Estimation of discrete distributions in relative entropy, and the deviations of the missing mass

Jaouad Mourtada

May 2, 2025

Abstract

We study the problem of estimating a distribution over a finite alphabet from an i.i.d. sample, with accuracy measured in relative entropy (Kullback-Leibler divergence). While optimal expected risk bounds are known, high-probability guarantees remain less wellunderstood. First, we analyze the classical Laplace (add-1) estimator, obtaining matching upper and lower bounds on its performance and showing its optimality among confidenceindependent estimators. We then characterize the minimax-optimal high-probability risk, which is attained via a simple confidence-dependent smoothing technique. Interestingly, the optimal non-asymptotic risk exhibits an additional logarithmic factor over the ideal asymptotic risk. Next, motivated by scenarios where the alphabet exceeds the sample size, we investigate methods that adapt to the sparsity of the distribution at hand. We introduce an estimator using data-dependent smoothing, for which we establish a high-probability risk bound depending on two effective sparsity parameters. As part of the analysis, we also derive a sharp high-probability upper bound on the missing mass.

Contents

1	Introduction	2
	1.1 Problem setting	2
	1.2 Existing guarantees	4
	1.3 Paper outline	5
	1.4 Related work	6
	1.5 Notation	7
2	Optimal guarantees for the Laplace estimator	7
	2.1 Upper bound for the Laplace estimator	7
	2.2 Lower bound for confidence-independent estimators	8
3	Minimax-optimal guarantees for confidence-dependent estimators	9
	3.1 Upper bound via confidence-dependent smoothing	9
	3.2 Lower bound for confidence-dependent estimators	10
4	Adaptation to the effective support size	11
	4.1 Minimax lower bounds for sparse distributions	12
	4.2 Upper bounds for adaptive estimators	13
5	High-probability bound on the missing mass	19

6	Proof of high-probability upper bounds	24				
	6.1 Risk decomposition	25				
	6.2 Upper bound in Hellinger distance	26				
	6.3 Control of the contribution of underestimated frequencies	28				
	6.4 Proof of Theorem 1	33				
	6.5 Proof of Theorem 3	33				
	6.6 Proof of Theorem 5	34				
7 Proofs of lower bounds		36				
	7.1 Proof of Theorem 2	36				
	7.2 Proof of Lemma 1 and Theorem 4	38				
	7.3 Proof of Proposition 1	39				
	7.4 Proof of Corollary 1	42				
8	8 Proof of Theorem 6 42					
9	Proof of Proposition 2					
10	10 Technical lemmata 48					

1 Introduction

1.1 Problem setting

Estimating a discrete probability distribution from a finite sample is a fundamental problem in statistics, machine learning, and information theory. In this work, we consider the following variant of this problem. Let P be an unknown probability distribution on the finite set $\{1, \ldots, d\}$ (identified with the vector (p_1, \ldots, p_d) , where p_j denotes the probability of the class j); given access to an i.i.d. sample X_1, \ldots, X_n from P, find a distribution $\hat{P}_n = (\hat{p}_1, \ldots, \hat{p}_d)$ such that the Kullback-Leibler divergence or relative entropy

$$\mathrm{KL}(P, \widehat{P}_n) = \sum_{j=1}^d p_j \log\left(\frac{p_j}{\widehat{p}_j}\right) \tag{1}$$

is small, with high probability over the random draw of the i.i.d. sample X_1, \ldots, X_n from P.

The relative entropy (1) is a natural loss function for estimating distributions, which is commonly used in several fields: in statistics [vdV98, vdG99], due to its connection with maximum likelihood estimation; in machine learning, as the excess risk for prediction under logarithmic (cross-entropy) loss [Vap00, Bac24]; in information theory, owing to its interpretation in terms of excess code-length in data compression [CT06, Gas18, PW23]; and in natural language processing, through its link with the "perplexity" metric for evaluating language models [JM25].

