.27, 0.54, and 0.79, respectively for $\alpha_0 = 0.2$, $\alpha_0 = 0.5$, and $\alpha_0 = 0.8$, which are coherent with our theoretical finding that $\hat{\alpha}_n$ converges at rate $n^{-\alpha_0}$. Figure 4.1 contains the histograms of $\hat{\alpha}_n$, showing that the distribution of $\hat{\alpha}_n$ does concentrate on α_0 , and this concentrations happens faster as α_0 gets larger.

Table 1. Monte Carlo estimates of $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2]$ based on 1000 replicates of n data points generated (i.i.d) from Zipf($1/\alpha_0$)

n	$\alpha_0 = 0.2$	$\alpha_0 = 0.5$	$\alpha_0 = 0.8$
1000	$2.75 \cdot 10^{-3}$	$1.28 \cdot 10^{-3}$	$1.92 \cdot 10^{-4}$
10000	$1.3 \cdot 10^{-3}$	$3.77 \cdot 10^{-4}$	$3.12 \cdot 10^{-5}$
100000	$7.39 \cdot 10^{-4}$	$1.19 \cdot 10^{-4}$	$4.80 \cdot 10^{-6}$
1000000	$4.28 \cdot 10^{-4}$	$3.02 \cdot 10^{-5}$	$8.18 \cdot 10^{-7}$

We consider the same MC setting to make an empirical validation of $\hat{U}_{n,\lambda n}$. We compute MC estimates of the risk

$$\tilde{R}_{n,\lambda n} := \frac{\mathbb{E}[(\hat{U}_{n,\lambda n} - U_{n,\lambda n})^2]}{\mathbb{E}[U_{n,\lambda n}]^2}.$$

In particular, we fix $n = 1000$, and increasing values of λ from $\lambda = 1.1$ to $\lambda = 40$. We compare our estimator with the state-of-the-art estimators of Orlicsky et al. [46], which are smoothed Efron-Thisted estimator based on Poisson smoothing or Binomial smoothing distributions. Estimated values of $\tilde{R}_{n,\lambda n}$ are reported in Table 1 and plotted in Figure 4.1. In Table 1, we only report the best value for smoothed Efron-Thisted estimators, which is not necessarily attained using the same smoothing in the various scenarios. Our results show that smoothed Efron-Thisted estimators can estimate $U_{n,\lambda n}$ with good accuracy only when λ is small, which is expected, whereas by estimating the tail-index the estimation of the unseen remains accurate even for large values of λ . In particular, $\hat{U}_{n,\lambda n}$ can estimate the unseen with more accuracy than smoothed Efron-Thisted estimators in the situations where the data come from




Figure 1. Monte Carlo estimates of $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2]$ based on 1000 replicates of n data points generated (i.i.d) from $\text{Zipf}(1/\alpha_0)$. The blue points represent the estimated values of $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2]$ for $n \in \{1000, 10000, 100000, 1000000\}$. The dashed line is a fit corresponding to $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2] = e^{-4.14} n^{-0.27}$ (left), $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2] = e^{-2.92} n^{-0.54}$ (center), and $\mathbb{E}[(\hat{\alpha}_n - \alpha_0)^2] = e^{-3.08} n^{-0.79}$ (right).




Figure 2. Histograms of $\hat{\alpha}_n$ based on 1000 replicates of n data points generated (i.i.d) from $\text{Zipf}(1/\alpha_0)$.

a distribution with a tail-index. It does not make clear, however, how large λ can be taken. To illustrate that the unseen can eventually be estimated with good accuracy even for large λ , we ran the same simulation for $\hat{U}_{n,\lambda n}$, $\alpha_0 = 0.8$, $n = 1000$, but with λ ranging from 100 to 10000. From Figure 4.1 it is clear that $\hat{R}_{n,\lambda n}$ grows logarithmically with λ , which agrees with our theoretical findings. Hence, in

this example, even for values of λ as large as $\lambda = 10000$ the accuracy of estimating the unseen remains satisfactory.

Table 2. Monte Carlo estimates of $\tilde{R}_{n,\lambda n}$ for various estimators of the unseen based on 1000 replicates of n data points generated (i.i.d) from $\text{Zipf}(1/\alpha_0)$. Simulations are made on the basis of $n = 1000$ observed samples. The column $\mathbb{E}[U_{n,\lambda n}]$ is a MC estimate of the expected number of unseen, the column FN gives a MC estimate of $\tilde{R}_{n,\lambda n}$ for $\hat{U}_{n,\lambda n}$, and the column O gives the best value of $\tilde{R}_{n,\lambda n}$ obtained using the three estimators proposed in Orilitsky et al. [46].

λ	$\alpha_0 = 0.2$			$\alpha_0 = 0.5$			$\alpha_0 = 0.8$		
	$\mathbb{E}[U_{n,\lambda n}]$	FN	O	$\mathbb{E}[U_{n,\lambda n}]$	FN	O	$\mathbb{E}[U_{n,\lambda n}]$	FN	O
1.1	1.20	0.99	2.35	62.53	$2.56 \cdot 10^{-2}$	$3.97 \cdot 10^{-2}$	1741	$9.66 \cdot 10^{-4}$	$1.35 \cdot 10^{-3}$
1.5	1.54	0.80	1.75	80.77	$2.24 \cdot 10^{-2}$	$4.89 \cdot 10^{-2}$	2323	$8.75 \cdot 10^{-4}$	$8.48 \cdot 10^{-3}$
2	1.86	0.69	1.64	101.63	$2.13 \cdot 10^{-2}$	$7.28 \cdot 10^{-2}$	3025	$8.30 \cdot 10^{-4}$	$4.12 \cdot 10^{-2}$
3	2.41	0.56	1.41	138.39	$1.92 \cdot 10^{-2}$	$1.51 \cdot 10^{-1}$	4364	$7.65 \cdot 10^{-4}$	$1.66 \cdot 10^{-1}$
4	2.86	0.49	1.21	170.81	$1.88 \cdot 10^{-2}$	$2.46 \cdot 10^{-1}$	5638	$7.61 \cdot 10^{-4}$	$2.87 \cdot 10^{-1}$
5	3.24	0.45	1.40	200.43	$1.82 \cdot 10^{-2}$	$2.20 \cdot 10^{-1}$	6861	$7.65 \cdot 10^{-4}$	$2.61 \cdot 10^{-1}$
6	3.55	0.43	1.98	227.71	$1.84 \cdot 10^{-2}$	$2.86 \cdot 10^{-1}$	8044	$7.79 \cdot 10^{-4}$	$3.43 \cdot 10^{-1}$
8	4.06	0.40	1.64	276.05	$1.83 \cdot 10^{-2}$	$4.04 \cdot 10^{-1}$	10312	$8.05 \cdot 10^{-4}$	$4.80 \cdot 10^{-1}$
10	4.54	0.38	1.37	319.49	$1.84 \cdot 10^{-2}$	$5.00 \cdot 10^{-1}$	12483	$8.30 \cdot 10^{-4}$	$5.83 \cdot 10^{-1}$
15	5.42	0.35	1.09	414.50	$1.85 \cdot 10^{-2}$	$6.61 \cdot 10^{-1}$	17598	$8.87 \cdot 10^{-4}$	$7.42 \cdot 10^{-1}$
20	6.15	0.33	0.99	494.98	$1.92 \cdot 10^{-2}$	$7.55 \cdot 10^{-1}$	22399	$9.52 \cdot 10^{-4}$	$8.25 \cdot 10^{-1}$
30	7.26	0.32	0.95	631.13	$2.00 \cdot 10^{-2}$	$8.52 \cdot 10^{-1}$	31377	$1.03 \cdot 10^{-3}$	$9.04 \cdot 10^{-1}$
40	8.12	0.31	0.95	746.16	$2.12 \cdot 10^{-2}$	$8.99 \cdot 10^{-1}$	39783	$1.10 \cdot 10^{-3}$	$9.39 \cdot 10^{-1}$

4.2. Synthetic data: less favorable scenario

We consider a collection of n data points modeled as i.i.d. random samples (X_1, \dots, X_n) coming from a distribution p such that $p_j \propto j^{-1/\alpha}$ for $j \in \{1, \dots, J\}$ and $p_j \propto j^{-1/\beta}$ for $j > J$; we shall refer to such distribution as “double Zipf distribution”. We assume that $0 < \beta < \alpha < 1$. Such a double Zipf distribution is less favorable for our estimator of the unseen, because it is constructed in such a way that it has a true tail-index of β , while if J is too large it will look like data are generated from a $\text{Zipf}(1/\alpha)$ distribution. Indeed, the Lemma 1 below, shows that the double Zipf distribution will always be in $\mathcal{P}(\beta, L, L')$ for some $L, L' > 0$, but the larger J is taken, the larger L' will be. Performances of our estimators at finite sample size n , however, are strongly dependent on the value L' (the smaller the better).

LEMMA 1. *For any $1 > \alpha > \beta > 0$ and $J \in \mathbb{N}$, let $p \propto \sum_{j=1}^J j^{-1/\alpha} \delta_j + \sum_{j>J} j^{-1/\beta} \delta_j$. There exist constants $L, C, J_0 > 0$ such that for all $J \geq J_0$,*

$$\sup_{x \in (0,1)} \frac{|x^\beta \bar{F}_p(x) - L|}{\sqrt{x^\beta \log(e/x)}} \leq C \cdot (\alpha - \beta) \sqrt{\frac{J^{2\alpha/\beta} \log(J)}{J}}.$$

The simulations displayed in Figure 4.2 illustrate how the risks for α and for the unseen deteriorate as a function of J for fixed n . We take $n = 1000$, $\alpha = 0.5$, $\beta = 0.4$, $\lambda = 50$ and J ranging from 10 to 290. For both $\hat{\alpha}_n$ and $U_{n,\lambda n}$ we compute MC estimates of the risks (as defined in the previous section) using 1000 samples for the risk of estimating β , and 100 samples for the unseen [to make the simulation run in short time]. We see that the risk $\mathbb{E}_p[(\hat{\alpha}_n - \beta)^2]$ behaves as expected; that is is small when J is reasonably small and increases with J , until the point where it gets of the order




Figure 3. Monte Carlo estimates of $\tilde{R}_{n, \lambda n}$ for various estimators of the unseen based on 1000 replicates of n data points generated (i.i.d) from $\text{Zipf}(1/\alpha_0)$. Simulations are made on the basis of $n = 1000$ observed samples. The red curve is the estimator $\hat{U}_{n, \lambda n}$ and the three other curves correspond to the three estimators proposed in Orlicsky et al. [46], respectively for a Poisson smoothing (blue) and two variants of a Binomial smoothing (green and magenta) of the Good-Toulmin estimator.




Figure 4. Monte Carlo estimates of $\tilde{R}_{n, \lambda n}$ for our estimator $\hat{U}_{n, \lambda n}$ based on 100 replicates of n data points generated (i.i.d) from $\text{Zipf}(1/\alpha_0)$, with $\alpha_0 = 0.8$. The growth of $\tilde{R}_{n, \lambda n}$ in λ is logarithmic, which is consistent with the theory.

of $(\alpha - \beta)^2 = 0.01$. The plot for $\tilde{R}_{n, \lambda n}$ might look more surprising at first, but it is indeed what is

expected. We see three different regions for the risk. First, if J is sufficiently small, then $\hat{\alpha}_n$ estimates β with good accuracy and the number of unseen is of order $(\lambda n)^\beta$; everything goes well in this region. Then, when J gets larger it becomes impossible to estimate β , and we rather have $\hat{\alpha}_n \approx \alpha$, but the number of unseen is still of order $(\lambda n)^\beta$; whence a very high risk. Finally, when J gets very large, we again have $\hat{\alpha}_n \approx \alpha$, but the number of unseen is also of order $(n\lambda)^\alpha$. Basically everything happens as if p was a Zipf($1/\alpha$) distribution; thus everything goes well again here. We emphasize that the reason the risk gets apparently better for large J is just an artifact of the simulation that uses a fixed value of n and λ .




Figure 5. On the left: Monte Carlo estimates of $\mathbb{E}[(\hat{\alpha}_n - \beta)^2]$ as a function of J based on 1000 replicates of n data points generated (i.i.d) from a double Zipf distributions with parameters $(1/0.5, 1/0.4, J)$. On the right: Monte Carlo estimates of $\tilde{R}_{n, \lambda n}$ for our estimator $\hat{U}_{n, \lambda n}(\hat{\alpha}_n)$ as a function of J based on 100 replicates of n data points generated (i.i.d) from a double Zipf distributions with parameters $(1/0.5, 1/0.4, J)$.

4.3. Real data

We consider text data arising from Gutenberg books and Wikipedia pages, and data from humans' electronic activities in email communication and Twitter posts. The email dataset (eMail) consists of a collection of data from the email communication activity of a Department at Università degli Studi di Padova in the years 2012 and 2013 (Formentin et al. [23]); the dataset is in the form of a table, where each row contains information of (sender, receiver, timestamp); for our analysis we set the sender identity to be a species, and the number of emails sent from a sender to be the frequency of the species. The Twitter dataset (Twitter) consists of a collection data from the communication activity of Twitter (Tria et al. [52], Monechi et al. [42]); the dataset is in the form of a table, where each row contains information of (timestamp, hashtag, user); for our analysis we set an hashtag to be a species, and the number of tweets containing an hashtag to be the frequency of the species. The Wikipedia (Wiki) and Gutenberg (Gut) datasets consist of a collection of data from Wikipedia pages and Gutenberg books (Tria et al. [52], Monechi et al. [42]); for our analysis we set a word to be a species, and the number of occurrence of a word to be the frequency of the species. Summaries of these datasets are reported in Table 3.

For each dataset, we make an empirical evaluation of $\hat{U}_{n, \lambda n}$ as follows. Given that the dataset has size N , see Table 3, for each value of λ we obtain a sample of size $n = \lfloor N/(1 + \lambda) \rfloor$ from the dataset by sampling individuals without replacement. Then, from this sample, we estimate $U_{n, \lambda n}$ in

the dataset and we report the normalized error $E = (\hat{U}_{n,\lambda n} - U_{n,\lambda n})/U_{n,\lambda n}$. We repeat this process 100 times in order to obtain 100 copies E_1, \dots, E_{100} and we compute the average squared error $\widehat{MSE} = \frac{1}{100} \sum_{1 \leq i \leq 100} E_i^2$. We also apply the same procedure to the three smoothed Efron-Thisted estimators proposed in Orlicsky et al. [46]. Results of our experiments are available in Figure 4.3, from which we see that for all datasets considered there is an interest in leveraging the power-law behaviour and estimating the tail-index to predict the unseen, as opposed to a fully agnostic approach *à la* Efron-Thisted.

Table 3. Summary statistics of the real datasets.

	Gut	eMail	Wiki	Twitter
Size N of the dataset	6314483	345469	1480277	1734849
Number of species	147801	69462	137030	608644




Figure 6. \widehat{MSE} for various estimators of the unseen on real datasets (exact procedure detailed in Section 4.3). The red curve is the estimator $\hat{U}_{n,\lambda n}$ and the three other curves correspond to the three estimators proposed in Orlicsky et al. [46], respectively for a Poisson smoothing (blue) and two variants of a Binomial smoothing (green and magenta) of the Good-Toulmin estimator.

5. Discussion

The unseen-species problem dates back to the seminal works of Fisher et al. [22], Good and Toulmin [27] and Efron and Thisted [19], and it has been the subject of recent breakthrough studies by Orlicsky

et al. [46], Wu and Yang [56, 57] and Polyanskiy and Wu [49]. In this paper, we offered new advances on these recent studies, presenting the first work on the estimation of $U_{n,\lambda n}$ under the assumption that the underlying unknown distribution p has regularly varying tails of index $\alpha \in (0, 1)$. This is motivated by the ubiquitous power-law type distributions, which nowadays occur in many natural and social phenomena (Clauset et al. [14]). Under a semi-parametric assumption of regular variation for the tail behaviour of p , we introduced an estimator of $U_{n,\lambda n}$ and we showed that it has the following theoretical guarantees: i) it is minimax near-optimal up to a power of $\log n$ factor; ii) it is consistent all of the way up to $\log \lambda \asymp n^{\alpha/2}/\sqrt{\log n}$, and this range is the best possible. The proposed estimator is simple, linear in the sampling information, computationally efficient and scalable to massive datasets, and its provable guarantees hold uniformly over the class of regularly varying tail distributions. From an empirical perspective, it is shown that our estimator outperforms existing ones on several synthetic and real datasets. To the best of our knowledge, this is the first work on the estimation of $U_{n,\lambda n}$ under assumptions on the underlying unknown distribution p . Our results rely on a novel approach, of independent interest, which combines the method of the two fuzzy hypotheses for minimax estimation of discrete functionals (Tsybakov [53], Wu and Yang [56, 57]) with Bayesian arguments under Poisson-Kingman priors for p .

Our semi-parametric assumption entails that $|\bar{F}_p(x) - Lx^{-\alpha}| \lesssim x^{-\alpha/2} \log(e/x)^{1/2}$, which guarantees that $\bar{F}_p(x)$ approaches $Lx^{-\alpha}$ fast enough when $x \rightarrow 0$. An interesting problem is to study the case where $x^\alpha \bar{F}_p(x) \rightarrow L$ is true but convergence happens at a slower rate than $x^{-\alpha/2} \log(e/x)^{1/2}$; say at rate $x^{-\beta}$ for $\beta \in (0, \alpha/2)$. In general, we can adapt the proof of Theorem 2 to obtain the following weaker Proposition.

PROPOSITION 2. *Suppose there exist $\alpha \in (0, 1)$, $\beta \in (0, \alpha)$, and $L > 0$ such that $\bar{F}_p(x) = Lx^{-\alpha} + O(x^{-\beta})$. Then, for every $\varepsilon > 0$ there is a constant $D > 0$ such that*

$$\bar{\mathbb{P}}_p\left(|\hat{\alpha}_n - \alpha|^2 > D \max(n^{-\alpha}, n^{-2\beta})\right) \leq \varepsilon.$$

According to Proposition 2, when convergence of $x^\alpha \bar{F}_p(x) \rightarrow L$ happens at the rate $x^{-\beta}$, $\beta \in (0, \alpha/2)$, we can still estimate the tail-index α , but at a deteriorated rate of $n^{-2\beta}$ in place of $n^{-\alpha} \log(n)$. Rates for the estimation of the unseen $U_{n,\lambda n}$ can be adapted in the same way. It is however unclear how to obtain matching minimax lower bounds under this weaker modeling assumption. We cannot exclude that using Poisson-Kingman partitions as fuzzy hypotheses produces too “nice” hypotheses to get matching rates. Another obvious, but more challenging question, would be to investigate the necessity of regular variation itself. We know that regular variations are not necessary, as shown in the works by Chee and Wang [13], Giguelay and Huet [25], Balabdaoui and Kulagina [3], but it would be of interest to know if weaker and non-parametric power-law modeling assumptions are feasible, assuming for instance that $x^{-\alpha_1} \lesssim \bar{F}_p(x) \lesssim x^{-\alpha_2}$ for some $0 < \alpha_1 < \alpha_2 < 1$. We conjecture that this is only possible if λ is small and/or the loss function weak enough, such as the NMSE considered in Orlitsky et al. [46].