An important feature of the Kullback-Leibler divergence, compared to other common divergences between probability distributions such as the total variation and Hellinger distances, is that it penalizes significant underestimation of true frequencies. To consider an extreme case, if the estimator \hat{P} assigns a probability $\hat{p}_j = 0$ to a class j with probability $p_j \neq 0$, then the Kullback-Leibler divergence $\text{KL}(P, \hat{P}_n)$ is infinite. This aligns with the needs of various applications: for instance, in a forecasting context where classes correspond to different outcomes, assigning a probability of 0 to outcomes that are in fact possible would constitute a severe underestimation of the underlying uncertainty. In addition, in the context of language models, one may be interested in generating new sentences that are not present in the training corpus; this requires assigning positive probabilities to sequence of words that have not been observed. **Empirical distribution.** Perhaps the most natural estimator is the empirical distribution $\overline{P}_n = (N_j/n)_{1 \leq j \leq d}$, where for $j = 1, \ldots, d$, we let

$$N_j = \sum_{i=1}^n \mathbf{1}(X_i = j) \tag{2}$$

denote the number of occurrences of the class j in the sample X_1, \ldots, X_n . This estimator happens to coincide with the maximum likelihood estimator (MLE) over the class \mathcal{P}_d of all probability distributions on $\{1, \ldots, d\}$. As such, it enjoys rather strong optimality properties in the "low-dimensional" asymptotic regime, where the number d of classes and distribution $P \in \mathcal{P}_d$ are fixed, while the sample n goes to infinity [LCY00].

In particular, if all classes j = 1, ..., d have nonzero probability, then \overline{P}_n converges to Pas $n \to \infty$ at a rate of $1/\sqrt{n}$ in distribution, is asymptotically normal and efficient [vdV98, §5.2–5.6]. As a result, $2n \cdot \text{KL}(P, \overline{P}_n)$ converges in distribution to a χ^2 distribution with d-1degrees of freedom. Together with a standard tail bound on χ^2 distributions, this implies the following guarantee: for any fixed $d \ge 2$, $P \in \mathcal{P}_d$, and $\delta \in (0, 1)$, one has

$$\limsup_{n \to +\infty} \mathbb{P}_P\Big(\mathrm{KL}(P, \overline{P}_n) \ge \frac{d + 2\log(1/\delta)}{n}\Big) \le \delta.$$
(3)

This guarantee features the optimal dependence on the dimension d, confidence level $1 - \delta$ and sample size, which may serve as a benchmark for an ideal upper bound.

On the other hand, a significant limitation of the guarantee (3) is that it is purely asymptotic, in that it holds in the limit of sufficiently large sample size with all other parameters being kept fixed. This is at odds with the modern paradigm of high-dimensional models, where the dimension d may be large and possibly comparable to the sample size n. Likewise, one may be interested in high-confidence bounds (that is, in small values in δ), as well as in guarantees that hold uniformly over all distributions $P \in \mathcal{P}_d$. All of these considerations highlight the limitations of purely asymptotic guarantees, and instead call for a quantitative, non-asymptotic analysis.

It should be emphasized that the lack of uniformity over the distribution P of the pointwise asymptotic guarantee (3) is not merely an artifact of its formulation, but instead reflects a fundamental limitation of the estimator \overline{P}_n itself. Indeed, the empirical distribution \overline{P}_n assigns a probability of 0 to classes that do not appear; such a configuration may occur in a finite sample, especially in the presence of rare classes. Hence, the MLE is generally inadequate for the purpose of density estimation in relative entropy, due to its propensity to underestimate uncertainty and to produce overly sharp probability estimates.