The unseen-species problem belongs to a broad class of discrete functional estimation problems, commonly referred to as “species-sampling problems” (Balocchi et al. [4]). These problems uses n random samples from an unknown underlying distribution p , and call for estimating features of p or features of λn additional unobservable samples. Recent noteworthy works on “species sampling problems” concerned with the estimation of, e.g., the support size (Valiant and Valiant [54], Wu and Yang [57]), the entropy (Jiao et al. [36], Wu and Yang [56]), the missing mass (Ohannessian and Dahleh [45], Mossel and Ohannessian [43], Ben-Hamou et al. [6], Ayed et al. [2]) and the risk of disclosure (Camerlenghi et al. [10]). Interest in these quantities first appeared in ecology, and it has grown in the

recent years driven by challenging applications in biological and physical sciences, statistical machine learning, engineering, theoretical computer science, information theory and forensic DNA analysis. Most of the works of “species sampling problems” do not rely on any assumption on the underlying distribution p . The sole exceptions are Ohannessian and Dahleh [45] and Ayed et al. [2], which contain a preliminary study of the problem of consistent estimation of the missing mass under the assumption of regularly varying tails for p , and leave open the problem of minimax optimal estimation. Our work paves the way to study “species sampling problems”, beyond the unseen-species problem, under the assumption that p has regularly varying tails. We believe that the techniques developed in our work are of direct application to the problem of disclosure risk assessment considered in Camerlenghi et al. [10] and to the problem of estimating the number of unseen species with multiplicity considered in Hao and Li [30].

Acknowledgements

The authors are grateful to the Editor (Professor Mark Podolsky), the Associate Editor and three anonymous Referees for all their comments, corrections, and numerous suggestions that improved remarkably the paper. Stefano Favaro and Zacharie Naulet received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 817257. Stefano Favaro gratefully acknowledge the financial support from the Italian Ministry of Education, University and Research (MIUR), “Dipartimenti di Eccellenza” grant 2018-2022.

Appendix A: Upper bounds for the estimators of the tail-index and the unseen

This appendix contains the proof of the Theorems 1, 2, 3 and Proposition 2.

A.1. Proof of Theorem 1

It is a well-known fact that α -stable Poisson Kingman partition model is equivalent to the *two parameters Chinese Restaurant process* with choice of parameters $(\alpha, 0)$; see (Pitman [48, see Section 3.2 for the definition of the two parameters process, and Section 4.2 for the equivalence]). Let write $M_{n,\ell}$ for the numbers of blocks of size ℓ in the partition Π_n . Obviously $\mathbf{M}_n := (M_{n,1}, \dots, M_{n,n})$ entirely determines the law of the partition Π_n , so we might as well assume we observe \mathbf{M}_n . Then by [48, Theorem 3.2], the probability of observing $\mathbf{M}_n = (m_1, \dots, m_n)$ under the parameter α is given by

$$p_\alpha(m_1, \dots, m_n) = \frac{\alpha^{k-1} \prod_{i=0}^{k-2} (1+i) \prod_{j=1}^n (\prod_{i=0}^{j-2} (1-\alpha+i))^{m_j}}{\prod_{i=0}^{n-2} (1+i)},$$

where $k = \sum_{j=1}^n m_j$. So any maximizer of $\alpha \mapsto p_\alpha(\mathbf{M}_n)$ is equivalently a maximizer of

$$F(\alpha) := (K_n - 1) \log(\alpha) + \sum_{j=1}^n \sum_{i=0}^{j-2} M_{n,j} \log(1 - \alpha + i).$$

Differentiating the previous with respect to α gives that $\hat{\alpha}_n$ if it exists must be solution to

$$0 = \frac{K_n - 1}{\alpha} - \sum_{j=1}^n \sum_{i=0}^{j-2} \frac{M_{n,j}}{1 - \alpha + i} = \frac{K_n - 1}{\alpha} - \sum_{i=1}^{n-1} \frac{C_{n,i+1}}{i - \alpha}.$$

This proves that $\hat{\alpha}_n$ must be solution to $\varphi(\alpha) = 0$. Differentiating F one more times gives

$$F''(\alpha) = -\frac{K_n - 1}{\alpha^2} - \sum_{i=1}^{n-1} \frac{C_{n,i+1}}{(i - \alpha)^2}.$$

This establish that $F''(\alpha) < 0$ for all $\alpha \in (0, 1)$, so if there is a solution to $F'(\alpha) = 0$, equivalently to $\varphi_n(\alpha) = 0$, then it is a unique maximiser. We observe that $C_{n,1} \geq C_{n,2} \geq \dots$. Unless $C_{n,2} = 0$ or $K_n = 1$, it is true that $\lim_{\alpha \rightarrow 0} \varphi_n(\alpha) = -(K_n - 1) < 0$ and $\lim_{\alpha \rightarrow 1} \varphi_n(\alpha) = +\infty$. Then it must be that $\varphi_n(\alpha) = 0$ has exactly one solution provided $C_{n,2} \neq 0$ and $K_n \neq 1$. Finally, we note that $C_{n,1} = K_n$, so $C_{n,2} = 0$ is equivalent $K_n = n$, which is equivalent to Π_n being the partition consisting only on singletons. Obviously $K_n \neq 1$ is equivalent to Π_n not being the partition with a single block.

A.2. Proof of Theorem 2

We remark that $\alpha \mapsto \alpha^{-1} \varphi_n(\alpha)$ is the score function of the model $\{\mathbb{P}_{\rho_\alpha} : \alpha \in (0, 1)\}$. We define $\Phi_{n,p}(\alpha) := \mathbb{E}_p[\varphi_n(\alpha)]$ and $\alpha_n^*(p)$ the solution to the equation $\Phi_{n,p}(\alpha) = 0$. We note that the model is misspecified, such that there is no reason to have $\alpha_n^*(p) = \alpha$ when $p \in \mathcal{P}(\alpha, L, L')$. We show in a first time in Proposition A.1 that distributions in $\mathcal{P}(\alpha, L, L')$ are sufficiently well behaved to control the bias $|\alpha_n^*(p) - \alpha|$. In a second time, we establish in Proposition A.2 a uniform concentration result for $|\hat{\alpha}_n - \alpha_n^*(p)|$ [using classical arguments for maximum likelihood estimation; see [55]]. The proof of Theorem 2 follows by combining the results of Propositions A.1 and A.2.

PROPOSITION A.1. *For all $0 < \underline{\alpha} < \bar{\alpha} < 1$ and all \bar{L} ,*

$$\limsup_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \frac{|\alpha_n^*(p) - \alpha|}{n^{-\alpha/2} \sqrt{\log(n)}} \leq \frac{5L'}{\bar{L}}.$$

Proof. We first show that $|\Phi_{n,p}(\alpha)|$ must remain small if $p \in \mathcal{P}(\alpha, L, L')$. By Lemma E.2, for all p and all $\alpha \in (0, 1)$

$$\Phi_{n,p}(\alpha) = 1 + \sum_{k=2}^n \frac{k\alpha}{k - (1 + \alpha)} \binom{n}{k} \int_0^1 \bar{F}_p(x) x^{k-1} (1-x)^{n-k} dx - n \int_0^1 \bar{F}_p(x) (1-x)^{n-1} dx.$$

We decompose \bar{F}_p in two parts such that $\bar{F}_p(x) = Lx^{-\alpha} + \delta(x)$, where by assumption we have that $|\delta(x)| \leq L' \sqrt{x^{-\alpha} \log(e/x)}$. We recall that $\int_0^1 x^{-\alpha+k-1} (1-x)^{n-k} dx = \frac{\Gamma(k-\alpha)\Gamma(n-k+1)}{\Gamma(n+1-\alpha)}$ for all $k = 1, \dots, n$ and all $n \geq 1$. Therefore,

$$\Phi_{n,p}(\alpha) = 1 + \frac{L\alpha n!}{\Gamma(n+1-\alpha)} \sum_{k=2}^n \frac{\Gamma(k-\alpha-1)}{\Gamma(k)} - \frac{Ln! \Gamma(1-\alpha)}{\Gamma(n+1-\alpha)}$$

$$+ \sum_{k=2}^n \frac{k\alpha}{k-(1+\alpha)} \binom{n}{k} \int_0^1 \delta(x) x^{k-1} (1-x)^{n-k} dx - n \int_0^1 \delta(x) (1-x)^{n-1} dx.$$

We note that

$$\sum_{k=2}^n \frac{\Gamma(k-\alpha-1)}{\Gamma(k)} = \frac{\Gamma(1-\alpha)}{\alpha} - \frac{\Gamma(n-\alpha)}{\alpha\Gamma(n)}$$

which gives the estimate

$$\begin{aligned} |\Phi_{n,p}(\alpha)| &\leq 1 + \frac{L\alpha n! \Gamma(n-\alpha)}{\Gamma(n+1-\alpha)\Gamma(n)} + n \int_0^1 |\delta(x)| (1-x)^{n-1} dx \\ &\quad + \sum_{k=2}^n \frac{k\alpha}{k-(1+\alpha)} \binom{n}{k} \int_0^1 |\delta(x)| x^{k-1} (1-x)^{n-k} dx. \end{aligned}$$

Then, we remark that the function $x \mapsto x/[x-(1+\alpha)]$ is monotone decreasing on $(1+\alpha, \infty)$, which implies that $k/(k-(1+\alpha)) \leq 2/(1-\alpha)$ for all $k \geq 2$. It follows the estimate

$$\begin{aligned} &\sum_{k=2}^n \frac{k\alpha}{k-(1+\alpha)} \binom{n}{k} \int_0^1 |\delta(x)| x^{k-1} (1-x)^{n-k} dx \\ &\leq \frac{2\alpha}{1-\alpha} \int_0^1 |\delta(x)| \sum_{k=1}^n \binom{n}{k} x^{k-1} (1-x)^{n-k} dx \\ &= \frac{2\alpha}{1-\alpha} \int_0^1 |\delta(x)| \frac{1-(1-x)^n}{x} dx \end{aligned}$$

and then

$$\begin{aligned} |\Phi_{n,p}(\alpha)| &\leq 1 + \frac{L\alpha n! \Gamma(n-\alpha)}{\Gamma(n+1-\alpha)\Gamma(n)} + n \int_0^1 |\delta(x)| (1-x)^{n-1} dx \\ &\quad + \frac{2\alpha}{1-\alpha} \int_0^1 |\delta(x)| \frac{1-(1-x)^n}{x} dx. \end{aligned}$$

By Stirling's formula and by the Lemma A.1, we obtain that,

$$\limsup_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \frac{(1-\alpha) |\Phi_{n,p}(\alpha)|}{5L'\Gamma(1-\alpha/2) \cdot n^{\alpha/2} \sqrt{\log(n)}} \leq 1.$$

Finally, we show that $\Phi_{n,p}(\alpha') = 0$ implies that α' must be close to α when $p \in \mathcal{P}(\alpha, L, L')$. In fact, from the definition of φ_n we see that $\Phi'_{n,p}(\alpha') = \sum_{k=1}^{n-1} \frac{k}{(k-\alpha')^2} \bar{\mathbb{E}}_p[C_{n,k+1}] \geq \frac{\bar{\mathbb{E}}_p[C_{n,2}]}{(1-\alpha')^2}$. By Lemma E.2, for n large enough it will be true that $\bar{\mathbb{E}}_p[C_{n,2}] \gg 0$, so that $\Phi_{n,p}$ is monotone increasing on $(0, 1)$. It follows from a Taylor expansion that for any $\alpha' \leq \alpha - \varepsilon_{n,\alpha}$ [$\varepsilon_{n,\alpha}$ a decreasing sequence

of real numbers going to zero, to be chosen accordingly]

$$\Phi_{n,p}(\alpha') \leq \Phi_{n,p}(\alpha - \varepsilon_{n,\alpha}) \leq \Phi_{n,p}(\alpha) - \inf_{z \in (0, \varepsilon_{n,\alpha})} \Phi'_{n,p}(\alpha - z) \varepsilon_{n,\alpha} \leq \Phi_{n,p}(\alpha) - \frac{\bar{\mathbb{E}}_p[C_{n,2}] \varepsilon_{n,\alpha}}{(1 - \alpha)^2}. \quad (\text{A.1})$$

Similarly if $\alpha' \geq \alpha + \varepsilon_{n,\alpha}$,

$$\Phi_{n,p}(\alpha') \geq \Phi_{n,p}(\alpha) + \frac{\bar{\mathbb{E}}_p[C_{n,2}] \varepsilon_{n,\alpha}}{(1 - \alpha + \varepsilon_{n,\alpha})^2}.$$

It follows,

$$\inf_{|\alpha' - \alpha| \geq \varepsilon_{n,\alpha}} |\Phi_{n,p}(\alpha')| \geq \frac{\bar{\mathbb{E}}_p[C_{n,2}] \varepsilon_{n,\alpha}}{(1 - \alpha + \varepsilon_{n,\alpha})^2} - |\Phi_{n,p}(\alpha)|.$$

By the equation (A.1) and the Lemma E.3,

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \inf_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \inf_{|\alpha' - \alpha| \geq \varepsilon_{n,\alpha}} |\Phi_{n,p}(\alpha')| \\ & \geq \liminf_{n \rightarrow \infty} \inf_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \frac{5L'\Gamma(1 - \alpha/2) \cdot n^{\alpha/2} \sqrt{\log(n)}}{1 - \alpha} \left[\frac{Ln^{\alpha/2}\Gamma(1 - \alpha) \cdot \varepsilon_{n,\alpha}}{5L'\Gamma(1 - \alpha/2) \sqrt{\log(n)}} - 1 \right] \end{aligned}$$

which will be strictly greater than zero whenever $\varepsilon_{n,\alpha} \geq \frac{5L'}{L} n^{-\alpha/2} \sqrt{\log(n)}$. Since by construction $\Phi_{n,p}(\alpha_n^*(p)) = 0$, it is necessary that

$$\limsup_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \frac{|\alpha_n^*(p) - \alpha|}{n^{-\alpha/2} \sqrt{\log(n)}} \leq \frac{5L'}{L}.$$

□

PROPOSITION A.2. For every $\varepsilon > 0$ there exists a constant B , depending only on ε and \mathbb{A} , such that

$$\limsup_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p \left(|\hat{\alpha}_n - \alpha_n^*(p)| > Bn^{-\alpha/2} \right) \leq \varepsilon.$$

Proof. For all $0 < \delta \leq 1/2$ and all p , let define the events $\Omega_n(p; \delta) := \{\mathbf{X} : |C_{n,2} - \bar{\mathbb{E}}_p[C_{n,2}]| \leq \delta \bar{\mathbb{E}}_p[C_{n,2}]\}$. Then, by Chebychev's inequality and because $M_{n,2} \leq C_{n,2}$ almost-surely, the Lemma A.2 guarantees that

$$\bar{\mathbb{P}}_p(\Omega_n(p; \delta)^c) \leq \frac{\text{var}_p(C_{n,2})}{\delta^2 \bar{\mathbb{E}}_p[C_{n,2}]^2} \leq \frac{2\bar{\mathbb{E}}_p[M_{n,2}]}{\delta^2 \bar{\mathbb{E}}_p[C_{n,2}]^2} \leq \frac{2}{\delta^2 \bar{\mathbb{E}}_p[C_{n,2}]}.$$

We deduce from the last display and the Lemma E.3 that whenever $\delta^2 n^\alpha \rightarrow \infty$ as $n \rightarrow \infty$,

$$\lim_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p(\Omega_n(p; \delta)^c) = 0.$$

We define other events $\Omega'_n(p; \varepsilon) := \{\mathbf{X} : |\varphi(\alpha_n^*(p))|^2 \leq A_p \bar{\mathbb{E}}_p[C_{n,1}] \varepsilon^{-1}\}$, where we let $A_p := \max\{1, \alpha_n^*(p)/(1 - \alpha_n^*(p))^2\}$. We remark that $\bar{\mathbb{E}}_p[\varphi(\alpha_n^*(p))] = \Phi_p(\alpha_n^*(p)) = 0$. Hence by Chebyshev's inequality and the Proposition A.3, for all p ,

$$\bar{\mathbb{P}}_p(\Omega'_n(p; \varepsilon)^c) \leq \frac{\overline{\text{var}}_p(\varphi(\alpha_n^*(p)))}{A_p \bar{\mathbb{E}}_p[C_{n,1}] \varepsilon^{-1}} \leq \varepsilon.$$

We finish the proof by showing that on the event $\Omega_n(p; \delta) \cap \Omega'_n(p; \varepsilon)$, it must be that $|\hat{\alpha}_n - \alpha_n^*(p)| \leq x_n$ for some $x_n \geq 0$ to be determined at the end of the proof. In fact, observe that $\varphi'(\alpha) = \sum_{k=1}^{n-1} \frac{k}{(k-\alpha)^2} C_{n,k+1} \geq C_{n,2}/(1-\alpha)^2$. Therefore φ is strictly monotone increasing on $(0, 1)$ when we are on the event $\Omega_n(p; \delta)$. Consequently, for all $\alpha \geq \alpha_n^* + x_n$, by a Taylor expansion of φ [remarking that $\varphi'' \geq 0$ on $\Omega_n(p; \delta)$ with the same argument as before]

$$\varphi(\alpha) \geq \varphi(\alpha_n^*(p) + x_n) \geq \varphi(\alpha_n^*(p)) + \varphi'(\alpha_n^*(p))x_n \geq \varphi(\alpha_n^*(p)) + \frac{C_{n,2}}{(1 - \alpha_n^*(p))^2} x_n;$$

Similarly for all $\alpha \leq \alpha_n^* - x_n$,

$$\varphi(\alpha) \leq \varphi(\alpha_n^*(p)) - \frac{C_{n,2}}{(1 - \alpha_n^*(p) + x_n)^2} x_n.$$

It follows that on the event $\Omega_n(p; \delta) \cap \Omega'_n(p; \varepsilon)$, for all α such that $|\alpha - \alpha_n^*(p)| \geq x_n$ we must have

$$|\varphi(\alpha)| \geq \frac{(1 - \delta) \bar{\mathbb{E}}_p[C_{n,2}]}{(1 - \alpha_n^*(p) + x_n)^2} x_n - \sqrt{\frac{A_p \bar{\mathbb{E}}_p[C_{n,1}]}{\varepsilon}}.$$

By choosing $x_n = Bn^{-\alpha/2}$ for a sufficiently large constant B , the rhs of the last display will be strictly greater than 0 uniformly for all $p \in \mathcal{P}(\alpha, L, L')$ for all $(\alpha, L, L') \in \mathbb{A}$, because of Proposition A.1 [which guarantees that A_p remains bounded for n large enough] and because of Lemma E.3 [which controls $\bar{\mathbb{E}}_p[C_{n,k}]$ uniformly in p]. But, by construction $\varphi(\hat{\alpha}_n) = 0$, so on $\Omega_n(p; \delta) \cap \Omega'_n(p; \varepsilon)$ it is necessary that $|\hat{\alpha}_n - \alpha_n^*(p)| \leq Bn^{-\alpha/2}$. \square

LEMMA A.1. *For all $0 < \underline{\alpha} < \bar{\alpha} < 1$, the following uniform limit results are true:*

$$\lim_{n \rightarrow \infty} \sup_{(\alpha, L, L') \in \mathbb{A}} \left| \frac{\int_0^1 L' x^{-\alpha/2} \sqrt{\log(e/x)} \frac{1-(1-x)^n}{x} dx}{\frac{L' \Gamma(1-\alpha/2)}{\alpha/2} \cdot n^{\alpha/2} \sqrt{\log(n)}} - 1 \right| = 0.$$

and

$$\lim_{n \rightarrow \infty} \sup_{(\alpha, L, L') \in \mathbb{A}} \left| \frac{n \int_0^1 L' x^{-\alpha/2} \sqrt{\log(e/x)} (1-x)^{n-1} dx}{L' \Gamma(1-\alpha/2) \cdot n^{\alpha/2} \sqrt{\log(n)}} - 1 \right| = 0.$$