Laplace smoothing. In order to mitigate this shortcoming of the MLE, a natural approach is to "smooth out" or regularize the empirical distribution, by assigning some probability to all classes. Arguably the simplest and most classical method to achieve this is the *add-one smoothing* technique, also known as *Laplace rule of succession* [Lap25], which consists in adding 1 to the count of each class. The Laplace estimator is then given by $\hat{P}_n = (\hat{p}_1, \ldots, \hat{p}_d)$, where

$$\widehat{p}_j = \frac{N_j + 1}{n+d} \qquad \text{for } j = 1, \dots, d.$$
(4)

This method was first proposed (in the case d = 2) by Laplace [Lap25, p. 23] in his treatise on probability. Laplace deduced this estimator from what would now be called a Bayesian approach, of which it constitutes one of the earliest instances. Indeed, it coincides with the Bayes predictive posterior distribution under a uniform prior on the probability simplex \mathcal{P}_d . This classical method has found use in various fields, including universal coding (see, e.g., [CT06, p. 435]) and natural language processing (e.g., [JM25, p. 46]). We also note in passing that a closely related method (adding 1/2 to the count of each class) was proposed by Krichevsky and Trofimov [KT81].

The Laplace estimator P_n turns out to achieve an optimal bound in expectation [Cat97, MG22]: for any $P \in \mathcal{P}_d$ and $n, d \ge 2$, one has

$$\mathbb{E}_{P}\left[\mathrm{KL}(P,\widehat{P}_{n})\right] \leqslant \log\left(1 + \frac{d-1}{n+1}\right) \leqslant \frac{d}{n}.$$
(5)

This matches the asymptotic rate of the MLE, but now non-asymptotically and uniformly over all distributions on $\{1, \ldots, d\}$. At the same time, the in-expectation bound (5) falls short of constituting a non-asymptotic analogue of the asymptotic tail bound (3), as it only provides limited information on the tails of the estimation error $\text{KL}(P, \hat{P}_n)$.

Main questions. In this work, we investigate the best possible high-probability guarantees for estimating discrete distributions in relative entropy, either through the Laplace estimator or other procedures. Specifically, the previous discussion naturally raises the following questions:

1. Does there exist a constant c > 0 such that, for any $n \ge d \ge 2$ and $\delta \in (0, 1/2)$, there exists an estimator \widehat{P}_n for which

$$\sup_{P \in \mathcal{P}_d} \mathbb{P}_P\left(\mathrm{KL}(P, \widehat{P}_n) \ge c \, \frac{d + \log(1/\delta)}{n}\right) \le \delta \,? \tag{6}$$

If not, then what is the best possible uniform high-probability guarantee?

2. Does the Laplace estimator \widehat{P}_n achieve the ideal high-probability bound (6)? If not, then what is the best high-probability guarantee for the Laplace estimator?

1.2 Existing guarantees

The previous questions will be addressed in the following sections, but before this, we survey existing high-probability guarantees for estimation of discrete distributions in relative entropy.

High-probability guarantees for the Laplace estimator. As a starting point, the inexpectation bound (5) for the Laplace estimator \hat{P}_n implies (by Markov's inequality) the following bound: for every distribution $P \in \mathcal{P}_d$, with probability at least $1 - \delta$ one has

$$\mathrm{KL}(P,\widehat{P}_n) \leqslant \frac{d}{n\delta} \,. \tag{7}$$

However, this naïve bound is significantly worse than the ideal asymptotic bound (3). For instance, it only shows that a bound of order d/n holds with constant probability, rather than (asymptotically) with probability at least $1 - e^{-d}$ as in (6).

A sequence of recent works has progressively tightened the bound (7). First, Bhattacharyya, Gayen, Price, and Vinodchandran [BGPV21, Theorem 6.1] established a concentration inequality for $\text{KL}(P, \hat{P}_n)$, which implies the following bound: for every $P \in \mathcal{P}_d$, with probability $1 - \delta$,

$$\operatorname{KL}(P, \widehat{P}_n) \lesssim \frac{d \log(n) \log(d/\delta)}{n},$$
(8)

where we use the notation $A \leq B$ to mean that $A \leq cB$ for some universal constant c. While this guarantee significantly improves over the bound (7) for small values of δ , it falls short of the asymptotic bound (3) due to the fact that the dimension d and deviation term $\log(1/\delta)$ are multiplied, rather than decoupled as in (6). A significantly improved bound was established by Han, Jana and Wu [HJW23, Lemma 17], who showed that, with probability $1 - \delta$,

$$\operatorname{KL}(P, \widehat{P}_n) \lesssim \frac{d + \sqrt{d} \log^3(1/\delta)}{n} \,. \tag{9}$$