Proof. We prove the result for the first integral which is the most complicated. The proof for the second integral is similar. By a simple change of variable

$$\int_0^1 L' x^{-\alpha/2} \sqrt{\log(e/x)} \frac{1-(1-x)^n}{x} dx = L' n^{\alpha/2} \sqrt{\log(en)} \int_0^n \sqrt{\frac{\log(en/y)}{\log(en)}} \frac{1-(1-y/n)^n}{y^{1+\alpha/2}} dy.$$

We decompose the integral in the rhs of the last display as $I_1 + I_2 + I_3 + I_4$, where

$$\begin{aligned} I_1 &:= \int_0^\infty \frac{1 - e^{-y}}{y^{1+\alpha/2}} dy \\ I_2 &:= - \int_n^\infty \frac{1 - e^{-y}}{y^{1+\alpha/2}} dy \\ I_3 &:= \int_0^n \left[\frac{1 - (1 - y/n)^n}{y^{1+\alpha/2}} - \frac{1 - e^{-y}}{y^{1+\alpha/2}} \right] dy \\ I_4 &:= \int_0^n \left[\sqrt{\frac{\log(en/y)}{\log(en)}} - 1 \right] \frac{1 - (1 - y/n)^n}{y^{1+\alpha/2}} dy. \end{aligned}$$

By direct computations it is found that $I_1 = 2\Gamma(1 - \alpha/2)/\alpha$. Then,

$$|I_2| \leq \int_n^\infty \frac{1}{y^{1+\alpha/2}} dy = \frac{2n^{-\alpha/2}}{\alpha}.$$

Regarding I_3 , we can rewrite,

$$I_3 = \int_0^n \frac{e^{-y}[1 - e^{n \log(1-y/n)+y}]}{y^{1+\alpha/2}} dy$$

Since $n \log(1 - y/n) \leq y$ for all $0 < y < n$, we see that $I_3 \geq 0$. Furthermore, $n \log(1 - y/n) + y \geq \frac{-y}{1-y/n} + y = \frac{-y^2/n}{1-y/n}$. This establishes that

$$\begin{aligned} |I_3| &\leq \frac{1}{n} \int_0^{n/2} \frac{y^{1-\alpha/2} e^{-y}}{1 - y/n} dy + \int_{n/2}^n \frac{e^{-y}}{y^{1+\alpha/2}} dy \\ &\leq \frac{2}{n} \int_0^{n/2} y^{1-\alpha/2} e^{-y} dy + e^{-n/2} \int_{n/2}^\infty \frac{1}{y^{1+\alpha/2}} dy \\ &\leq \frac{2\Gamma(2 - \alpha/2)}{n} + \frac{2\alpha/2+1 e^{-n/2}}{\alpha n^{\alpha/2}}. \end{aligned}$$

Now regarding I_4 ,

$$\begin{aligned} I_4 &= \frac{1}{\sqrt{\log(en)}} \int_0^n \left[\sqrt{\log(en) - \log(y)} - \sqrt{\log(en)} \right] \frac{1 - (1 - y/n)^n}{y^{1+\alpha/2}} dy \\ &= \frac{1}{\sqrt{\log(en)}} \int_0^n \frac{-\log(y)}{\sqrt{\log(en/y)} + \sqrt{\log(en)}} \frac{1 - (1 - y/n)^n}{y^{1+\alpha/2}} dy. \end{aligned}$$

We deduce that,

$$\begin{aligned} |I_4| &\leq \frac{1}{\log(en)} \int_0^1 \log\left(\frac{1}{y}\right) \frac{1 - (1 - y/n)^n}{y^{1+\alpha/2}} dy + \frac{1}{\log(en)} \int_1^n \frac{\log(y)}{y^{1+\alpha/2}} dy \\ &\leq \frac{1}{(1 - 1/n) \log(en)} \int_0^1 \log\left(\frac{1}{y}\right) y^{-\alpha/2} dy + \frac{1}{\log(en)} \int_1^n \frac{\log(y)}{y^{1+\alpha/2}} dy, \end{aligned}$$

where the second line follows because $1 - (1 - y/n)^n = 1 - \exp[n \log(1 - y/n)] \leq 1 - \exp[-y/(1 - y/n)] \leq y/[1 - y/n]$ for all $y \in [0, 1]$. Finally,

$$|I_4| \leq \frac{1}{\log(en)} \frac{1}{1 - 1/n} \frac{1}{(1 - \alpha/2)^2} + \frac{1}{\log(en)} \frac{4}{\alpha^2}.$$

Finally, the conclusion follows by gathering all the previous estimates. \square

LEMMA A.2. *For all p , all $n \geq 1$ and all $1 \leq k \leq n$,*

$$\overline{\text{var}}_p(C_{n,k}) \leq k \cdot \bar{\mathbb{E}}_p[M_{n,k}].$$

Proof. We bound the variance of $C_{n,k}$ using an Efron-Stein argument. By the Theorem 3.1 in [8], it is the case that

$$\overline{\text{var}}_p(C_{n,k}) \leq \sum_{i=1}^n \inf_{Z_i} \bar{\mathbb{E}}_p[(C_{n,k} - Z_i)^2]$$

where the infimum is over all $(X_1, \dots, X_{i-1}, X_{i+1}, X_n)$ -measurable random variables. In particular, defining $Y_{n,j}^{(i)} := Y_{n,j} - \mathbf{1}_{X_i=j}$, and $C_{n,k}^{(i)} := \sum_{j \geq 1}^n \mathbf{1}_{Y_{n,j}^{(i)} \geq k}$, we see that $C_{n,k}^{(i)}$ does not depend on X_i . Hence the last display entails that

$$\overline{\text{var}}_p(C_{n,k}) \leq \sum_{i=1}^n \bar{\mathbb{E}}_p[(C_{n,k} - C_{n,k}^{(i)})^2].$$

But,

$$\begin{aligned} C_{n,k} - C_{n,k}^{(i)} &= \sum_{j \geq 1} (\mathbf{1}_{Y_{n,j} \geq k} - \mathbf{1}_{Y_{n,j}^{(i)} \geq k}) \\ &= \sum_{j \geq 1} \mathbf{1}_{Y_{n,j}=k} \mathbf{1}_{X_i=j}, \end{aligned}$$

where the second line follows because $Y_{n,j} \geq Y_{n,j}^{(i)} \geq Y_{n,j} - 1$ almost-surely, hence $\mathbf{1}_{Y_{n,j} \geq k} - \mathbf{1}_{Y_{n,j}^{(i)} \geq k}$ is either zero or one, and it is one iff $Y_{n,j} = k$ and $Y_{n,j}^{(i)} = k - 1$; that is iff $Y_{n,j} = k$ and $X_i = j$. It follows,

$$\begin{aligned} \overline{\text{var}}_p(C_{n,k}) &\leq \sum_{i=1}^n \bar{\mathbb{E}}_p \left[\sum_{j \geq 1} \mathbf{1}_{Y_{n,j}=k} \mathbf{1}_{X_i=j} \sum_{j' \geq 1} \mathbf{1}_{Y_{n,j'}=k} \mathbf{1}_{X_i=j'} \right] \\ &= \sum_{j \geq 1} \bar{\mathbb{E}}_p \left[\mathbf{1}_{Y_{n,j}=k} \sum_{i=1}^n \mathbf{1}_{X_i=j} \right] \\ &= \sum_{j \geq 1} \bar{\mathbb{E}}_p [\mathbf{1}_{Y_{n,j}=k} Y_{n,j}] \\ &= k \bar{\mathbb{E}}_p \left[\sum_{j \geq 1} \mathbf{1}_{Y_{n,j}=k} \right] \end{aligned}$$

$$= k \cdot \mathbb{E}_p[M_{n,k}].$$

□

PROPOSITION A.3. *The following is true for all $\alpha \in (0, 1)$ and all p*

$$\overline{\text{var}}_p(\varphi_n(\alpha)) \leq \max \left\{ 1, \left(\frac{\alpha}{1-\alpha} \right)^2 \right\} \mathbb{E}_p[C_{n,1}].$$

Proof. We bound the variance of $\varphi_n(\alpha')$ using the same Efron-Stein argument as in the proof of Lemma A.2. For simplicity we write $\varphi \equiv \varphi_n(\alpha)$; and we let $C_{n,k}^{(i)}$ having the same definition as in Lemma A.2. Then defining $\varphi_i := \sum_{k=1}^{n-1} \frac{\alpha}{k-\alpha} C_{n,k+1}^{(i)} - C_{n,1}^{(i)}$, following the arguments in Lemma A.2

$$\overline{\text{var}}_p(\varphi) \leq \sum_{i=1}^n \mathbb{E}_p[(\varphi - \varphi_i)^2]$$

and

$$\begin{aligned} (\varphi - \varphi_i)^2 &= \left(\sum_{k=1}^{n-1} \frac{\alpha'}{k-\alpha'} (C_{n,k+1} - C_{n,k+1}^{(i)}) - (C_{n,1} - C_{n,1}^{(i)}) \right)^2 \\ &= \left\{ \sum_{j \geq 1} \mathbf{1}_{X_i=j} \left(\sum_{k=1}^{n-1} \frac{\alpha'}{k-\alpha'} \mathbf{1}_{Y_{n,j}=k+1} - \mathbf{1}_{Y_{n,j}=1} \right) \right\}^2. \end{aligned}$$

It follows,

$$\begin{aligned} \sum_{i=1}^n (\varphi - \varphi_i)^2 &= \sum_{j \geq 1} \left(\sum_{i=1}^n \mathbf{1}_{X_i=j} \right) \left(\sum_{k=1}^{n-1} \frac{\alpha'}{k-\alpha'} \mathbf{1}_{Y_{n,j}=k+1} - \mathbf{1}_{Y_{n,j}=1} \right)^2 \\ &= \sum_{j \geq 1} Y_{n,j} \left(\sum_{k=1}^{n-1} \frac{\alpha'}{k-\alpha'} \mathbf{1}_{Y_{n,j}=k+1} - \mathbf{1}_{Y_{n,j}=1} \right)^2 \\ &= \sum_{j \geq 1} \mathbf{1}_{Y_{n,j}=1} + \sum_{j \geq 1} Y_{n,j} \mathbf{1}_{Y_{n,j} \neq 1} \left(\sum_{k=1}^{n-1} \frac{\alpha'}{k-\alpha'} \mathbf{1}_{Y_{n,j}=k+1} \right)^2 \\ &= M_{n,1} + \sum_{k=1}^{n-1} \sum_{k'=1}^{n-1} \frac{\alpha'}{k-\alpha'} \frac{\alpha'}{k'-\alpha'} \sum_{j \geq 1} Y_{n,j} \mathbf{1}_{Y_{n,j} \neq 1} \mathbf{1}_{Y_{n,j}=k+1} \mathbf{1}_{Y_{n,j}=k'+1} \\ &= M_{n,1} + \sum_{k=1}^{n-1} \left(\frac{\alpha'}{k-\alpha'} \right)^2 \sum_{j \geq 1} Y_{n,j} \mathbf{1}_{Y_{n,j}=k+1}. \end{aligned}$$

Hence we have shown,

$$\sum_{i=1}^n \mathbb{E}_p[(\varphi - \varphi_i)^2] = \mathbb{E}_p[M_{n,1}] + \sum_{k=1}^{n-1} \left(\frac{\alpha'}{k-\alpha'} \right)^2 (k+1) \mathbb{E}_p[M_{n,k+1}]$$

$$\leq \max \left\{ 1, 2 \left(\frac{\alpha'}{1 - \alpha'} \right)^2 \right\} \bar{\mathbb{E}}_p[C_{n,1}],$$

where the second line follows because $(k+1)/(k-\alpha')^2 \leq 2/(1-\alpha')^2$ for all $k \geq 1$ and because $\sum_{\ell=1}^n M_{n,\ell} = C_{n,1}$. \square

A.3. Proof of Theorem 3

We want to bound the probability of the event

$$E_n := \left\{ \ell_\alpha(\hat{U}_{n,\lambda n}, U_{n,\lambda n}) \leq D \min \left(1, \frac{\log(n) \log(\lambda)^2}{n^\alpha} \right) \right\}$$

for $(\alpha, L, L') \in \mathbb{A}$ and $p \in \mathcal{P}(\alpha, L, L')$ be arbitrary. First consider the case where $\log(\lambda) > \frac{2Cn^{\alpha/2}}{\sqrt{\log(n)}}$. Then we see that $\log(\lambda) > \frac{Cn^{\hat{\alpha}_n/2}}{\sqrt{\log(n)}} \cdot 2n^{(\alpha - \hat{\alpha}_n)/2}$. We deduce that on the event that $|\hat{\alpha}_n - \alpha| \leq \frac{2\log(2)}{\log(n)}$ is must be that $\hat{U}_{n,m} = 0$, that is $\ell_\alpha(\hat{U}_{n,\lambda n}, U_{n,\lambda n}) = \frac{U_{n,\lambda n}^2}{(\lambda n)^{2\alpha}} \leq \frac{K_{n+\lambda n}^2}{(\lambda n)^{2\alpha}}$. It follows for $D' = D \min(1, 4C^2)$,

$$\begin{aligned} \bar{\mathbb{P}}_p(E_n^c) &\leq \bar{\mathbb{P}}_p(K_{n+\lambda n} \geq (\lambda n)^\alpha \sqrt{D'}) + \bar{\mathbb{P}}_p\left(|\hat{\alpha}_n - \alpha| > \frac{2\log(2)}{\log(n)}\right) \\ &\leq \frac{\overline{\text{var}}_p(K_{n+\lambda n})}{(\lambda n)^{2\alpha} D'} + \bar{\mathbb{P}}_p\left(|\hat{\alpha}_n - \alpha| > \frac{2\log(2)}{\log(n)}\right). \end{aligned}$$

The first term of the rhs of the last display can be made smaller than $\varepsilon/2$ by choosing D (and thus D') large enough, in virtue of Lemma A.2 and E.3 [remarking that $K_n = C_{n,1}$]. The second term of the rhs of the last display is controlled via the Theorem 2 and will be smaller than $\varepsilon/2$ whenever n gets large enough.

Now consider the case where $\log(\lambda) \leq 2Cn^{\alpha/2}/\sqrt{\log(n)}$. We define the event $A_n := \{\log(\lambda) \leq Cn^{\hat{\alpha}_n/2}/\sqrt{\log(n)}\}$, such that $\hat{U}_{n,m} = 0$ on the complement A_n^c . With the exact same argument as above

$$\bar{\mathbb{P}}_p\left(E_n^c \cap A_n^c\right) \leq \bar{\mathbb{P}}_p(K_{n+\lambda n} \geq (\lambda n)^\alpha \sqrt{D'}) \leq \frac{\varepsilon}{3}$$

by choosing D sufficiently large. While if we are on the event A_n , it is the case that $\hat{U}_{n,\lambda n} = K_n((1+\lambda)^{\hat{\alpha}_n} - 1)$. Equivalently, it is the case that $\hat{U}_{n,\lambda n} - U_{n,\lambda n} = K_n(1+\lambda)^{\hat{\alpha}_n} - K_{n+\lambda n}$. Define another event $B_n := \{|\hat{\alpha}_n - \alpha|^2 \leq B \log(n) n^{-\alpha}\}$ for a large enough constant $B > 0$. By choosing B large enough, by the Theorem 2 it will be that $\bar{\mathbb{P}}_p(B_n^c) \leq \varepsilon/3$. So it is enough to analyze what happens on the event $A_n \cap B_n$. But on $A_n \cap B_n$ for n large enough,

$$\begin{aligned} |\hat{U}_{n,\lambda n} - U_{n,\lambda n}| &= |K_n(1+\lambda)^\alpha [(1+\lambda)^{\hat{\alpha}_n - \alpha} - 1] + [K_n(1+\lambda)^\alpha - K_{n+\lambda n}]| \\ &\leq 2e^{2C\sqrt{B}} \cdot K_n(1+\lambda)^\alpha |\hat{\alpha}_n - \alpha| \log(1+\lambda) + |K_n(1+\lambda)^\alpha - K_{n+\lambda n}| \\ &\lesssim \frac{K_n \lambda^\alpha \log(\lambda) \sqrt{\log(n)}}{n^{\alpha/2}} + |\bar{\mathbb{E}}_p[K_n](1+\lambda)^\alpha - \bar{\mathbb{E}}_p[K_{n+\lambda n}]| \\ &\quad + \lambda^\alpha |K_n - \bar{\mathbb{E}}_p[K_n]| + |K_{n+\lambda n} - \bar{\mathbb{E}}_p[K_{n+\lambda n}]|. \end{aligned}$$

Again remarking that $K_n = C_{n,1}$, we deduce from Lemma E.3 and a couple applications of Chebychev's inequality that for D large enough,

$$\bar{\mathbb{P}}_p(E_n^c \cap A_n \cap B_n) \leq \frac{\varepsilon}{3}.$$

Combining all the estimates above, we have demonstrated that in any cases $\bar{\mathbb{P}}_p(E_n^c) \leq \varepsilon$ by choosing D sufficiently large. Inspection of the proof shows that this last estimate holds uniformly over $(\alpha, L, L') \in \mathbb{A}$ and $p \in \mathcal{P}(\alpha, L, L')$.

A.4. Proof of Proposition 2

The proof of the Proposition is nearly identical to the proof of the Theorem 2. In particular, the Proposition A.2 is unchanged, and to finish the proof it is enough to adapt the proof of the Proposition A.1 to establish that

$$\limsup_{n \rightarrow \infty} n^\beta |\alpha_n^*(p) - \alpha| \leq C$$

for some $C > 0$ if $\bar{F}_p(x) = Lx^{-\alpha} + O(x^{-\beta})$ as $x \rightarrow 0$. This can be established by adapting the proofs of Proposition A.1 and Lemma A.1 line by line.

Appendix B: Minimax lower bound for estimating the tail-index

This appendix contains the proof of Theorem 5.

B.1. Guidelines and construction of the hypotheses

The lower bound is obtained from the result of Theorem 4. For $0 < \alpha_0 < \alpha_1 < 1$ to be chosen accordingly, we consider the priors ν_0 and ν_1 consisting on PK distributions with respective Lévy measures

$$\rho_0(dx) = \frac{\alpha_0 x^{-1-\alpha_0} dx}{\Gamma(1-\alpha_0)}, \quad \rho_1(dx) = \frac{\alpha_1 x^{-1-\alpha_1} dx}{\Gamma(1-\alpha_1)}.$$

Then, we wish to apply Theorem 4 to the functional $\delta(\Pi, p) \equiv \delta(p) = \lim_{x \rightarrow 0} \frac{\log \bar{F}_p(x)}{-\log(x)}$. It is a well-known fact (Kingman [38], Gneden et al. [26]) that $\nu_0(\delta(P) = \alpha_0) = 1$ and $\nu_1(\delta(P) = \alpha_1) = 1$. Hence, taking $2\varepsilon = \alpha_1 - \alpha_0$, we can take $\kappa = 0$ and it suffices to show that for suitable choice of constants we have

$$\min_{j=0,1} \nu_j(P \in \mathcal{P}^*) \geq 1 - \xi, \tag{B.1}$$

where we let,

$$\mathcal{P}^* := \bigcup_{(\alpha, L, L') \in \mathbb{A}} \mathcal{P}(\alpha, L, L'),$$

and also,

$$\text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n) \leq \frac{1}{2}. \tag{B.2}$$

Combining the results of the two next sections shows that this will be the case if \bar{L} is sufficiently large, if $\alpha_0 = \underline{\alpha}$, and if for $c > 0$ small enough but constant we take

$$\varepsilon = \frac{c}{n^{\alpha/2} \log(n)}.$$

Summarizing, (B.1) and (B.2) combined with the Theorem 4 guarantee that whenever n is large enough

$$\inf_{\hat{\alpha}_n} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p \left(|\hat{\alpha}_n - \alpha| \geq \varepsilon \right) \geq \frac{1}{4}.$$

The final result follows because $\varepsilon \geq \frac{c}{n^{\alpha/2} \log(n)}$ for all $(\alpha, L, L') \in \mathbb{A}$.