This implies an optimal bound of d/n when $\log(1/\delta) \leq d^{1/6}$, but leads to (presumably) suboptimal guarantees in the regime $\log(1/\delta) \gg d^{1/6}$.

Finally, the best available high-probability guarantees on the Laplace estimator are due to Canonne, Sun and Suresh [CSS23], who showed that for some absolute constant c > 0, with probability at least $1 - \delta$,

$$\mathrm{KL}(P,\widehat{P}_n) \leqslant \mathbb{E}_P\left[\mathrm{KL}(P,\widehat{P}_n)\right] + c \,\frac{\sqrt{d}\log^{5/2}(d/\delta)}{n} \lesssim \frac{d + \sqrt{d}\log^{5/2}(1/\delta)}{n} \tag{10}$$

(where we used (5) and that $\sqrt{d} \log^{5/2} d \leq d$). In particular, this removes a $\sqrt{\log(1/\delta)}$ factor in the deviation term compared to (9), leading to optimal guarantees in the larger range $\log(1/\delta) \leq d^{1/5}$. Nevertheless, the bound still deteriorates in the regime $\log(1/\delta) \gg d^{1/5}$.

Upper bound for an alternative estimator. Recently, van der Hoeven, Zhivotovskiy and Cesa-Bianchi [vdHZCB23] proposed an alternative estimator \hat{P}_n^{HZC} , based on a high-probability online-to-batch conversion scheme, which achieves the following high-probability guarantee: for any distribution P, with probability at least $1 - \delta$ one has

$$\mathrm{KL}(P, \widehat{P}_n^{\mathrm{HZC}}) \lesssim \frac{d + \log(n) \log(1/\delta)}{n} \,. \tag{11}$$

In many regimes of interest, this constitutes the best known high-probability guarantee in the literature, for any estimator. On the other hand, this result raises important questions. First, the bound (11) features an additional log *n* factor compared to the asymptotic rate (3), which can be avoided in some regimes (e.g., in light of (10)), leaving open the question of the best possible statistical guarantees. Second, the estimator \hat{P}_n^{HZC} is computationally expensive: it requires integrating non-log-concave functions over the probability simplex, the cost of which a priori scales exponentially in the dimension *d* via a grid-based approach. This raises the question of whether or not the problem of high-probability estimation of discrete distributions exhibits a statistical-computational trade-off, that is, if a computational cost exponential in *d* is necessary to achieve a guarantee as strong as (11).

1.3 Paper outline

This paper is organized as follows. In Section 2, we describe the best possible high-probability guarantee on the Laplace estimator (Theorems 1 and 2), and show in particular that this method is optimal among "confidence-independent" estimators. In Section 3, we characterize the best possible uniform guarantees for any estimator (Theorem 3 and 4); the upper bound is achieved by a simple modification of the Laplace estimator using a confidence-dependent smoothing level.

In Section 4, in order to handle situations where the total number of classes is very large, we study guarantees that depend on the "effective sparsity" of the distribution at hand. We establish in particular a minimax lower bound for estimating sparse distributions that holds with high probability (Proposition 1), and then propose simple estimators using data-dependent smoothing that achieve high-probability guarantees (Theorem 5); these guarantees adapt to two natural "effective sparsity" parameters of the distribution. In Section 5, we present a sharp

high-probability bound on the missing and underestimated masses (Theorem 6), which is used in our analysis of the sparse case but may also be of independent interest.