B.2. Bounding the KL divergence between hypotheses

In this section we establish that for the choice of hypotheses and constants made in Section B.1, it is true that (B.2) holds.

PROPOSITION B.1. *There is a constant $C > 0$ depending only on α_0 such that,*

$$\text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n) \leq C \varepsilon^2 n^{\alpha_0} \log(n)^2.$$

Proof. By direct computations, we obtain that $\mu_0(u, \ell) = \frac{\alpha_0 \Gamma(\ell - \alpha_0)}{\Gamma(1 - \alpha_0)} u^{-\ell + \alpha_0}$ and $\mu_1(u, \ell) = \frac{\alpha_0 \Gamma(\ell - \alpha_1)}{\Gamma(1 - \alpha_0)} u^{-\ell + \alpha_1}$. It follows,

$$\begin{aligned} -\log \frac{\mu_1(u, \ell)}{\mu_0(u, \ell)} + \frac{\mu_1(u, \ell)}{\mu_0(u, \ell)} - 1 &= -\log \frac{\Gamma(\ell - \alpha_1)}{\Gamma(\ell - \alpha_0)} + \frac{\Gamma(\ell - \alpha_1)}{\Gamma(\ell - \alpha_0)} - 1 \\ &\quad + \frac{\Gamma(\ell - \alpha_1)}{\Gamma(\ell - \alpha_0)} \left(u^{\alpha_1 - \alpha_0} - (\alpha_1 - \alpha_0) \log u - 1 \right) \\ &\quad + \left\{ \frac{\Gamma(\ell - \alpha_1)}{\Gamma(\ell - \alpha_0)} - 1 \right\} (\alpha_1 - \alpha_0) \log u. \end{aligned}$$

We let ψ denote the digamma function; that is the derivative of the map $z \mapsto \log \Gamma(z)$. Then, by a Taylor expansion argument

$$\begin{aligned} \sup_{\ell=1, \dots, n} \left| \frac{\Gamma(\ell - \alpha_1)}{\Gamma(\ell - \alpha_0)} - 1 \right| &\leq \sup_{\ell=1, \dots, n} \sup_{\alpha \in (\alpha_0, \alpha_1)} \left| e^{\psi(\ell - \alpha)(\alpha_0 - \alpha_1)} - 1 \right| \\ &\leq (1 + o(1)) \log(n) |\alpha_0 - \alpha_1|, \end{aligned}$$

where the second line follows because $\psi(z) \sim \log(z)$ as $z \rightarrow \infty$ and because by construction $\log(n)(\alpha_0 - \alpha_1) = o(1)$. Similarly,

$$0 \leq u^{\alpha_1 - \alpha_0} - (\alpha_1 - \alpha_0) \log u - 1 \leq (\alpha_1 - \alpha_0)^2 \log(u)^2 \cdot \max(1, u^{\alpha_1 - \alpha_0}).$$

Therefore,

$$\begin{aligned} 0 &\leq -\log \frac{\mu_1(u, \ell)}{\mu_0(u, \ell)} + \frac{\mu_1(u, \ell)}{\mu_0(u, \ell)} - 1 \\ &\leq (1 + o(1)) (\alpha_0 - \alpha_1)^2 \left(\frac{\log(n)^2}{2} + \log(u)^2 \cdot \max(1, u^{2\varepsilon}) + \log(n) \cdot |\log(u)| \right) \end{aligned}$$

Thus by Proposition 1, and by the fact that $\sum_{\ell=1}^n M_{n, \ell} = K_n$,

$$\begin{aligned} \text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n) &\leq 4\varepsilon^2 (1 + o(1)) \left(\frac{\mathbb{E}_{\rho_0}[K_n] \log(n)^2}{2} \right. \\ &\quad \left. + \mathbb{E}_{\rho_0}[K_n \log(U_n)^2 \cdot \max(1, U_n^{2\varepsilon})] + \log(n) \cdot \mathbb{E}_{\rho_0}[K_n |\log(U_n)|] \right). \end{aligned}$$

It is well known that under ρ_0 it holds $\mathbb{E}_{\rho_0}[K_n] \asymp n^{\alpha_0}$; see for instance Pitman [48]. Furthermore,

$$\mathbb{E}_{\rho_0}[K_n \log(U_n)^2 \cdot \max(1, U_n^{2\varepsilon})] \leq \mathbb{E}_{\rho_0}[K_n \log(U_n)^2] + \mathbb{E}_{\rho_0}[K_n \cdot U_n^{2\varepsilon} \log(U_n)^2],$$

and by Cauchy-Schwarz' inequality

$$\mathbb{E}_{\rho_0}[K_n \cdot U_n^{2\varepsilon} \log(U_n)^2] \leq \mathbb{E}_{\rho_0}[K_n^2]^{1/2} \mathbb{E}_{\rho_0}[U_n^{4\varepsilon} \log(U_n)^4]^{1/2}.$$

Recall that $U_n \stackrel{d}{=} G_n/T$ where $G_n \sim \text{Gamma}(n, 1)$ is independent of T , so

$$\begin{aligned} \mathbb{E}_{\rho_0}[U_n^{4\varepsilon} \log(U_n)^4] &= \mathbb{E}_{\rho_0}[G_n^{4\varepsilon} T^{-4\varepsilon} \log(G_n/T)^4] \\ &\leq 8\mathbb{E}_{\rho_0}[G_n^{4\varepsilon} T^{-4\varepsilon} \{\log(G_n)^4 + \log(T)^4\}] \\ &= 8\mathbb{E}_{\rho_0}[G_n^{4\varepsilon} \log(G_n)^4] \mathbb{E}_{\rho_0}\left[\frac{1}{T^{4\varepsilon}}\right] + 8\mathbb{E}_{\rho_0}[G_n^{4\varepsilon}] \mathbb{E}_{\rho_0}\left[\frac{\log(T)^4}{T^{4\varepsilon}}\right]. \end{aligned}$$

Under ρ_0 , the random variable T has a α_0 -stable distribution (Pitman [48]). By a classical analysis of the stable distributions (Samorodnitsky and Taqqu [51]), $T^{-4\varepsilon}$ and $T^{-4\varepsilon} \log(T)^4$ have finite expectation, and they can be bounded by a constant depending only on α_0 as $\varepsilon = o(1)$. Since, $\varepsilon = o(n^{-\alpha_0/2})$ we obtain

$$\mathbb{E}_{\rho_0}[U_n^{4\varepsilon} \log(U_n)^4] \lesssim \log(n)^4.$$

Finally, $\mathbb{E}_{\rho_0}[K_n^2] = O(n^{2\alpha_0})$ (Pitman [48]), so

$$\mathbb{E}_{\rho_0}[K_n \log(U_n)^2 \cdot \max(1, U_n^{2\varepsilon})] \lesssim n^{\alpha_0} \log(n)^2.$$

The remaining term $\mathbb{E}_{\rho_0}[K_n |\log(U_n)|]$ can be bounded using the same arguments, giving the conclusion. \square

B.3. Priors generate regular probability vectors with large probability

In this section we establish that for the choice of hypotheses and constants made in Section B.1, it is true that equation (B.1) holds. We proceed by proving that with large ν_j -probability, there exists L, L'

such that $p \in \mathcal{P}(\alpha_j, L, L')$. First, observe that by direct computations, we have,

$$\bar{\rho}_j(x) = \frac{\alpha_0/\alpha_j}{\Gamma(1-\alpha_0)} x^{-\alpha_j}. \quad (\text{B.3})$$

We define two random variables $V_j := \lim_{x \rightarrow 0} x^{\alpha_j} \bar{\rho}_j(xT)$. Then, the Proposition B.2 below guarantees that for all $\epsilon > 0$ there exists L' such that $\min_{j=1,2} \nu_j \{P \in \mathcal{P}(\alpha_j, V_j, L')\} \geq 1 - \epsilon$.

PROPOSITION B.2. *For all $j = 0, 1$ and for all $\epsilon > 0$ there exists a constant $L' > 0$ depending only on ϵ and α_j such that with ν_j -probability at least $1 - \epsilon$*

$$\sup_{x \in (0,1)} \frac{|x^{\alpha_j} \bar{F}_P(x) - V_j|}{\sqrt{x^{\alpha_j} \log(e/x)}} \leq L'.$$

Proof. Let $j = 0$ or $j = 1$. We remark that $V_j = x^{\alpha_j} \bar{\rho}_j(xT)$ by equation (B.3). Also, recall that $\bar{F}_P(x) \stackrel{d}{=} \sum_{i \geq 1} \mathbf{1}_{J_i > xT}$ conditional on $(J_i)_{i \geq 1}$. Then, in distribution

$$\begin{aligned} \sup_{x \in (0,1)} \frac{|x^{\alpha_j} \bar{F}_P(x) - V_j|}{\sqrt{x^{\alpha_j} \log(e/x)}} &\leq \sup_{x \in (0,1)} \sqrt{\frac{x^{\alpha_j}}{\log(e/x)}} \left| \sum_{i \geq 1} \mathbf{1}_{J_i > xT} - \bar{\rho}_j(xT) \right| \\ &= \sup_{x \in (0,T)} \sqrt{\frac{(x/T)^{\alpha_j}}{\log(eT/x)}} \left| \sum_{i \geq 1} \mathbf{1}_{J_i > x} - \bar{\rho}_j(x) \right|. \end{aligned}$$

Let $\bar{\rho}_j^{-1}(u) := \inf\{x > 0 : \bar{\rho}_j(x) < u\}$ and let $(\Gamma_i)_{i \geq 1}$ be a unit rate Poisson process on \mathbb{R}_+ . By (Ferguson and Klass [21]), it is the case that $(J_i)_{i \geq 1} \stackrel{d}{=} (\bar{\rho}_j^{-1}(\Gamma_i))_{i \geq 1}$. Furthermore, $\bar{\rho}_j^{-1}(x) > y \iff x < \bar{\rho}_j(y)$. Hence, in distribution,

$$\begin{aligned} \sup_{x \in (0,1)} \frac{|x^{\alpha_j} \bar{F}_P(x) - V_j|}{\sqrt{x^{\alpha_j} \log(e/x)}} &= \sup_{x \in (0,T)} \sqrt{\frac{(x/T)^{\alpha_j}}{\log(eT/x)}} \left| \sum_{i \geq 1} \mathbf{1}_{\bar{\rho}_j^{-1}(\Gamma_i) > x} - \bar{\rho}_j(x) \right| \\ &= \sup_{x \in (0,T)} \sqrt{\frac{(x/T)^{\alpha_j}}{\log(eT/x)}} \left| \sum_{i \geq 1} \mathbf{1}_{\Gamma_i < \bar{\rho}_j(x)} - \bar{\rho}_j(x) \right| \\ &= \sup_{y \in (\bar{\rho}_j(T), \infty)} \sqrt{\frac{(\bar{\rho}_j^{-1}(y)/T)^{\alpha_j}}{\log(eT/\bar{\rho}_j^{-1}(y))}} \left| \sum_{i \geq 1} \mathbf{1}_{\Gamma_i < y} - y \right|. \end{aligned}$$

We deduce from equation (B.3) that $\bar{\rho}_j^{-1}(y) = C_1(\alpha_j) y^{-1/\alpha_j}$, where $C_1(\alpha_j) := [(\alpha_0/\alpha_j)/\Gamma(1-\alpha_0)]^{1/\alpha_j}$. Also $T = \bar{\rho}_j^{-1}(\bar{\rho}_j(T))$; ie $\bar{\rho}_j^{-1}(y)/T = (y/\bar{\rho}_j(T))^{-1/\alpha_j}$. Since the function $z \mapsto z^{\alpha_j}/\log(e/z)$ is monotone increasing in $(0, \infty)$, we have on the event that $\bar{\rho}_j(T) \leq A$

$$\sup_{x \in (0,1)} \frac{|x^{\alpha_j} \bar{F}_P(x) - V_j|}{\sqrt{x^{\alpha_j} \log(e/x)}} \leq \sup_{y \in (0, \infty)} \sqrt{\frac{A/y}{\log[e(y/A)^{1/\alpha_j}]}} \left| \sum_{i \geq 1} \mathbf{1}_{\Gamma_i < y} - y \right|$$

$$\lesssim \sup_{y \in (0, \infty)} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y|}{\sqrt{\max\{1, y \log y\}}}$$

The random variable T has a α_j stable distribution, so it is always possible to choose A large enough so that the event $\bar{\rho}_j(T) \leq A$ occur with probability $1 - \epsilon/2$. Furthermore, by the Lemma B.1, for B sufficiently large the rhs of the last display will be smaller than B with probability at least $1 - \epsilon/2$. This concludes the proof. \square

LEMMA B.1. *Let $(\Gamma_i)_{i \geq 1}$ be an homogeneous Poisson process on the half real line. Then, for every $\epsilon > 0$ there exists $D > 0$ such that with probability at least $1 - \epsilon$*

$$\sup_{y \in (0, \infty)} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y|}{\sqrt{\max\{1, y \log y\}}} \leq D.$$

Proof. Let y_0 be the unique non-negative solution of $y \log y = 1$. Obviously,

$$\sup_{y \in (0, \infty)} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y|}{\sqrt{\max\{1, y \log y\}}} \leq \max \left\{ \sup_{y \in (0, y_0)} \left| \sum_{j \geq 1} \mathbf{1}_{\Gamma_j \leq y} - y \right|, \sup_{y \geq y_0} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y|}{\sqrt{y \log y}} \right\}.$$

By Doob's maximal inequality,

$$\mathbb{E} \left[\sup_{y \in (0, y_0)} \left| \sum_{j \geq 1} \mathbf{1}_{\Gamma_j \leq y} - y \right|^2 \right] \leq 4 \mathbb{E} \left[\left| \sum_{j \geq 1} \mathbf{1}_{\Gamma_j \leq y_0} - y_0 \right|^2 \right] = 4y_0.$$

Hence, by Chebychev's inequality we have that $\sup_{y \in (0, y_0)} |\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y| \leq C$ with probability at least $1 - 4y_0/C^2$. For $C > 0$ large enough this will happen with probability $1 - \epsilon/2$. For the other term, we decompose,

$$\begin{aligned} \sup_{y \geq y_0} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y|}{\sqrt{y \log y}} &= \sup_{k \geq 1} \sup_{y \in [ky_0, (k+1)y_0]} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y|}{\sqrt{y \log y}} \\ &\leq \sup_{k \geq 1} \frac{1}{\sqrt{ky_0 \log(ky_0)}} \sup_{y \in [ky_0, (k+1)y_0]} \left| \sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y \right|. \end{aligned}$$

Note that for all $y \in [ky_0, (k+1)y_0]$ we have,

$$\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y \leq \sum_{j \geq 1} \mathbf{1}_{\Gamma_j < (k+1)y_0} - ky_0 = \sum_{j \geq 1} \mathbf{1}_{\Gamma_j < (k+1)y_0} - (k+1)y_0 + y_0,$$

and,

$$\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y \geq \sum_{j \geq 1} \mathbf{1}_{\Gamma_j < k} - (k+1)y_0 = \sum_{j \geq 1} \mathbf{1}_{\Gamma_j < k} - ky_0 - y_0.$$

Therefore,

$$\sup_{y \in [ky_0, (k+1)y_0]} \left| \sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y \right| \leq y_0 + \max_{\ell=k, k+1} \left\{ \left| \sum_{j \geq 1} \mathbf{1}_{\Gamma_j < \ell} - \ell \right| \right\}.$$

It follows, because $(k+1)y_0 \log((k+1)y_0) \geq ky_0 \log(ky_0)$,

$$\sup_{y \geq y_0} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < y} - y|}{\sqrt{y \log y}} \leq \sqrt{\frac{y_0}{\log(y_0)}} + \sup_{k \geq 1} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < k} - k|}{\sqrt{ky_0 \log(ky_0)}}.$$

We finish the proof using Chernoff's bound on the Poisson distribution. Indeed, for any $t > 0$

$$\begin{aligned} \mathbb{P}\left\{\sup_{k \geq 1} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < k} - k|}{\sqrt{ky_0 \log(ky_0)}} > t\right\} &\leq \sum_{k \geq 1} \mathbb{P}\left\{\frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < k} - k|}{\sqrt{ky_0 \log(ky_0)}} > t\right\} \\ &\leq 2 \sum_{k \geq 1} \exp\left\{-\frac{t^2 ky_0 \log(ky_0)/2}{k + t\sqrt{ky_0 \log(ky_0)}}\right\} \\ &= 2 \sum_{k \geq 1} \exp\left\{-\frac{t^2 y_0 \log(ky_0)/2}{1 + ty_0 \sqrt{\log(ky_0)/(ky_0)}}\right\} \\ &\leq 2 \sum_{k \geq 1} \exp\left\{-\frac{t^2 y_0 \log(ky_0)/2}{1 + ty_0 e^{-1/2}}\right\}. \end{aligned}$$

Letting $g(t) := \frac{y_0 t^2/2}{1 + ty_0 e^{-1/2}}$, and remarking that $g(t) > 1$ if we take t sufficiently large,

$$\begin{aligned} \mathbb{P}\left\{\sup_{k \geq 1} \frac{|\sum_{j \geq 1} \mathbf{1}_{\Gamma_j < k} - k|}{\sqrt{ky_0 \log(ky_0)}} > t\right\} &\leq 2 \sum_{k \geq 1} (y_0 k)^{-g(t)} \\ &= 2y_0^{-g(t)} + 2 \sum_{k \geq 2} (y_0 k)^{-g(t)} \\ &\leq 2y_0^{-g(t)} + 2y_0^{-g(t)} \int_1^\infty x^{-g(t)} dx \\ &= \frac{2g(t)}{g(t) - 1} y_0^{-g(t)}. \end{aligned}$$

Observe that $y_0 > 1$ (this is obvious since for $y \leq 1$ it must be that $y \log(y) \leq 0$, hence any positive solution of $y \log(y) = 1$ must strictly greater than 1; by numerical computations, $y_0 \approx 1.76$). Then, the previous display can be made smaller than $\epsilon/2$ by taking t sufficiently large, which concludes the proof of the lemma. \square

Appendix C: Minimax lower bound for estimating the unseen

This appendix contains the proof of Theorem 6. We prove the theorem by showing that predicting the unseen over the class $\mathcal{P}(\alpha, L, L')$ is essentially a harder problem than estimating α . That is, we propose to reduce the problem to the estimation of α . In the whole proof, for $G_1, G_2 > 0$ to be chosen accordingly, we let

$$\varepsilon := \min\left\{\frac{G_1 \log(1 + \lambda)}{n^{\alpha/2} \log(n)}, G_2\right\}.$$

C.1. Reduction to the problem of estimating the tail-index

The first step of the reduction consists on noticing that estimating $U_{n,m}$ and K_{n+m} are equivalent problems. Indeed $U_{n,m} = K_{n+m} - K_n$ implies that for all estimator \hat{U} and all $p \in \mathcal{P}(\alpha, L, L')$, writing $\hat{K} = K_n + \hat{U}$,

$$\begin{aligned} \bar{\mathbb{P}}_p(\ell_\alpha(\hat{U}, U_{n,m}) > \varepsilon^2) &= \bar{\mathbb{P}}_p(|\hat{U} - U_{n,m}|^2 > (n\lambda)^2 \alpha \varepsilon^2) \\ &= \bar{\mathbb{P}}_p(|\hat{K} - K_{n+m}| > (n\lambda)^\alpha \varepsilon) \end{aligned}$$

The second step consists on showing that predicting K_{n+m} is always as much as hard as predicting its conditional expectation $\tilde{K}_{n,m}(p) := \bar{\mathbb{E}}_p[K_{n+m} | \Pi_n]$. This trick was also used by Orlitsky et al. [46] where they obtain a similar result in expectation using the convexity of the loss function and Jensen's inequality. Here we need a result that holds in probability, but the idea is similar. In fact, let consider the probability of the events $\Omega_1 := \{\Pi : \bar{\mathbb{E}}_p[|\hat{K} - K_{n+m}| | \Pi_n] > (n\lambda)^\alpha \varepsilon\}$ and $\Omega_2 := \{\Pi : |\hat{K} - K_{n+m}| > (n\lambda)^\alpha \varepsilon\}$. Then,

$$\bar{\mathbb{P}}_p(\Omega_1) = \bar{\mathbb{P}}_p(\Omega_1 \cap \Omega_2) + \bar{\mathbb{P}}_p(\Omega_1 \cap \Omega_2^c) = \bar{\mathbb{P}}_p(\Omega_1 \cap \Omega_2) \leq \bar{\mathbb{P}}_p(\Omega_2),$$

where the second equality follows because it must be the case that $\Omega_1 \cap \Omega_2^c = \emptyset$. In other words,

$$\bar{\mathbb{P}}_p(|\hat{K} - K_{n+m}| > (n\lambda)^\alpha \varepsilon) \geq \bar{\mathbb{P}}_p(\bar{\mathbb{E}}_p[|\hat{K} - K_{n+m}| | \Pi_n] > (n\lambda)^\alpha \varepsilon).$$

Furthermore, by Jensen's inequality and by the fact that \hat{K} depends only on Π_n , it is true that $\bar{\mathbb{E}}_p[|\hat{K} - K_{n+m}| | \Pi_n] \geq |\hat{K} - \bar{\mathbb{E}}_p[K_{n+m} | \Pi_n]|$, which entails

$$\bar{\mathbb{P}}_p(|\hat{K} - K_{n+m}| > (n\lambda)^\alpha \varepsilon) \geq \bar{\mathbb{P}}_p(|\hat{K} - \tilde{K}_{n,m}(p)| > (n\lambda)^\alpha \varepsilon).$$

The third step of the reduction trades $\tilde{K}_{n,m}(p)$ for its expectation $\bar{K}_{n,m}(p) := \bar{\mathbb{E}}_p[\tilde{K}_{n,m}(p)]$; observe that we also have $\bar{K}_{n,m}(p) = \bar{\mathbb{E}}_p[K_{n+m}]$. In the sequel we will also write $\bar{K}_n(p) := \bar{K}_{n,0}(p)$ for the expectation of K_n . In fact, rather than reducing to the problem of estimating $\bar{K}_{n,m}(p)$, it is more convenient to reduce to the problem of estimating $R_{n,m}(p) := \bar{K}_{n,m}(p)/\bar{K}_n(p)$. From \hat{K} we can build the estimator $\hat{R} := \hat{K}/K_n$ for $R_{n,m}(p)$, then by triangular inequality

$$\begin{aligned} |\hat{K} - \tilde{K}_{n,m}(p)| &= |K_n \hat{R} - K_n R_{n,m}(p) + K_n R_{n,m}(p) - \tilde{K}_{n,m}(p)| \\ &\geq K_n |\hat{R} - R_{n,m}(p)| - |K_n R_{n,m}(p) - \tilde{K}_{n,m}(p)|. \end{aligned}$$

So we have that $K_n |\hat{R} - R_{n,m}(p)| > 2(n\lambda)^\alpha \varepsilon$ and $|\tilde{K}_{n,m}(p) - K_n R_{n,m}(p)| \leq (n\lambda)^\alpha \varepsilon$ implies $|\hat{K} - \tilde{K}_{n,m}(p)| > (n\lambda)^\alpha \varepsilon$, which means that $\bar{\mathbb{P}}_p(\ell_\alpha(\hat{U}, U_{n,m}) > \varepsilon^2)$ is always at least

$$\bar{\mathbb{P}}_p\left(|\hat{R} - R_{n,m}(p)| > \frac{2(n\lambda)^\alpha \varepsilon}{K_n}, |\tilde{K}_{n,m}(p) - K_n R_{n,m}(p)| \leq (n\lambda)^\alpha \varepsilon\right). \quad (\text{C.1})$$

But by triangular inequality again,

$$\begin{aligned} |\tilde{K}_{n,m}(p) - K_n R_{n,m}(p)| &\leq |\tilde{K}_{n,m}(p) - \bar{K}_n(p) R_{n,m}(p)| + |\bar{K}_n(p) R_{n,m}(p) - K_n R_{n,m}(p)| \\ &= |\tilde{K}_{n,m}(p) - \bar{K}_{n,m}(p)| + R_{n,m}(p) |K_n - \bar{K}_n(p)|. \end{aligned}$$

The last display, together with Chebychev's inequality and Lemma C.1, implies that for any $p \in \mathcal{P}(\alpha, L, L')$

$$\begin{aligned} & \bar{\mathbb{P}}_p(|\tilde{K}_{n,m} - K_n R_{n,m}(p)| > (n\lambda)^\alpha \varepsilon) \\ & \leq \bar{\mathbb{P}}_p\left(|\tilde{K}_{n,m} - \bar{K}_{n,m}(p)| > \frac{(n\lambda)^\alpha \varepsilon}{2}\right) + \bar{\mathbb{P}}_p\left(|K_n - \bar{K}_n(p)| > \frac{(n\lambda)^\alpha \varepsilon}{2R_{n,m}(p)}\right) \\ & \leq \frac{4\overline{\text{var}}_p(\tilde{K}_{n,m})}{(n\lambda)^{2\alpha} \varepsilon^2} + \frac{4R_{n,m}(p)^2 \overline{\text{var}}_p(K_n)}{(n\lambda)^{2\alpha} \varepsilon^2} \\ & \lesssim \frac{1}{\lambda(n\lambda)^\alpha \varepsilon^2} + \frac{R_{n,m}(p)^2}{\lambda^{2\alpha} n^\alpha \varepsilon^2}. \end{aligned}$$

When $p \in \mathcal{P}(\alpha, L, L')$, we get by Lemma E.3 that $R_{n,m}(p) = \lambda^\alpha(1 + o(1))$ as $n \rightarrow \infty$ [recall that $K_n = C_{n,1}$]. Hence, there is a constant $c_1 > 0$ such that

$$\bar{\mathbb{P}}_p(|\tilde{K}_{n,m} - K_n R_{n,m}(p)| > (n\lambda)^\alpha \varepsilon) \leq \frac{c_1}{n^\alpha \varepsilon^2}. \quad (\text{C.2})$$

Since by assumption $\log(\lambda) \geq b \log(n)$, if the constant b is large enough we can make ε sufficiently large so that the previous display smaller than $1/16$. We note that in order to reduce the problem to the estimation of α it is necessary to take $G_1 \propto \sqrt{D}$ where D is the constant in Theorem 5, so that our only degree of freedom to make the rhs of (C.2) smaller than $1/16$ is to have b large enough. Since for any events A, B we have $\bar{\mathbb{P}}_p(A \cap B) = \bar{\mathbb{P}}_p(A) + \bar{\mathbb{P}}_p(B) - \bar{\mathbb{P}}_p(A \cup B) \geq \bar{\mathbb{P}}_p(A) + \bar{\mathbb{P}}_p(B) - 1 = \bar{\mathbb{P}}_p(A) - \bar{\mathbb{P}}_p(B^c)$, we deduce from equations (C.1) and (C.2) that for all \hat{U} and all $p \in \mathcal{P}(\alpha, L, L')$

$$\bar{\mathbb{P}}_p(\ell_\alpha(\hat{U}, U_{n,m}) > \varepsilon^2) \geq \bar{\mathbb{P}}_p\left(|\hat{R} - R_{n,m}(p)| > \frac{2(n\lambda)^\alpha \varepsilon}{K_n}\right) - \frac{1}{16}.$$

Also, we already know that $K_n \sim L\Gamma(1 - \alpha)n^\alpha$ almost-surely if $p \in \mathcal{P}(\alpha, L, L')$. So in fact if n is large enough there is a constant $c_2 > 0$ such that for all \hat{U} and all $p \in \mathcal{P}(\alpha, L, L')$

$$\bar{\mathbb{P}}_p(\ell_\alpha(\hat{U}, U_{n,m}) > \varepsilon^2) \geq \bar{\mathbb{P}}_p(|\hat{R} - R_{n,m}(p)| > 2c_2 \lambda^\alpha \varepsilon) - \frac{1}{8}. \quad (\text{C.3})$$

The fourth step consists on reducing to the problem of estimating the functional $p \mapsto \lambda^{\alpha(p)}$. Indeed, we have already shown that $R_{n,m}(p) = \lambda^\alpha(1 + o(1))$ which is a consequence of Lemma E.3. Indeed, by being more careful we see that in fact $R_{n,m}(p) = (1 + \lambda)^\alpha(1 + O(n^{-\alpha/2} \sqrt{\log(n)}))$ whenever $p \in \mathcal{P}(\alpha, L, L')$. The $\sqrt{\log(n)}$ is a bit problematic as it would make the reduction scheme work only for $\varepsilon \geq n^{-\alpha/2} \sqrt{\log(n)}$, that is only for $\log(\lambda) \gg \log(n)^{3/2}$. In order to weaken to $\log(\lambda) \gg \log(n)$, we leverage that the maximum risk over $\mathcal{P}(\alpha, L, L')$ has to be greater than the maximum risk over $\mathcal{P}(\alpha, L, L') \cap \mathcal{R}$, where we define

$$\mathcal{R} := \left\{ p : |R_{n,m}(p) - (1 + \lambda)^\alpha| \leq c_2 \lambda^{\alpha(p)} \varepsilon \right\}.$$

Later we will have to take additional care to handle the fact that we restrict the class of parameters to a better behaved subclass, but for now it is always true that from (C.3) and the definition of \mathcal{R} that

$$\inf_{\hat{U}} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p(\ell_\alpha(\hat{U}, U_{n,m}) > \varepsilon^2) \geq \inf_{\hat{R}} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L') \cap \mathcal{R}}} \bar{\mathbb{P}}_p(|\hat{R} - (1 + \lambda)^\alpha| > c_2 \lambda^\alpha \varepsilon) - \frac{1}{8}. \quad (\text{C.4})$$

Observe that, for every $c_3 > 0$,

$$\begin{aligned} |\log \hat{R} - \alpha \log(1 + \lambda)| &> c_3 \varepsilon \\ \implies \log \hat{R} - \alpha \log(1 + \lambda) &> c_3 \varepsilon \quad \text{or} \quad \log \hat{R} - \alpha \log(1 + \lambda) < -c_3 \varepsilon \\ \implies |\hat{R} - (1 + \lambda)^\alpha| &> \min\{(1 + \lambda)^\alpha (e^{c_3 \varepsilon} - 1), (1 + \lambda)^\alpha (1 - e^{-c_3 \varepsilon})\} \\ \implies |\hat{R} - (1 + \lambda)^\alpha| &> (1 + \lambda)^\alpha c_3 \varepsilon e^{-c_3 \varepsilon} \\ \implies |\hat{R} - (1 + \lambda)^\alpha| &> (1 + \lambda)^\alpha c_3 e^{-c_3 G_2 \varepsilon}, \end{aligned}$$

where the last line follows because by assumption $\varepsilon \leq G_2$. Hence, by taking G_2 sufficiently small and c_3 sufficiently large, we have

$$|\log \hat{R} - \alpha \log(1 + \lambda)| > c_3 \varepsilon \implies |\hat{R} - (1 + \lambda)^\alpha| > c_2 \lambda^\alpha \varepsilon.$$

Therefore,

$$\begin{aligned} \bar{\mathbb{P}}_p(|\hat{R} - (1 + \lambda)^\alpha| > c_2 \lambda^\alpha \varepsilon) &\geq \bar{\mathbb{P}}_p(|\log(\hat{R}) - \alpha \log(1 + \lambda)| > c_3 \varepsilon) \\ &\geq \bar{\mathbb{P}}_p\left\{\left|\frac{\log(\hat{R})}{\log(1 + \lambda)} - \alpha\right| > \frac{c_3 \varepsilon}{\log(1 + \lambda)}\right\} \\ &\geq \bar{\mathbb{P}}_p\left\{\left|\frac{\log(\hat{R})}{\log(1 + \lambda)} - \alpha\right| > \frac{c_3 G_1}{n^{\alpha/2} \log(n)}\right\}, \end{aligned}$$

where the last inequality follows because $\varepsilon \leq G_1 \log(1 + \lambda) n^{-\alpha/2} / \log(n)$ by construction. Then by (C.4), we have obtained that

$$\inf_{\hat{U}} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \bar{\mathbb{P}}_p(\ell_\alpha(\hat{U}, U_{n,m}) > \varepsilon^2) \geq \inf_{\hat{\alpha}} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L') \cap \mathcal{R}}} \bar{\mathbb{P}}_p\left\{|\hat{\alpha} - \alpha| > \frac{c_3 G_1}{n^{\alpha/2} \log(n)}\right\} - \frac{1}{8}.$$

We have finally established that estimating α is an easier task than predicting $U_{n,m}$. However, we cannot immediately conclude using the Theorem 5, as we have restricted the maximum over $\mathcal{P}(\alpha, L, L') \cap \mathcal{R}$ instead of $\mathcal{P}(\alpha, L, L')$. But, in view of the Theorem 4, it is obvious that the conclusion of Theorem 5 remains true over $\mathcal{P}(\alpha, L, L') \cap \mathcal{R}$ if we establish that $\nu_0(\mathcal{R})$ and $\nu_1(\mathcal{R})$ can be made arbitrarily close to 1, which is the case is $\log(\lambda) \geq b \log(n)$ for suitably large b by Lemma C.2.

C.2. Auxiliary results

LEMMA C.1. *For any $p \in \mathcal{P}(\alpha, L, L')$ it holds*

$$\bar{\mathbb{E}}_p[\tilde{K}_{n,m}] = L\Gamma(1 - \alpha)n^\alpha(1 + \lambda)^\alpha + O((n\lambda)^{\alpha/2}\sqrt{\log(n\lambda)}),$$

and,

$$\overline{\text{var}}_p(\tilde{K}_{n,m}) \leq \frac{\bar{\mathbb{E}}_p[K_{n+2m}]}{1+2\lambda} \lesssim \frac{(n\lambda)^\alpha}{\lambda}.$$

Proof. We observe that $\bar{\mathbb{E}}_p[\tilde{K}_{n,m}] = \bar{\mathbb{E}}_p[K_{n+m}]$, so that the first relation simply follows from Lemma E.3. For the second relation, we proceed using an Efron-Stein argument, as in the Lemma A.2. We write $\Delta_j := \sum_{i=n+1}^{n+m} \mathbf{1}_{X_i=j}$, and we observe that $Y_{n+m,j} \geq Y_{n,j}$ almost-surely. Then,

$$\begin{aligned} K_{n+m} &= \sum_{j \geq 1} \mathbf{1}_{Y_{n+m,j} \geq 1} \\ &= \sum_{j \geq 1} \mathbf{1}_{Y_{n,j} \geq 1} \mathbf{1}_{Y_{n+m,j} \geq 1} + \sum_{j \geq 1} \mathbf{1}_{Y_{n,j} = 0} \mathbf{1}_{Y_{n+m,j} \geq 1} \\ &= K_n + \sum_{j \geq 1} \mathbf{1}_{Y_{n,j} = 0} \mathbf{1}_{\Delta_j \geq 1} \\ &= K_n + \sum_{j \geq 1} \mathbf{1}_{Y_{n,j} = 0} (1 - \mathbf{1}_{\Delta_j = 0}). \end{aligned}$$

It follows,

$$\begin{aligned} \tilde{K}_{n,m} &= \mathbb{E}[K_{n+m} \mid \Pi_n] = K_n + \sum_{j \geq 1} \mathbf{1}_{Y_{n,j} = 0} (1 - (1 - p_j)^m) \\ &= K_n + \sum_{j \geq 1} (1 - \mathbf{1}_{Y_{n,j} \geq 1}) (1 - (1 - p_j)^m) \\ &= K_n + \sum_{j \geq 1} (1 - (1 - p_j)^m) - \sum_{j \geq 1} \mathbf{1}_{Y_{n,j} \geq 1} (1 - (1 - p_j)^m) \\ &= \sum_{j \geq 1} (1 - (1 - p_j)^m) + \sum_{j \geq 1} \mathbf{1}_{Y_{n,j} \geq 1} (1 - p_j)^m \end{aligned}$$

Then, with the same notations as in Lemma A.2, we can bound

$$\begin{aligned} \overline{\text{var}}_p(\tilde{K}_{n,m}) &= \overline{\text{var}}_p\left(\sum_{j \geq 1} \mathbf{1}_{Y_{n,j} \geq 1} (1 - p_j)^m\right) \\ &\leq \sum_{i=1}^n \bar{\mathbb{E}}_p \left[\left(\sum_{j \geq 1} (\mathbf{1}_{Y_{n,j} \geq 1} - \mathbf{1}_{Y_{n,j}^{(i)} \geq 1}) (1 - p_j)^m \right)^2 \right] \\ &= \sum_{i=1}^n \bar{\mathbb{E}}_p \left[\left(\sum_{j \geq 1} \mathbf{1}_{Y_{n,j} = 1} \mathbf{1}_{X_i=j} (1 - p_j)^m \right)^2 \right] \\ &= \sum_{i=1}^n \bar{\mathbb{E}}_p \left[\sum_{j \geq 1} \mathbf{1}_{Y_{n,j} = 1} \mathbf{1}_{X_i=j} (1 - p_j)^m \sum_{k \geq 1} \mathbf{1}_{Y_{n,k} = 1} \mathbf{1}_{X_i=k} (1 - p_k)^m \right] \\ &= \bar{\mathbb{E}}_p \left[\sum_{j \geq 1} \mathbf{1}_{Y_{n,j} = 1} \sum_{i=1}^n \mathbf{1}_{X_i=j} (1 - p_j)^{2m} \right] \end{aligned}$$

$$= \sum_{j \geq 1} \bar{\mathbb{E}}_p[\mathbf{1}_{Y_{n,j}=1}](1-p_j)^{2m},$$

where the last line follows because $Y_{n,j} = \sum_{i=1}^n \mathbf{1}_{X_i=j}$ by construction. Now $Y_{n,j} = 1$ if and only if all the X_1, \dots, X_n are not equal to j but one, which can happen only in n ways (either $X_1 = j$ and all the others are not equal to j , or $X_2 = j$ and the other are not equal to j , etc..). We deduce that

$$\begin{aligned} \overline{\text{var}}_p(\tilde{K}_{n,m}) &\leq \sum_{j \geq 1} np_j(1-p_j)^{n+2m-1} \\ &= \frac{n}{n+2m} \bar{\mathbb{E}}_p[M_{n+2m,1}] \\ &\leq \frac{\bar{\mathbb{E}}_p[K_{n+2m}]}{1+2\lambda}, \end{aligned}$$

which concludes the proof. \square

LEMMA C.2. *For every $\eta > 0$ there exists $C > 0$ and $n_0 > 0$ such that if $\varepsilon \geq Cn^{\alpha_j/2}$ and $n \geq n_0$, then*

$$\min\{\nu_0(\mathcal{R}), \nu_1(\mathcal{R})\} \geq 1 - \eta.$$

Proof. Let $j = 0$ or $j = 1$ arbitrary. Recall that under ν_j we have $\alpha(P) = \alpha_j$ almost-surely. Then, we observe that

$$|R_{n,m}(p) - (1+\lambda)^{\alpha_j}| \leq c_2 \lambda^{\alpha_j} \varepsilon \iff |\bar{K}_{n,m}(p) - (1+\lambda)^{\alpha_j} \bar{K}_n(p)| \leq c_2 \lambda^{\alpha_j} \bar{K}_n(p) \varepsilon. \quad (\text{C.5})$$