The proof of our high-probability upper bounds for estimation are provided in Section 6, while Section 7 contains the proofs of lower bounds for estimation. Section 8 is devoted to the proof of Theorem 6 on the missing mass, and Section 9 to the elementary proof of the in-expectation guarantee of Proposition 2. Finally, Section 10 gathers various technical lemmata.

1.4 Related work

High-probability guarantees in relative entropy. As discussed in Section 1.2, our contribution belongs to a series of recent works [BGPV21, HJW23, CSS23, vdHZCB23] on high-probability guarantees for estimation of discrete distributions in relative entropy. The best known guarantee for the Laplace estimator is the upper bound (10) from [CSS23], while the best guarantee (in many regimes) for any estimator is the upper bound (11) from [vdHZCB23].

Other aspects of estimation of discrete distributions. Naturally, estimation of discrete distributions is a basic problem, which has been investigated from various other perspectives in the literature. First, one may consider different loss functions than the relative entropy; we refer to [KOPS15, Can20] (and references therein) for an overview of existing guarantees under various losses. Second, even in relative entropy, the minimax-optimal in-expectation bound (5) has been refined in various ways. Braess and Sauer [BS04] characterize asymptotically optimal numerical constants in the minimax expected relative entropy risk, in the regime where d is fixed while $n \to \infty$. In another direction, Orlitsky and Suresh [OS15] consider a more demanding competitive optimality criterion, in the spirit of the empirical Bayes paradigm [Rob51, Goo53].

Concentration properties of the empirical distribution. A relevant but distinct question concerns the concentration properties of the empirical distribution \overline{P}_n . As discussed above, the relative entropy $\operatorname{KL}(P, \overline{P}_n)$ does not enjoy distribution-free concentration properties, since \overline{P}_n may assign zero probability to classes with positive true probability. On the other hand, the opposite configuration cannot happen: a class j with $p_j = 0$ cannot appear in the sample, which qualitatively suggests that the *reverse* relative entropy $\operatorname{KL}(\overline{P}_n, P)$ may be well-behaved. This is indeed the case: the theory of large deviations suggests that the reverse relative entropy $\operatorname{KL}(\overline{P}_n, P)$ sharply encodes the concentration properties of the empirical distribution. Specifically, a classical inequality [CT06, Theorem 11.2.1 p. 356] established by the so-called "method of types" [Csi98] states that, for any $n, d \ge 2, P \in \mathcal{P}_d$ and $\delta \in (0, 1)$, one has

$$\mathbb{P}_P\left(\mathrm{KL}(\overline{P}_n, P) \ge \frac{d\log(n+1) + \log(1/\delta)}{n}\right) \le \delta.$$
(12)

While this bound is distribution-free, it features a presumably suboptimal $\log(n+1)$ factor. This classical bound has been tightened in a series of recent works [MJT⁺20, Agr20, GR20, BP23, Agr22] on concentration of the reverse relative entropy. In particular, it follows from [Agr22, Corollary 1.7] (although this could also be deduced up to constants from the earlier work [Agr20]) that, for any $n, d \ge 2$, $P \in \mathcal{P}_d$ and $\delta \in (0, 1)$,

$$\mathbb{P}_{P}\left(\mathrm{KL}(\overline{P}_{n}, P) \geqslant \frac{6d + 6\log(1/\delta)}{n}\right) \leqslant \delta.$$
(13)

This deviation bound effectively settles the probabilistic question of optimal concentration of the empirical distribution. In this work, we study the complementary statistical question of optimal high-probability estimation guarantees in relative entropy $\text{KL}(P, \hat{P}_n)$.

Missing mass. As part of our analysis, we study the tail behavior of the "missing mass". We refer to Section 5 for a discussion of existing high-probability bounds on this quantity.

1.5 Notation

Throughout this work, we let $n \ge 1$ denote the sample size and $d \ge 2$ the number of classes. If A is a finite set, we denote by |A| its cardinality. We identify the set of probability distributions on $[d] = \{1, \ldots, d\}$ with the set of probability vectors $\mathcal{P}_d = \{(p_1, \ldots, p_d) \in \mathbb{R}^d_+ : \sum_{j=1}^d p_j = 1\}$, where for $1 \le j \le d$ we let p_j denote the probability of the class j. For $j \in \{1, \ldots, d\}$, we let $\delta_j \in \mathcal{P}_d$ denote the Dirac mass at j, identified with the j-th basis vector in \mathbb{R}^d . Given two probability distributions $P = (p_1, \ldots, p_d) \in \mathcal{P}_d$ and $Q = (q_1, \ldots, q_d) \in \mathcal{P}_d$, we define the Kullback-Leibler divergence or relative entropy between P and Q by

$$\operatorname{KL}(P,Q) = \sum_{j=1}^{d} p_j \log\left(\frac{p_j}{q_j}\right),$$

with the convention that $p \log(p/q)$ equals 0 if p = 0, and $+\infty$ if p > 0 but q = 0. For $u, v \in \mathbb{R}^+$, we let $D(u, v) = u \log(\frac{u}{v}) - u + v$ with similar conventions. Since $\sum_{j=1}^d p_j = \sum_{j=1}^d q_j = 1$, we have

$$\operatorname{KL}(P,Q) = \sum_{j=1}^{d} D(p_j, q_j).$$
(14)

Additionally we define the function $h : \mathbb{R}^+ \to \mathbb{R}^+$ by $h(x) = x \log x - x + 1$ for x > 0 and h(0) = 1, so that $D(u, v) = v \cdot h(u/v)$ for $u, v \in \mathbb{R}^+$.

Given a distribution $P \in \mathcal{P}_d$, the sample X_1, \ldots, X_n is comprised of n i.i.d. random variables with distribution P. We use the notations \mathbb{P}_P and \mathbb{E}_P to respectively denote probabilities and expectations when the distribution of (X_1, \ldots, X_n) is $P^{\otimes n}$. For $j = 1, \ldots, d$, we defined the count of the class j as its number of occurrences in the sample X_1, \ldots, X_n , namely

$$N_j = N_{j,n} = \sum_{i=1}^n \mathbf{1}(X_i = j).$$
(15)

An estimator is a map $\Phi : [d]^n \to \mathcal{P}_d$, which we identify (following a standard convention) with the random variable $\widehat{P}_n = \Phi(X_1, \ldots, X_n)$ taking values in \mathcal{P}_d .

Finally, for any $\lambda \in \mathbb{R}^+$, we denote by $\mathsf{Poisson}(\lambda)$ the Poisson distribution with intensity λ , which assigns a probability of $e^{-\lambda}\lambda^k/k!$ to any non-negative integer $k \in \mathbb{N}$.

2 Optimal guarantees for the Laplace estimator

In this section, we consider the question of optimal high-probability guarantees for the classical Laplace (add-one) estimator, defined by (4). In Section 2.1, we state our main upper bound, while in Section 2.2 we provide a matching lower bound for a large class of estimators that includes the Laplace estimator.

2.1 Upper bound for the Laplace estimator

Our first main result is a finite-sample high-probability bound for the Laplace estimator.

Theorem 1. For any $n \ge 12, d \ge 2$ and $P \in \mathcal{P}_d$, the Laplace estimator \widehat{P}_n defined by (4) achieves the following guarantee: for any $\delta \in (e^{-n/6}, e^{-2})$,

$$\mathbb{P}_P\left(\mathrm{KL}(P, \widehat{P}_n) \ge 110000 \, \frac{d + \log(1/\delta) \log \log(1/\delta)}{n}\right) \le 4\delta \,. \tag{16}$$

The proof of Theorem 1 is provided in Section 6.4.