Recall that $\bar{K}_n(p) = n \int_0^1 \bar{F}_p(x)(1-x)^{n-1} dx$; see for instance the proof of Proposition A.2. Recall that conditional on $(J_i)_{i \geq 1}$ we have $\bar{F}_P(x) = \sum_{i \geq 1} \mathbf{1}_{J_i > xT}$ with $T = \sum_{i \geq 1} J_i$. Then, we can decompose $\bar{K}_n(P)$ as

$$\bar{K}_n(P) = n \int_0^1 \bar{\rho}_j(xT)(1-x)^{n-1} dx + \underbrace{n \int_0^1 \{\bar{F}_P(x) - \bar{\rho}_j(xT)\}(1-x)^{n-1} dx}_{\Delta_n(P)}. \quad (\text{C.6})$$

We now show that $\Delta_n(P) = O_p(n^{\alpha_j/2})$ where the O_p is understood with respect to ν_j . Indeed,

$$\begin{aligned} |\Delta_n(P)| &\leq n \int_0^1 \left| \sum_{i \geq 1} \mathbf{1}_{J_i > xT} - \bar{\rho}_j(xT) \right| (1-x)^{n-1} \\ &= \frac{n}{T} \int_0^T \left| \sum_{i \geq 1} \mathbf{1}_{J_i > x} - \bar{\rho}_j(x) \right| e^{-(n-1)x/T} dx \end{aligned}$$

Recall that T has an α_j -stable distribution. Then, for any $\eta > 0$ we can find $t_0, t_1 > 0$ such that with ν_j -probability at least $1 - \eta$ we have $t_0 \leq T \leq t_1$. Moreover, on the event that $t_0 \leq T \leq t_1$, we have

$$|\Delta_n(P)| \leq \frac{n}{t_0} \int_0^\infty \left| \sum_{i \geq 1} \mathbf{1}_{J_i > x} - \bar{\rho}_j(x) \right| e^{-(n-1)x/t_1} dx.$$

So it is enough to show that the above integral is a $O_p(n^{\alpha_j/2})$. But, under ν_j , the variable $\sum_{i \geq 1} \mathbf{1}_{J_i > x}$ has a Poisson distribution with mean $\bar{\rho}_j(x)$. Therefore,

$$\begin{aligned} \mathbb{E}_{\nu_j} \left[\int_0^\infty \left| \sum_{i \geq 1} \mathbf{1}_{J_i > x} - \bar{\rho}_j(x) \right| e^{-(n-1)x/t_1} dx \right] &= \int_0^\infty \mathbb{E}_{\nu_j} \left[\left| \sum_{i \geq 1} \mathbf{1}_{J_i > x} - \bar{\rho}_j(x) \right| \right] e^{-(n-1)x/t_1} dx \\ &\leq \int_0^\infty \bar{\rho}_j(x)^{1/2} e^{-(n-1)x/t_1} dx. \end{aligned}$$

But $\bar{\rho}_j(x) \lesssim x^{-\alpha_j} \mathbf{1}_{0 < x \leq 1}$, which follows from (B.3). This establish the claim that $\Delta_n(P) = O_p(n^{\alpha_j/2})$, because

$$\int_0^1 x^{-\alpha_j/2} e^{-(n-1)x/t_1} dx \leq \left\{ \frac{n-1}{t_1} \right\}^{\alpha_j/2-1} \Gamma(1 + \alpha_j/2).$$

On the other hand, writing $C_j := \frac{\alpha_0/\alpha_j}{\Gamma(1-\alpha_j)}$, we have again by (B.3) that

$$\begin{aligned} n \int_0^1 \bar{\rho}_j(xT) (1-x)^{n-1} dx &= C_j n \int_0^1 (xT)^{-\alpha_j} (1-(xT)^{\alpha_j}) \mathbf{1}_{0 < xT \leq 1} (1-x)^{n-1} dx \\ &= C_j n \int_0^1 (xT)^{-\alpha_j} \mathbf{1}_{xT \leq 1} (1-x)^{n-1} dx + C_j n \int_0^1 \mathbf{1}_{xT \leq 1} (1-x)^{n-1} dx \\ &= C_j n \int_0^1 (xT)^{-\alpha_j} \mathbf{1}_{xT \leq 1} (1-x)^{n-1} dx + O(1). \end{aligned}$$

We consider two cases, according to whether $T \leq 1$ or not. If $T \leq 1$, then

$$\begin{aligned} n \int_0^1 (xT)^{-\alpha_j} \mathbf{1}_{xT \leq 1} (1-x)^{n-1} dx &= \frac{n}{T^{\alpha_j}} \int_0^1 x^{-\alpha_j} (1-x)^{n-1} dx \\ &= \frac{\Gamma(1-\alpha_j)}{T^{\alpha_j}} \frac{n!}{\Gamma(n+1-\alpha)} \\ &= \frac{\Gamma(1-\alpha_j) n^{\alpha_j}}{T^{\alpha_j}} + O(T^{-\alpha_j}). \end{aligned}$$

If $T > 1$, then

$$\begin{aligned} n \int_0^1 (xT)^{-\alpha_j} \mathbf{1}_{xT \leq 1} (1-x)^{n-1} dx &= \frac{n}{T^{\alpha_j}} \int_0^{1/T} x^{-\alpha_j} (1-x)^{n-1} dx \\ &= \frac{n}{T^{\alpha_j}} \int_0^1 x^{-\alpha_j} (1-x)^{n-1} dx - \frac{n}{T^{\alpha_j}} \int_{1/T}^1 x^{-\alpha_j} (1-x)^{n-1} dx \\ &= \frac{\Gamma(1-\alpha_j) n^{\alpha_j}}{T^{\alpha_j}} + O(T^{-\alpha_j}) - \frac{n}{T^{\alpha_j}} \int_{1/T}^1 x^{-\alpha_j} (1-x)^{n-1} dx \\ &= \frac{\Gamma(1-\alpha_j) n^{\alpha_j}}{T^{\alpha_j}} + O(1) - \frac{n}{T^{\alpha_j}} \int_{1/T}^1 x^{-\alpha_j} (1-x)^{n-1} dx \end{aligned}$$

On the event that $t_0 \leq T \leq t_1$, it is clear that the integral in the rhs of the last display is a $o(1)$ as $n \rightarrow \infty$, so that we have shown that in both the cases we have

$$\begin{aligned} n \int_0^1 \bar{\rho}_j(xT)(1-x)^{n-1} dx &= \frac{C_j \Gamma(1-\alpha_j) n^{\alpha_j}}{T^{\alpha_j}} + O_p(1) \\ &= \frac{\alpha_0/\alpha_j}{T^{\alpha_j}} n^{\alpha_j} + O_p(1). \end{aligned}$$

Combining this result with (C.6), it follows that under ν_j ,

$$\bar{K}_n(P) = \frac{\alpha_0/\alpha_j}{T^{\alpha_j}} n^{\alpha_j} + O_p(n^{\alpha_j/2}).$$

Since $\bar{K}_{n,m}(P) = \bar{K}_{n+m}(P)$ and since $m = n\lambda$, we deduce that under ν_j

$$\begin{aligned} \bar{K}_{n,m}(P) - (1+\lambda)^{\alpha_j} \bar{K}_n(P) &= \frac{\alpha_0/\alpha_j}{T^{\alpha_j}} \left\{ (n+m)^{\alpha_j} - (1+\lambda)^{\alpha_j} n^{\alpha_j} \right\} + O_p(n^{\alpha_j/2} \lambda^{\alpha_j}) \\ &= O_p\left(\frac{\lambda^{\alpha_j} \bar{K}_n(P)}{n^{\alpha_j/2}}\right) \end{aligned}$$

It follows that if $\varepsilon \geq C n^{\alpha_j/2}$ for a large enough constant $C > 0$, then with ν_j probability at least $1 - \eta$ it will be true that $|\bar{K}_{n,m}(P) - (1+\lambda)^{\alpha_j} \bar{K}_n(P)| \leq c_2 \lambda^{\alpha_j} \bar{K}_n(P) \varepsilon$, and then the conclusion follows from (C.5). \square

Appendix D: Method of the two fuzzy hypotheses for partition models

This section contains the proof of Theorem 4 and of Proposition 1.

D.1. Proof of Theorem 4

The proof is a minor adaptation of [53, Theorems 2.14 and 2.15]. For any estimator $\hat{\delta}$,

$$\begin{aligned} \sup_{p \in \mathcal{P}'} \bar{\mathbb{P}}_p \left(|\hat{\delta}(\Pi_n) - \delta(\Pi, p)| \geq \varepsilon \right) &\geq \frac{1}{2} \sum_{j=0}^1 \int_{\mathcal{P}'} \bar{\mathbb{P}}_p \left(|\hat{\delta}(\Pi_n) - \delta(\Pi, p)| \geq \varepsilon \right) \frac{\nu_j(dp)}{\nu_j(\mathcal{P}')} \\ &\geq \frac{1}{2} \sum_{j=0}^1 \int \bar{\mathbb{P}}_p \left(|\hat{\delta}(\Pi_n) - \delta(\Pi, p)| \geq \varepsilon \right) \nu_j(dp) - \xi \\ &= \frac{1}{2} \sum_{j=0}^1 F_j \left(|\hat{\delta}(\Pi_n) - \delta(\Pi, P)| \geq \varepsilon \right) - \xi \end{aligned}$$

But, because for any events A, B we have $F_j(A \cap B) = F_j(A) + F_j(B) - F_j(A \cup B) \geq F_j(A) + F_j(B) - 1 = F_j(A) - F_j(B^c)$, we get

$$\begin{aligned} F_0\left(|\hat{\delta}(\Pi_n) - \delta(\Pi, P)| \geq \varepsilon\right) &\geq F_0(\hat{\delta}(\Pi_n) \geq c + \varepsilon, \delta(\Pi, P) \leq c) \\ &\geq F_0(\hat{\delta}(\Pi_n) \geq c + \varepsilon) - F_0(\delta(\Pi, P) > c) \\ &\geq F_0(\hat{\delta}(\Pi_n) \geq c + \varepsilon) - \kappa \\ &= \nu_0(\mathbb{P}_P^n)(\hat{\delta}(\Pi_n) \geq c + \varepsilon) - \kappa, \end{aligned}$$

and,

$$\begin{aligned} F_1\left(|\hat{\delta}(\Pi_n) - \delta(\Pi, P)| \geq \varepsilon\right) &\geq F_1(\hat{\delta}(\Pi_n) < c + \varepsilon, \delta(\Pi, P) \geq c + 2\varepsilon) \\ &\geq F_1(\hat{\delta}(\Pi_n) < c + \varepsilon) - F_1(\delta(\Pi, P) < c + 2\varepsilon) \\ &\geq \nu_1(\mathbb{P}_P^n)(\hat{\delta}(\Pi_n) < c + \varepsilon) - \kappa. \end{aligned}$$

Then,

$$\begin{aligned} \sup_{p \in \mathcal{P}'} \bar{\mathbb{P}}_p(|\hat{\delta}(\Pi_n) - \delta(\Pi, p)| \geq \varepsilon) &\geq \frac{1}{2} \left(\nu_0(\mathbb{P}_P^n)(\hat{\delta} \geq c + \varepsilon) + \nu_1(\mathbb{P}_P^n)(\hat{\delta} < c + \varepsilon) \right) - \kappa - \xi \\ &\geq \frac{1}{2} \left(1 - \|\nu_0(\mathbb{P}_P^n) - \nu_1(\mathbb{P}_P^n)\|_{\text{TV}} \right) - \kappa - \xi. \end{aligned}$$

D.2. Proof of Proposition 1

We note that $\sum_{j=1}^n M_{n,j} \leq n$ and $\sum_{j=1}^n j M_{n,j} = n$ and we will write \mathcal{M}_n the set of vectors satisfying these constraints. By [48, Chapter 2] the statistic \mathbf{M}_n is sufficient for ρ , so that the KL divergence between the laws of Π_n is the same as the KL divergence between the laws of \mathbf{M}_n . We will abusively write \mathbb{P}_ρ^n to denote either the law of Π_n or the law of \mathbf{M}_n under the model ρ . We write \mathbb{Q}_ρ^n the joint-distribution of (\mathbf{M}_n, U_n) under the model ρ . We obtain our first inequality using the chain-rule for relative entropies

$$\text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n) \leq \text{KL}(\mathbb{Q}_{\rho_0}^n; \mathbb{Q}_{\rho_1}^n).$$

Let define the so-called Laplace exponents $\psi_j(u) := \int_0^\infty (1 - e^{-ux}) \rho_j(dx)$, $j = 0, 1$. In view of Propositions D.1 and D.2, for all $\mathbf{m} \in \mathcal{M}_n$,

$$\mathbb{Q}_{\rho_0}^n(\mathbf{M}_n = (m_1, \dots, m_n), U_n \in du) = \frac{n!}{\prod_{i=1}^n (i!)^{m_i} m_i!} \frac{u^{n-1} e^{-\psi_0(u)}}{\Gamma(n)} \prod_{\ell=1}^n \mu_0(u, \ell)^{m_\ell} du,$$

and similarly under ρ_1 . Therefore,

$$\text{KL}(\mathbb{P}_{\rho_0}^n; \mathbb{P}_{\rho_1}^n) \leq \mathbb{E}_{\rho_0}^n \left[-\psi_0(U_n) + \psi_1(U_n) + \sum_{\ell=1}^n M_{n,\ell} \log \frac{\mu_0(U_n, \ell)}{\mu_1(U_n, \ell)} \right].$$

Hence, to finish the proof it is enough to show that

$$\mathbb{E}_{\rho_0}^n[\psi_j(U_n)] = \sum_{\ell=1}^n \mathbb{E}_{\rho_0}^n \left[M_{n,\ell} \frac{\mu_j(U_n, \ell)}{\mu_0(U_n, \ell)} \right], \quad j = 0, 1. \quad (\text{D.1})$$

We now establish (D.1). The starting point is the following equation, which can be obtained from [48, Lemma 4.3].

$$\mathbb{P}_{\rho}^n(\Pi_n = \{A_1, \dots, A_k\} | T) = \int_{\mathbb{R}_+^k} \frac{f_T(\rho; T - \sum_{j=1}^k y_j) \mathbf{1}\{\sum_{j=1}^k y_j < T\}}{T^n f_T(\rho; T)} \prod_{j=1}^k y_j^{|A_j|} \prod_{j=1}^k \rho(dy_j), \quad (\text{D.2})$$

where $f_T(\rho; \cdot)$ is the density (wrt Lebesgue measure) of the variable T under the model ρ . From here, the strategy runs as follows. We first obtain from (D.2) the law of $\mathbf{M}_n | U_n$ (Proposition D.1), then the laws of $T | U_n$ and U_n (Proposition D.2), from which we determine $\mathbb{E}_{\rho}^n[M_{n,\ell} | U_n]$ (Proposition D.3). Finally, we show in Proposition D.4 that $\mathbb{E}_{\rho_0}^n[\psi_j(U_n)]$ can be conveniently reexpressed in term of $\mathbb{E}_{\rho_0}^n[M_{n,\ell} | U_n]$, leading to the equation (D.1). An important quantity that will show up everywhere in this section is the function h_n such that

$$h_n(\rho; u) := \int_0^{\infty} x^n f_T(\rho; x) e^{-ux} dx, \quad u \geq 0.$$

PROPOSITION D.1. For all $\mathbf{m} \in \mathcal{M}_n$, $j = 0, 1$,

$$\mathbb{P}_{\rho_j}^n(\mathbf{M}_n = (m_1, \dots, m_k) | U_n) = \frac{n!}{\prod_{i=1}^n (i!)^{m_i} m_i!} \frac{e^{-\psi_j(U_n)}}{h_n(\rho_j; U_n)} \prod_{\ell=1}^n \mu(U_n, \ell)^{m_{\ell}}.$$

Proof. By equation (D.2) and by Proposition D.2,

$$\begin{aligned} & \mathbb{P}_{\rho_j}^n(\Pi_n = \{A_1, \dots, A_k\} | U_n) \\ &= \int_{\mathbb{R}_+^k} \int_{\mathbb{R}_+} \frac{f_T(\rho_j; t - \sum_{j=1}^k y_k) \mathbf{1}\{t - \sum_{j=1}^k y_j > 0\}}{h_n(\rho_j; U_n)} e^{-U_n t} dt \prod_{j=1}^k y_j^{|A_j|} \prod_{j=1}^k \rho_j(dy_j) \\ &= \int_{\mathbb{R}_+^k} \int_{\mathbb{R}_+} \frac{f_T(\rho_j; t) e^{-U_n t}}{h_n(\rho_j; U_n)} dt \prod_{j=1}^k y_j^{|A_j|} e^{-U_n y_j} \prod_{j=1}^k \rho_j(dy_j) \\ &= \frac{e^{-\psi_j(U_n)}}{h_n(\rho_j; U_n)} \prod_{j=1}^k \mu_j(U_n, |A_j|). \end{aligned}$$

The conclusion follows from the previous result by taking a partition $\{A_1, \dots, A_k\}$ with m_1 blocks of size 1, m_2 blocks of size 2, etc.. and multiplying by the number of such partitions. It can be seen that there are exactly $\frac{n!}{\prod_{i=1}^n (i!)^{m_i} m_i!}$ such partitions [see for instance 48, Section 2]. \square

PROPOSITION D.2. The following are true.

$$\mathbb{P}_{\rho_j}^n(T \in dt | U_n) = \frac{t^n f_T(\rho_j; t) e^{-U_n t} dt}{h_n(\rho_j; U_n)},$$

and,

$$\mathbb{P}_{\rho_j}^n(U_n \in du) = \frac{u^{n-1} h_n(\rho_j; u) du}{\Gamma(n)}.$$

Proof. We have the first trivial relation,

$$\begin{aligned} \mathbb{P}_{\rho_j}^n(0 < T < a, 0 < U_n < b) &= \mathbb{P}_{\rho_j}^n(0 < T < a, 0 < G_n < bT) \\ &= \int_0^a f_T(\rho_j; t) \int_0^{bt} \frac{u^{n-1} e^{-u}}{\Gamma(n)} du dt \\ &= \int_0^a \int_0^b t^n f_T(\rho_j; t) \frac{u^{n-1} e^{-ut}}{\Gamma(n)} du dt. \end{aligned}$$

Whence the result after conditioning. \square

PROPOSITION D.3. *It holds*

$$\mathbb{E}_{\rho_j}^n[M_{n,\ell} | U_n] = \frac{h_{n-\ell}(\rho_j; U_n)}{h_n(\rho_j; U_n)} \binom{n}{\ell} \mu_j(U_n, \ell).$$