We note in passing that the condition $\delta > e^{-n/6}$ in Theorem 1 is not restrictive, as it constitutes the nontrivial regime. Indeed, when $n \ge d$ and $\delta = e^{-n/6}$, the upper bound (16) is of order log *n*. But an upper bound $\operatorname{KL}(P, \widehat{P}_n) \le \log(2n)$ actually holds deterministically (thus for $\delta = 0$) since for every $j = 1, \ldots, d$ one has $\widehat{p}_j \ge 1/(n+d) \ge 1/(2n)$, so that $p_j/\widehat{p}_j \le 2n$.

Theorem 1 improves the previously best known upper bound (10) on the Laplace estimator from [CSS23], which is of order $\{d + \sqrt{d}\log^{5/2}(1/\delta)\}/n$. Since $\delta > e^{-n/6}$ and thus $\log \log(1/\delta) \leq \log n$, it also improves the previously best known upper bound (11) of order $\{d + \log(n)\log(1/\delta)\}/n$ for this problem, achieved by the (computationally involved) estimator from [vdHZCB23]. This shows in particular that such guarantees can be achieved in a computationally efficient manner, specifically in time linear in n.

A curious feature of the upper bound (16) is that it exhibits non-standard tails, in the form of the $\log(1/\delta) \log \log(1/\delta)/n$ deviation term. This should be contrasted with the more common quantiles of exponential and Poisson variables, respectively of order $\log(1/\delta)$ and $\frac{\log(1/\delta)}{\log \log(1/\delta)}$. In particular, this tail bound is *super-exponential*, which points to a technical difficulty in its proof: it cannot be established by the standard Chernoff method based on the moment generating function (m.g.f.). Indeed, super-exponential tails lead to an infinite m.g.f.¹, and conversely a finite m.g.f. would lead to sub-exponential tails. For this reason, the proof does not use the m.g.f. and instead proceeds by controlling raw moments (L^p norms).

We refer to Section 6 for a description of the main tools in the proof of Theorem 1, which also serve for high-probability upper bounds stated in subsequent sections. Roughly speaking, the key step in the analysis is to control the contribution to the error $\text{KL}(P, \hat{P}_n)$ of classes $j = 1, \ldots, d$ for which the Laplace estimate significantly underestimates the true probability, as such classes may lead to a potentially large errors.

2.2 Lower bound for confidence-independent estimators

It should be noted that the uniform non-asymptotic high-probability bound of Theorem 1 for the Laplace estimator exceeds the asymptotic tail bound (3) of the MLE (or Laplace estimator) by a factor of $\log \log(1/\delta)$ in the deviation term. This raises the question of whether this extra factor is necessary, or whether it can be removed by a more precise analysis.

As it turns out, the extra $\log \log(1/\delta)$ factor is necessary, not only for the Laplace estimator but in fact for any "confidence-independent" estimator $\hat{P}_n = \Phi(X_1, \ldots, X_n)$ that does not depend on the desired confidence level $1 - \delta$.

Theorem 2. Let $n \ge d \ge 4000$ and $\kappa \ge 1$. Let $\Phi : [d]^n \to \mathcal{P}_d$ be an estimator such that, denoting $\widehat{P}_n = \Phi(X_1, \ldots, X_n)$ we have for any $P \in \mathcal{P}_d$:

$$\mathbb{P}_P\Big(\mathrm{KL}(P,\widehat{P}_n) \leqslant \frac{\kappa d}{n}\Big) > 0.$$
(17)

Then, for any $\delta \in (e^{-n}, e^{-16\kappa^2})$, there exists a distribution $P \in \mathcal{P}_d$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \geqslant \frac{d + \log(1/\delta)\log\log(1/\delta)}{5000\,n}\right) \geqslant \delta\,.$$
(18)

The proof of Theorem 2 (based on Lemma 12) can be found in Section 7.1. The main idea of the proof is that, in order to achieve the guarantee (17), the estimator \hat{P}_n cannot be "too

¹Technically speaking, the m.g.f. of the variable $\operatorname{KL}(P, \widehat{P}_n)$ is actually finite, due to the deterministic bound $\operatorname{KL}(P, \widehat{P}_n) \leq \log(2n)$. However, this (loose) bound depends on n, hence bounds based on the m.g.f. would lead to deviation terms with an additional dependence on n compared to (16).

far" from the empirical distribution \overline{P}_n . But in this case, it may significantly underestimate the probability of some classes with small probability δ , leading to the lower bound (18).