Proof. We obtain the result in three steps. We compute $\mathbb{E}_{\rho_j}^n[M_{n,\ell} | (J_j)_{j \geq 1}]$, then $\mathbb{E}_{\rho_j}^n[M_{n,\ell} | T]$, and finally $\mathbb{E}_{\rho_j}^n[M_{n,\ell} | U_n]$. Obviously,

$$\begin{aligned} \mathbb{E}_{\rho_j}^n[M_{n,\ell} | (J_j)_{j \geq 1}] &= \sum_{j \geq 1} \mathbb{E}_{\rho_j}^n[\mathbf{1}_{Y_{n,j}=\ell} | (J_j)_{j \geq 1}] \\ &= \binom{n}{\ell} \sum_{i \geq 1} \left(\frac{J_i}{T}\right)^\ell \left(1 - \frac{J_i}{T}\right)^{n-\ell} \\ &= T^{-n} \binom{n}{\ell} \sum_{i \geq 1} J_i^\ell (T - J_i)^{n-\ell}. \end{aligned}$$

It follows using Slivnyak-Mecke's theorem Daley and Vere-Jones [16], that for any $t > 0$

$$\begin{aligned} \mathbb{E}_{\rho_j}^n[M_{n,\ell} \mathbf{1}\{0 < T < t\}] &= \binom{n}{\ell} \mathbb{E}_{\rho_j}^n \left[\sum_{i \geq 1} \frac{J_i^\ell (\sum_{k \geq 1} J_k - J_i)^{n-\ell} \mathbf{1}\{0 < \sum_{k \geq 1} J_k < t\}}{(\sum_{k \geq 1} J_k)^n} \right] \\ &= \binom{n}{\ell} \int_{\mathbb{R}_+} x^\ell \mathbb{E}_{\rho_j}^n \left[\frac{T^{n-\ell} \mathbf{1}\{0 < T + x < t\}}{(T + x)^n} \right] \rho_j(dx) \\ &= \binom{n}{\ell} \int_{\mathbb{R}_+} f_T(\rho_j; u) u^{n-\ell} \int_{\mathbb{R}_+} \frac{x^\ell \mathbf{1}\{0 < u + x < t\}}{(u + x)^n} \rho_j(dx) du \\ &= \binom{n}{\ell} \int_0^t \int_0^v \frac{f_T(\rho_j; v-x)}{v^n} x^\ell (v-x)^{n-\ell} \rho_j(dx) dv. \end{aligned}$$

So, by Lebesgue's differentiation theorem, we deduce that,

$$\begin{aligned}\mathbb{E}_{\rho_j}^n [M_{n,\ell} | T = t] &= \lim_{\varepsilon \rightarrow 0} \frac{\mathbb{E}_{\rho_j}^n [M_{n,\ell} \mathbf{1}\{t - \varepsilon < T < t + \varepsilon\}]}{\mathbb{P}_{\rho_j}^n (t - \varepsilon < T < t + \varepsilon)} \\ &= \lim_{\varepsilon \rightarrow 0} \frac{\binom{n}{\ell} \int_{t-\varepsilon}^{t+\varepsilon} \int_0^v \frac{f_T(\rho_j; v-x)}{v^n} x^\ell (v-x)^{n-\ell} \rho_j(dx) dv}{\int_{t-\varepsilon}^{t+\varepsilon} f_T(\rho_j; v) dv} \\ &= \binom{n}{\ell} \int_0^t \frac{f_T(\rho_j; t-x)}{t^n f_T(\rho_j; t)} x^\ell (t-x)^{n-\ell} \rho_j(dx).\end{aligned}$$

So it remains to compute $\mathbb{E}_{\rho_j}^n [M_{n,\ell} | U_n] = \mathbb{E}_{\rho_j}^n [\mathbb{E}_{\rho_j}^n [M_{n,\ell} | T] | U_n]$. The law of $T | U_n$ has been determined in Proposition D.2. Using this result, we have

$$\begin{aligned}\mathbb{E}_{\rho_j}^n [M_{n,\ell} | U_n] &= \frac{1}{h_n(\rho_j; U_n)} \binom{n}{\ell} \int_{\mathbb{R}_+} \int_0^t f_T(\rho_j; t-x) e^{-U_n t} x^\ell (t-x)^{n-\ell} \rho_j(dx) dt \\ &= \frac{1}{h_n(\rho_j; U_n)} \binom{n}{\ell} \int_{\mathbb{R}_+} v^{n-\ell} f_T(\rho_j; v) e^{-U_n v} dv \int_{\mathbb{R}_+} x^\ell e^{-U_n x} \rho_j(dx) \\ &= \frac{h_{n-\ell}(\rho_j; U_n)}{h_n(\rho_j; U_n)} \binom{n}{\ell} \mu_j(U_n, \ell).\end{aligned}$$

□

PROPOSITION D.4. For $j = 0, 1$,

$$\mathbb{E}_{\rho_0}^n [\psi_j(U_n)] = \sum_{\ell=1}^n \binom{n}{\ell} \mathbb{E}_{\rho_0}^n \left[\frac{h_{n-\ell}(\rho_0; U_n)}{h_n(\rho_0; U_n)} \mu_j(U_n; \ell) \right] = \sum_{\ell=1}^n \mathbb{E}_{\rho_0}^n \left[M_{n,\ell} \frac{\mu_j(U_n, \ell)}{\mu_0(U_n, \ell)} \right].$$

Proof. We start from the rhs, and we will show that it is equal to the lhs. Using the expression of the density of U_n in Proposition D.2, we have that

$$\begin{aligned}\mathbb{E}_{\rho_0}^n \left[\frac{h_{n-\ell}(\rho_0; U_n)}{h_n(\rho_0; U_n)} \mu_j(U_n, \ell) \right] &= \frac{1}{\Gamma(n)} \int_{\mathbb{R}_+} u^{n-1} h_{n-\ell}(\rho_0; u) \mu_j(u, \ell) du \\ &= \frac{1}{\Gamma(n)} \int_{\mathbb{R}_+^3} u^{n-1} x^{n-\ell} f_T(\rho_0; x) e^{-ux} y^\ell e^{-uy} \rho_j(dy) dx du.\end{aligned}$$

Consequently,

$$\begin{aligned}\sum_{\ell=1}^n \binom{n}{\ell} \mathbb{E}_{\rho_0}^n \left[\frac{h_{n-\ell}(\rho_0; U_n)}{h_n(\rho_0; U_n)} \mu_j(U_n, \ell) \right] \\ = \frac{1}{\Gamma(n)} \int_{\mathbb{R}_+^3} u^{n-1} e^{-u(x+y)} f_T(\rho_0; x) \left\{ (x+y)^n - x^n \right\} \rho_j(dy) dx du.\end{aligned}$$

We remark that,

$$\frac{1}{\Gamma(n)} \int_{\mathbb{R}_+} f_T(\rho_0; x) (x+y)^n \int_{\mathbb{R}_+} u^{n-1} e^{-u(x+y)} du dx = \int_{\mathbb{R}_+} f_T(\rho_0; x) dx = \mathbb{E}_{\rho_0}^n [1].$$

and,

$$\begin{aligned} \frac{1}{\Gamma(n)} \int_{\mathbb{R}_+} u^{n-1} \int_{\mathbb{R}_+} x^n e^{-u(x+y)} f_T(\rho_0; x) dx du &= \frac{1}{\Gamma(n)} \int_{\mathbb{R}_+} u^{n-1} e^{-uy} h_n(\rho_0; u) du \\ &= \mathbb{E}_{\rho_0}^n [e^{-U_n y}]. \end{aligned}$$

So by Fubini,

$$\begin{aligned} \sum_{\ell=1}^n \binom{n}{\ell} \mathbb{E}_{\rho_0}^n \left[\frac{h_{n-\ell}(\rho_0; U_n)}{h_n(\rho_0; U_n)} \mu_j(U_n, \ell) \right] &= \int_{\mathbb{R}_+} \mathbb{E}_{\rho_0}^n [1 - e^{-U_n y}] \rho_j(dy) \\ &= \mathbb{E}_{\rho_0}^n \left[\int_{\mathbb{R}_+} (1 - e^{-U_n y}) \rho_j(dy) \right] \\ &= \mathbb{E}_{\rho_0}^n [\psi_j(U_n)]. \end{aligned}$$

□

Appendix E: Results on regular variations

E.1. Folklore results

In this section, we state some classical results about the behaviour of K_n , $\{M_{n,\ell} : \ell = 1, \dots, n\}$ and $\{C_{n,k} : k = 1, \dots, n\}$ when \bar{F}_p has regularly varying tails. Those results are used multiple times in the proofs of our theorems. The following lemma is a restatement of Corollary 21 in [26].

LEMMA E.1. *If there is $\alpha \in (0, 1)$ and a non-negative slowly varying function at infinity ℓ such that $\bar{F}_p(x) \sim x^{-\alpha} \ell(1/x)$ when $x \rightarrow 0$, then almost-surely as $n \rightarrow \infty$*

$$K_n \sim \Gamma(1 - \alpha) \cdot n^\alpha \ell(n),$$

and, for all $\ell \in \mathbb{N}$ almost-surely as $n \rightarrow \infty$,

$$M_{n,\ell} \sim \frac{\alpha \Gamma(\ell - \alpha)}{\ell!} \cdot n^\alpha \ell(n).$$

With the same arguments [26], we can easily establish the following for the random variables $\{C_{n,k} : k = 1, \dots, n\}$. It is to be noted that equation (66) in [37] establishes that $\lim_{n \rightarrow \infty} \frac{C_{n,k}}{\mathbb{E}_p[C_{n,k}]} = 1$ almost-surely under no condition on p . Surprisingly [37] does not explicit the asymptotic behaviour of $\bar{\mathbb{E}}_p[C_{n,k}]$ under regular variations; we fill this gap in the following lemma.

LEMMA E.2. *For all p and all $k \in \mathbb{N}$, $\lim_{n \rightarrow \infty} \frac{C_{n,k}}{\mathbb{E}_p[C_{n,k}]} \rightarrow 1$ almost-surely, where*

$$\bar{\mathbb{E}}_p[C_{n,k}] = k \binom{n}{k} \int_0^1 \bar{F}_p(x) x^{k-1} (1-x)^{n-k} dx.$$

Furthermore, if there is $\alpha \in (0, 1)$ and a non-negative slowly varying function at infinity ℓ such that $\bar{F}_p(x) \sim x^{-\alpha} \ell(1/x)$ when $x \rightarrow 0$, then for all $k \in \mathbb{N}$

$$\bar{\mathbb{E}}_p[C_{n,k}] \sim n^{\alpha} \ell(n) \cdot \frac{\Gamma(k - \alpha)}{\Gamma(k)}$$

as $n \rightarrow \infty$.

Proof. The almost-sure convergence result is proved in [37]. Next, the starting point is

$$\begin{aligned} \bar{\mathbb{E}}_p[M_{n,\ell}] &= \sum_{j \geq 1} \bar{\mathbb{E}}_p[\mathbf{1}_{Y_{n,j}=\ell}] \\ &= \binom{n}{\ell} \sum_{j \geq 1} p_j^\ell (1 - p_j)^{n-\ell} \\ &= \binom{n}{\ell} \sum_{j \geq 1} \int_0^{p_j} \left\{ \ell x^{\ell-1} (1-x)^{n-\ell} - (n-\ell) x^\ell (1-x)^{n-\ell-1} \right\} dx. \end{aligned}$$

Since $\bar{F}_p(x) = \sum_{j \geq 1} \mathbf{1}_{p_j > x}$

$$\bar{\mathbb{E}}_p[M_{n,\ell}] = \binom{n}{\ell} \int_0^1 \bar{F}_p(x) x^{\ell-1} (1-x)^{n-\ell-1} (\ell - nx) dx.$$

The first expression for the expectation follows after remarking that

$$\sum_{\ell=k}^n \binom{n}{\ell} x^{\ell-1} (1-x)^{n-\ell-1} (\ell - nx) = k \binom{n}{k} x^{k-1} (1-x)^{n-k}.$$

Furthermore,

$$\int_0^1 \bar{F}_p(x) x^{k-1} (1-x)^{n-k} dx = \frac{1}{n^k} \int_0^\infty \bar{F}_p(un^{-1}) u^{k-1} (1-un^{-1})^{n-k} \mathbf{1}_{u \leq n} du$$

But if $\bar{F}_p(x) \sim x^{-\alpha} \ell(x)$, by Potter's bound [7, Theorem 1.5.6] for all $\delta > 0$ we can find $A, \varepsilon > 0$ such that $\bar{F}_p(un^{-1})/[n^{\alpha} \ell(n)] \leq Au^{-\alpha-\delta}$ for all $un^{-1} \leq \varepsilon$. If $un^{-1} > \varepsilon$, then $\bar{F}_p(un^{-1}) \leq C$ for some universal C , and $(1-un^{-1})^{n-k} \leq e^{-u(n-k)/n}$. Since $n^{\alpha} \ell(n) \rightarrow \infty$, this shows that $u \mapsto \bar{F}_p(un^{-1}) u^{k-1} (1-un^{-1})^{n-k} \mathbf{1}_{u \leq n} / [n^{\alpha} \ell(n)]$ is dominated by an integrable function on \mathbb{R}_+ . Therefore by the dominated convergence theorem,

$$\begin{aligned} &\lim_{n \rightarrow \infty} \frac{1}{n^{\alpha} \ell(n)} \int_0^\infty \bar{F}_p(un^{-1}) u^{k-1} (1-un^{-1})^{n-k} \mathbf{1}_{u \leq n} du \\ &= \int_0^\infty \lim_{n \rightarrow \infty} \frac{\bar{F}_p(un^{-1})}{n^{\alpha} \ell(n)} u^{k-1} (1-un^{-1})^{n-k} \mathbf{1}_{u \leq n} du \\ &= \int_0^\infty u^{-\alpha+k-1} e^{-u} du \\ &= \Gamma(k - \alpha). \end{aligned}$$

The conclusion follows by Stirling's formula. \square

E.2. Uniform results

LEMMA E.3. For all $0 < \underline{\alpha} < \bar{\alpha} < 1$ and all $\bar{L} > 0$, and for all $k \geq 1$,

$$\limsup_{n \rightarrow \infty} \sup_{\substack{(\alpha, L, L') \in \mathbb{A} \\ p \in \mathcal{P}(\alpha, L, L')}} \frac{|\mathbb{E}_p[C_{n,k}] - Ln^\alpha \frac{\Gamma(k-\alpha)}{\Gamma(k)}|}{L'n^{\alpha/2} \sqrt{\log(n)} \Gamma(k-\alpha/2)} \leq \Gamma(k).$$

Proof. If $p \in \mathcal{P}(\alpha, L, L')$, then we can rewrite $\bar{F}_p(x) = Lx^{-\alpha} + \delta(x)$ where δ is a function satisfying $|\delta(x)| \leq L' \sqrt{x^{-\alpha} \log(e/x)}$. Then, by the Lemma E.2

$$\begin{aligned} \bar{\mathbb{E}}_p[C_{n,k}] &= Lk \binom{n}{k} \int_0^1 x^{k-\alpha-1} (1-x)^{n-k} dx + k \binom{n}{k} \int_0^1 \delta(x) x^{k-1} (1-x)^{n-k} dx \\ &= L \cdot \frac{n!}{\Gamma(n+1-\alpha)} \cdot \frac{\Gamma(k-\alpha)}{\Gamma(k)} + k \binom{n}{k} \int_0^1 \delta(x) x^{k-1} (1-x)^{n-k} dx. \end{aligned}$$

On one hand, by Stirling's formula, there exists n_0 such that for all $n \geq n_0$

$$\left| \frac{Ln!}{\Gamma(n+1-\alpha)} - Ln^\alpha \right| \leq 2Ln^{-1+\alpha}.$$

On the other hand, since $|\delta(x)| \leq L' \sqrt{x^{-\alpha} \log(e/x)}$,

$$\left| \int_0^1 \delta(x) x^{k-1} (1-x)^{n-k} dx \right| \leq L' \int_0^1 \sqrt{\log(e/x)} x^{-\alpha/2+k-1} (1-x)^{n-k} dx.$$

We remark that for all $B > 0$, the function $B \log(x) + \frac{1}{2} \log \log(e/x)$ has a unique maximum at $x = e^{1-1/(2B)}$. From this observation, we deduce that $\sqrt{\log(e/x)} \leq \frac{e^B}{e^{1/2} \sqrt{2B}} x^{-B}$ for all $B > 0$. Therefore taking $B > 0$ small enough, we find that

$$\begin{aligned} \left| \int_0^1 \delta(x) x^{k-1} (1-x)^{n-k} dx \right| &\leq \frac{L' e^B}{e^{1/2} \sqrt{2B}} \int_0^1 x^{-B-\alpha/2+k-1} (1-x)^{n-k} dx \\ &= \frac{L' e^B}{e^{1/2} \sqrt{2B}} \frac{\Gamma(n-k+1) \Gamma(k-\alpha/2-B)}{\Gamma(n+1-\alpha/2-B)}. \end{aligned}$$

By an asymptotic expansion of $\log \Gamma(z-B)$ as $B \rightarrow 0$ and using that $\frac{\Gamma'(z)}{\Gamma(z)} \sim \log(z)$ as $z \rightarrow \infty$, we find that

$$\begin{aligned} \limsup_{n \rightarrow \infty} \sup_{(\alpha, L, L') \in \mathbb{A}} \frac{|k \binom{n}{k} \int_0^1 \delta(x) x^{k-1} (1-x)^{n-k} dx|}{L'n^{\alpha/2} \sqrt{\log(n)} \Gamma(k-\alpha/2) / \Gamma(k)} &\leq \limsup_{n \rightarrow \infty} \frac{e^B \cdot e^{B \log(n)}}{e^{1/2} \sqrt{2B} \cdot \log(n)} \\ &\leq 1, \end{aligned}$$

by choosing $2B \cdot \log(n) = 1$. □

Appendix F: Additional results regarding simulations

F.1. Proof of Lemma 1

Recall that $p = C(\alpha, \beta, J) [\sum_{j=1}^J j^{-1/\alpha} \delta_j + \sum_{j>J} j^{-1/\beta} \delta_j]$, with $\beta < \alpha$. Now suppose $x < C(\alpha, \beta, J) J^{-1/\beta}$. We have that $C(\alpha, \beta, J) j^{-1/\alpha} \geq C(\alpha, \beta, J) j^{-1/\beta} \geq C(\alpha, \beta, J) J^{-1/\beta} > x$ for all $j \in \{1, \dots, J\}$. Therefore when $x < C(\alpha, \beta, J) J^{-1/\beta}$

$$\bar{F}_p(x) = \left\lfloor \frac{C(\alpha, \beta, J)^\beta}{x^\beta} \right\rfloor$$

from which we deduce that $\lim_{x \rightarrow 0} x^\beta \bar{F}_p(x) = C(\alpha, \beta, J)^\beta$ and using that $C(\alpha, \beta, J)^\beta - x^\beta \leq x^\beta [C(\alpha, \beta, J)^\beta / x^\beta] \leq C(\alpha, \beta, J)^\beta$

$$\frac{|x^\beta \bar{F}_p(x) - C(\alpha, \beta, J)^\beta|}{\sqrt{x^\beta \log(e/x)}} \leq \frac{x^\beta}{\sqrt{x^{\beta/2} \log(e/x)}} \leq \frac{C(\alpha, \beta, J)^\beta}{\sqrt{J - J \log C(\alpha, \beta, J) + (J/\beta) \log J}}.$$

We now consider the case where $x \geq C(\alpha, \beta, J) J^{-1/\beta}$. In this case we have that

$$\bar{F}_p(x) = \left\lfloor \frac{C(\alpha, \beta, J)^\alpha}{x^\alpha} \right\rfloor.$$

With the same reasoning as above,

$$|x^\beta \bar{F}_p(x) - C(\alpha, \beta, J)^\beta| \leq x^\beta + C(\alpha, \beta, J)^\beta \left| \frac{[C(\alpha, \beta, J)/x]^\alpha}{[C(\alpha, \beta, J)/x]^\beta} - 1 \right|.$$

Observe that $C(\alpha, \beta, J) \geq \sum_{j \geq 1} j^{-1/\beta} = \zeta(1/\beta)$ where ζ is Riemann's Zeta function, because $\alpha > \beta$. But a classical analysis of the Zeta function shows that $\zeta(1/\beta) > 1$ for all $0 < \beta < 1$. It follows that for all $x \in [0, 1]$ we must have $x \leq C(\alpha, \beta, J)$ and we only need to analyze the last display for $C(\alpha, \beta, J) J^{-1/\beta} \leq x \leq C(\alpha, \beta, J)$. In this case,

$$\frac{|x^\beta \bar{F}_p(x) - C(\alpha, \beta, J)^\beta|}{\sqrt{x^\beta \log(e/x)}} \leq \sqrt{\frac{x^\beta}{\log(e/x)}} + C(\alpha, \beta, J)^\alpha \frac{x^{\beta-\alpha} (\alpha - \beta) \log \frac{C(\alpha, \beta, J)}{x}}{\sqrt{x^\beta \log(e/x)}}$$

which completes the proof.

F.2. Elementary diagnoses for checking the power-law assumption

As mentioned in the introduction of the paper, it is impossible to guarantee that $p \in \mathcal{P}(\alpha, L, L')$ on the basis of the data alone. This is because one cannot extrapolate the behaviour of $\bar{F}_p(x)$ near zero from a finite amount of data. If that was possible, estimating the unseen would be feasible without extra modeling assumption. Consequently, without exterior information on \bar{F}_p , coming for instance from biological or physical insights, the result of the estimation cannot be, in general, trusted. That is said, power-law data have readily identifiable signatures that allows to assess if the power-law assumption is plausible or not. In other word, we cannot guarantee that the assumption holds, but it is often easy to diagnose when it does not hold.

One possible diagnostic consists in inspecting the behaviour of the statistic $j \mapsto C_{n,j}$. If the data is power-law, for small j the variables $C_{n,j}$ should not deviate too much from their expectation (in virtue of the Lemma E.3). Then, we expect that $C_{n,j} \propto \frac{\Gamma(j-\alpha)}{\Gamma(j)}$. To illustrate this, we have plotted in Figure 7 $\log(C_{n,j})$ as a function of $\log(j)$ for a partition generated from a Poisson-Kingman process with Lévy measure ρ_α , with $\alpha = 0.7$ and $n = 1000000$. We have added in blue the best fit (with respect to mean square error) obtained for $\log(C_{n,j}) = \log \frac{\Gamma(j-A)}{\Gamma(j)} + B$. This confirms that in the ideal scenario, the empirical cumulative frequencies are in good agreement with their expected values. In Figure 8, we did the same thing for the four real datasets used in Section 4.3. In all cases, the empirical cumulative frequencies seem compatible with the power-law assumption.

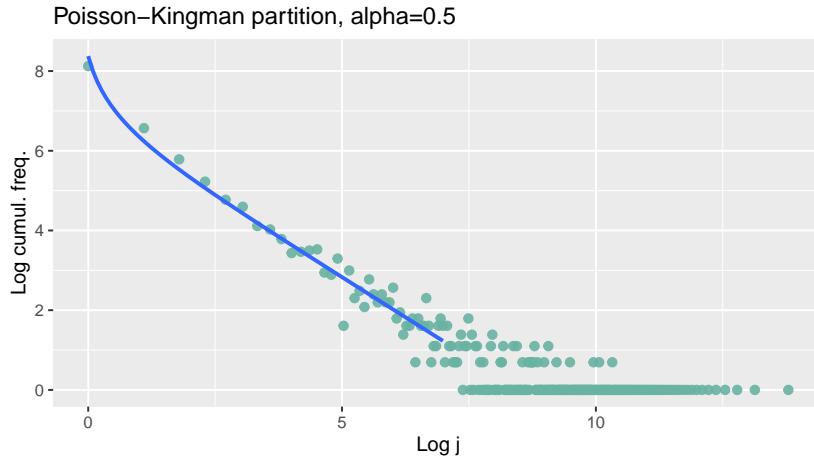


Figure 7. In green, $C_{n,j}$ as a function of j on a log-log scale. The cumulative frequencies $C_{n,j}$ come from a Poisson-Kingman partition with Lévy measure ρ_α , with $\alpha = 0.7$ and $n = 1000000$. In blue is the best fit for $\log(C_{n,j}) = \log \frac{\Gamma(j-A)}{\Gamma(j)} + B$, which is the expected theoretical behaviour of the cumulative frequencies for power-law data.

References

- [1] ANEVSKI, DRAGI AND GILL, RICHARD D AND ZOHREN, STEFAN (2017). Estimating a probability mass function with unknown labels. *The Annals of Statistics* **45**, 2708–2735.
- [2] AYED, F., BATTISTON, M., CAMERLENGHI, F. AND FAVARO, S. (2018). On consistent and rate optimal estimation of the missing mass. *Annales de l’Institut Henri Poincaré - Probabilités et Statistiques* **57**, 1476–1494.
- [3] BALABDAOUI, FADOUA AND KULAGINA, YULIA (2020). Completely monotone distributions: Mixing, approximation and estimation of number of species. *Computational Statistics & Data Analysis* **150**, 107014.
- [4] BALOCCHI, C., FAVARO, S. AND NAULET, Z. (2021) Bayesian nonparametric inference for “species-sampling problems”. *Preprint available upon request.*

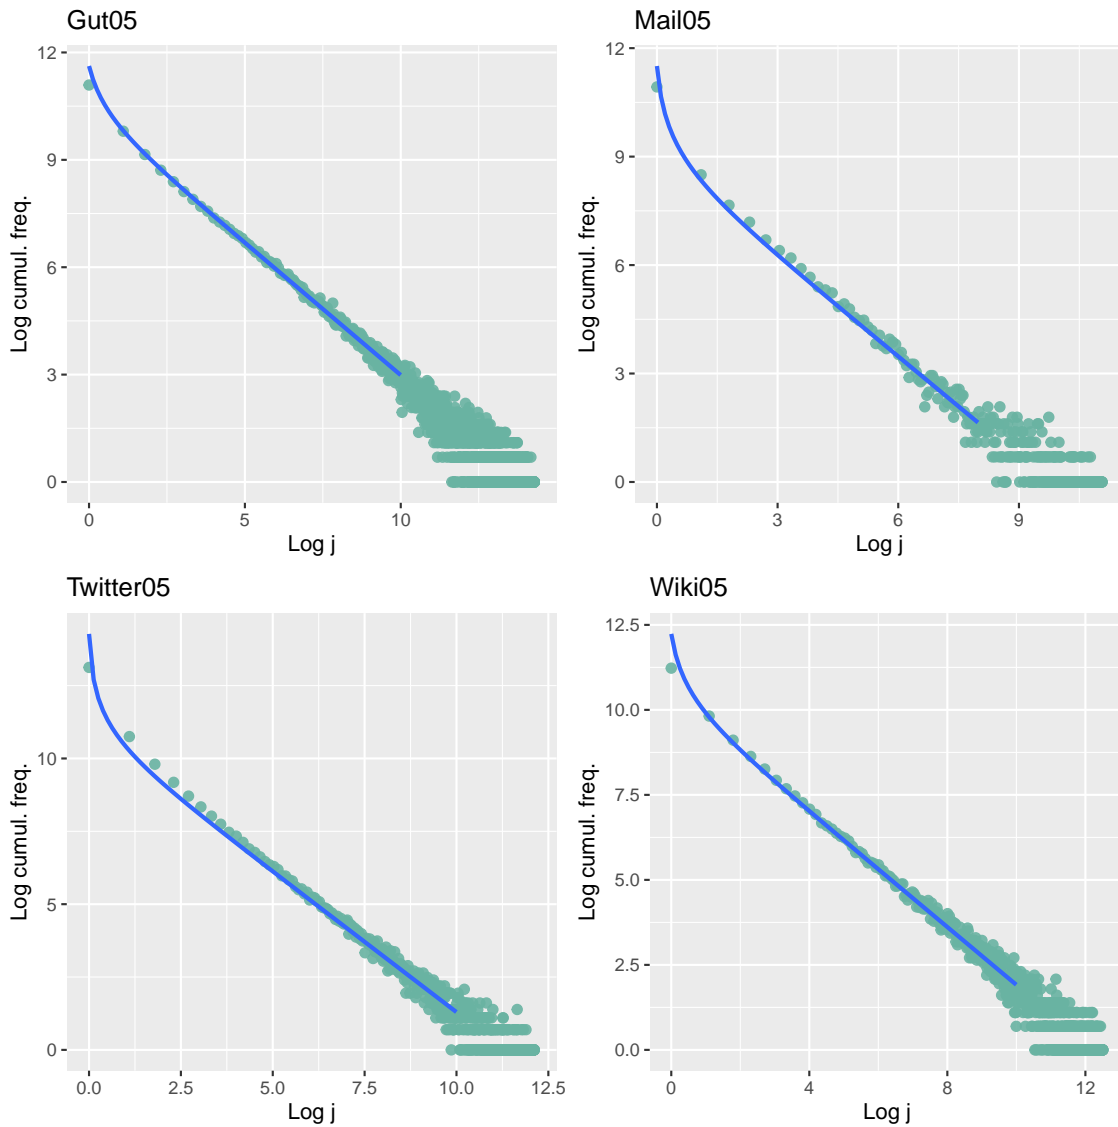


Figure 8. In green, $C_{n,j}$ as a function of j on a log-log scale for the four datasets used in the numerical illustration section. In blue are the corresponding best fit for $\log(C_{n,j}) = \log \frac{\Gamma(j-A)}{\Gamma(j)} + B$, which is the expected theoretical behaviour of the cumulative frequencies for power-law data.

- [5] BARABÁSI, A.L. (2005) The origin of bursts and heavy tails in human dynamics. *Nature* **435**, 227.
- [6] BEN-HAMOU, A., BOUCHERON, S. AND OHANNESSIAN, M.I. (2017). Concentration inequalities in the infinite urn scheme for occupancy counts and the missing mass, with applications. *Bernoulli* **23**, 249–287.

- [7] BINGHAM, N.H., GOLDIE, C.M. AND TEUGELS, J.L. (1989). *Regular variation*. Cambridge University Press.
- [8] BOUCHERON, S., LUGOSI, G. AND MASSART, P. (2013). *Concentration inequalities: a nonasymptotic theory of independence*. Oxford University Press.
- [9] BUNGE, J. AND FITZPATRICK, M. (1993) Estimating the number of species: a review. *Journal of the American Statistical Association* **88**, 364–373.
- [10] CAMERLENGHI, F., FAVARO, S., NAULET, Z. AND PANERO, F. (2020). Optimal disclosure risk assessment. *The Annals of Statistics* **49**, 723–744.
- [11] CANCHO, R.F. AND SOLÉ, R.V. (2003). Least effort and the origins of scaling in human language. *Proceeding of the National Academy of Sciences of USA* **100**, 788–791.
- [12] CHAO, A. AND LEE, S. (1992) Estimating the number of classes via sample coverage. *Journal of the American Statistical Association* **87**, 210–217.
- [13] CHEW-SENG CHEE AND YONG WANG (2016). Nonparametric estimation of species richness using discrete k -monotone distributions. *Computational Statistics and Data Analysis* **93**, 107–118.
- [14] CLAUSET, A., SHALIZI, C.R. AND NEWMAN, M.E.J. (2009). Power-law distributions in empirical data. *SIAM Review* **51**, 661–703.
- [15] DALEY, T. AND SMITH, A.D. (2013). Predicting the molecular complexity of sequencing libraries. *Nature Methods* **10**, 325–327.
- [16] DALEY, D.J. AND VERE-JONES, D. (2007). *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media.
- [17] DE HAAN, L. AND FERREIRA A. (2006). *Extreme value theory: an introduction* (Vol. 21). New York: Springer.
- [18] DREES, H. (1998). Optimal rates of convergence for estimates of the extreme value index. *The Annals of Statistics* **26**, 434–448
- [19] EFRON, B. AND THISTED, R. (1976). Estimating the number of unseen species: How many words did Shakespeare know? *Biometrika* **63**, 435–447.
- [20] FELLER, W. (1971). *An introduction to probability theory and its applications, Volume II*. Wiley.
- [21] FERGUSON T.S. AND MICHAEL J.K. (1972). A representation of independent increment processes without Gaussian components. *The Annals of Mathematical Statistics* **43**, 1634–1643.
- [22] FISHER, R.A., CORBET, A.S. AND WILLIAMS, C.B. (1943). The relation between the number of species and the number of individuals in a random sample of an animal population. *Journal of Animal Ecology* **12**, 42–58.
- [23] FORMENTIN, M., LOVISON, A., MARITAN, A., AND ZANZOTTO, G. (2014). Hidden scaling patterns and universality in written communication *Physical Review E* **90**, 012817
- [24] GAO, Z., TSENG, C.H., PEI, Z. AN BLASER, M.J. (2007). Molecular analysis of human forearm superficial skin bacterial biota. *Proceedings of the National Academy of Sciences of USA* **104**, 2927–2932.
- [25] GIGUELAY, J. AND HUET, S. Testing k -monotonicity of a discrete distribution. Application to the estimation of the number of classes in a population. *Computational Statistics and Data Analysis* **127**, 96–115.
- [26] GNEDIN, A., HANSEN, B. AND PITMAN, J. (2007). Notes on the occupancy problems with infinitely many boxes: general asymptotics and power law. *Probability Surveys* **4**, 146–171.
- [27] GOOD, I.J. AND TOULMIN, G.H. (1956). The number of new species, and the increase in population coverage, when a sample is increased. *Biometrika* **43**, 45–63.
- [28] HALL, P. AND WELSH, H. (1984). Best attainable rates of convergence for estimates of parameters of regular variation. *The Annals of Statistics* **12**, 1079–1084
- [29] HALL, P. AND WELSH, H. (1985). Adaptive estimates of parameters of regular variation. *The Annals of Statistics* **13**, 331–341

- [30] HAO, Y. AND LI, P. (2020). Optimal prediction of the number of unseen species with multiplicity. In *Advances in Neural Information Processing Systems*
- [31] HARALD, B.R. (2001). *Word Frequency Distributions*. Springer
- [32] HILL, B.M. (1975). A simple general approach to inference about the tail of a distribution. *The Annals of Statistics* **3**, 1163–1174
- [33] HUBERMAN, B.A. AND ADAMIC, L.A. (1999) Internet: growth dynamics of the World-Wide Web. *Nature* **401**, 131
- [34] IONITA-LAZA, I., LANGE, C. AND LAIRD, N.M. (2009). Estimating the number of unseen variants in the human genome. *Proceedings of the National Academy of Sciences of USA* **106**, 5008–5013.
- [35] JANA, SOHAM AND POLYANSKIY, YURI AND WU, YIHONG (2020). Extrapolating the profile of a finite population. *Conference on Learning Theory*, 2011–2033.
- [36] JIAO, J., VENKAT, K., HAN, Y. AND WEISSMAN, T. (2015). Minimax estimation of functionals of discrete distributions. *IEEE Transaction in Information Theory* **61**, 2835–2885.
- [37] KARLIN, SAMUEL (1967). Central Limit Theorems for Certain Infinite Urn Schemes. *Journal of Mathematics and Mechanics* **17**, 373–401
- [38] KINGMAN, J.F.C. (1975). Random discrete distributions. *Journal of the Royal Statistical Society Series B* **37**, 1–15.
- [39] KINGMAN, J.F.C. (1993). *Poisson processes*. Wiley Online Library.
- [40] KROES, I., LEPP, P.W. AND RELMAN, D.A. (1999). Bacterial diversity within the human sub-gingival crevice. *Proceeding of the National Academy of Sciences of USA* **96**, 14547–14552.
- [41] LIJOI, A. AND PRÜNSTER, I. (2010). Models beyond the Dirichlet process. In *Bayesian Non-parametrics*, Hjort, N.L., Holmes, C.C. Müller, P. and Walker, S.G. Eds. Cambridge University Press.
- [42] MONECHI, B., RUIZ-SERRANO, A., TRIA, F., AND LORETO, V. (2017). Waves of novelties in the expansion into the adjacent possible. *PloS ONE* **12**
- [43] MOSSEL, E. AND OHANNESSIAN, M.I. (2019) On the impossibility of learning the missing mass. *Entropy* **21**, 28.
- [44] MUCHNIK, L., PEI, S., PARRA, L.C., REIS, S.D.S, ANDRADE, J.S., HAVLIN, S. AND MAKSE, H.A. (2013). Origins of power-law degree distribution in the heterogeneity of human activity in social networks. *Nature Scientific Reports* **3**, 1783
- [45] OHANNESSIAN, M.I. AND DAHLEH, M.A. (2012). Rare probability estimation under regularly varying heavy tails. *Journal of Machine Learning Research* **23**, 1–24.
- [46] ORLITSKY, A., SURESH, A.T. AND WU, Y. (2017). Optimal prediction of the number of unseen species. *Proceeding of the National Academy of Sciences of USA* **113**, 13283–13288.
- [47] PITMAN, J. (2003). Poisson-Kingman partitions. In *Science and Statistics: A Festschrift for Terry Speed*, Goldstein, D.R. Eds. Institute of Mathematical Statistics.
- [48] PITMAN, J. (2006). *Combinatorial Stochastic Processes*. Ecole d’Eté de Probabilités de Saint-Flour XXXII. Lecture notes in mathematics, Springer - New York.
- [49] POLYANSKIY, Y. AND WE, Y. (2020). Dualizing Le Cam’s method for functional estimation, with applications to estimating the unseens. *Preprint: arXiv:1902.05616*
- [50] RYBSKI, D., BULDYREV, S.V., HAVLIN, S., LILJEROS, F. AND MAKSE, H A. (2016). Scaling laws of human interaction activity. *Proceeding of the National Academy of Sciences of USA* **106**, 12640.
- [51] SAMORODNITSKY, G. AND TAQQU M.S. (1994). *Stable Non-Gaussian Random Processes*. Chapman & Hall/CRC.
- [52] TRIA, F., LORETO, V., SERVEDIO, V.D.P AND STROGATZ, S.H. (2014). The dynamics of correlated novelties. *Nature Scientific Reports* **4**, 5890.

- [53] TSYBAKOV, A. (2009) *Introduction to nonparametric estimation*. Springer.
- [54] VALIANT, P. AND VALIANT, G. (2013). Estimating the unseen: improved estimators for entropy and other properties. In *Advances in Neural Information Processing Systems* **27**, 2157–2165
- [55] VAN DER VAART, A.W. (2000) *Asymptotic statistics*. Cambridge University Press.
- [56] WU, Y. AND YANG, P. (2016). Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transaction on Information Theory* **62**, 3702–3720.
- [57] WU, Y. AND YANG, P. (2019). Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *Annals of Statistics* **47**, 857–883.
- [58] ZIPF, G.K. (1949). *Human behaviour and the principle of least effort: an introduction to human ecology*. Addison-Wesley.