A few comments may help clarify the meaning of Theorem 2. First, one should think of the parameter $\kappa \ge 1$ as an absolute constant. The condition (17) states that the estimator \hat{P}_n achieves a bound of order d/n with positive probability for any distribution P. This holds in particular when the estimator \hat{P}_n achieves an optimal in-expectation bound $\mathbb{E}_P[\mathrm{KL}(P, \hat{P}_n)] \le \kappa d/n$, and more generally when the estimator achieves an optimal bound of $\kappa d/n$ in the regime of constant (bounded away from 0 and 1) confidence level. The first condition applies with $\kappa = 1$ to the Laplace estimator, in light of its in-expectation bound (5). We refer to such estimators, including those that are optimal in expectation, as "confidence-independent".

The content of Theorem 2 is that such an estimator must necessarily incur the $\log \log(1/\delta)$ factor for small values of δ . In other words, the extra $\log \log(1/\delta)$ factor in the high-confidence regime is a necessary price to pay for optimality in the constant-confidence regime.

When compared with the upper bound of Theorem 1, the lower bound of Theorem 2 implies that the Laplace estimator is optimal in a minimax sense, over a large class of estimators. This may be of interest in itself given the simplicity of this procedure.

3 Minimax-optimal guarantees for confidence-dependent estimators

An obvious restriction in the lower bound of Theorem 2 is that it only applies to "confidenceindependent" estimators—in particular, to those that achieve an optimal d/n guarantee with constant probability. This leaves open the possibility that, for a given $\delta \in (0, 1/2)$, improved guarantees with probability $1 - \delta$ may be achieved by an estimator $\hat{P}_{n,\delta} = \Phi_{\delta}(X_1, \ldots, X_n)$ that depends on δ ; that is, which is tuned for the desired confidence level $1 - \delta$, at the cost of being suboptimal at constant confidence levels. We call such an estimator "confidence-dependent".

In this section, we investigate optimal high-probability guarantees for confidence-dependent estimators; Section 3.1 is dedicated to the upper bound, and Section 3.2 to the lower bound.

3.1 Upper bound via confidence-dependent smoothing

Since the gap between the asymptotically ideal tail bound (3) and the non-asymptotic upper and lower bounds of Section 2 consists in an extra $\log \log(1/\delta)$ factor in the deviation term, the question is whether this factor can be improved by a confidence-dependent estimator. Theorem 3 below answers this question in the affirmative:

Theorem 3. For any $n \ge 12, d \ge 2$ and $\delta \in (e^{-n/6}, e^{-2})$, define the estimator $\widehat{P}_{n,\delta} = (\widehat{p}_1, \ldots, \widehat{p}_d)$ by, for $j = 1, \ldots, d$,

$$\widehat{p}_j = \frac{N_j + \lambda_\delta}{n + \lambda_\delta d} \qquad where \qquad \lambda_\delta = \max\left\{1, \frac{\log(1/\delta)}{d}\right\}.$$
(19)

Then, for any $P \in \mathcal{P}_d$, we have

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_{n,\delta}) \ge 110000 \, \frac{d + \log(d)\log(1/\delta)}{n}\right) \le 4\delta \,. \tag{20}$$

The estimator (19) in Theorem 3 can be seen as a confidence-dependent modification of the Laplace estimator. Specifically, in the low and moderate-confidence regime where $\delta \ge e^{-d}$, the estimator $\hat{P}_{n,\delta}$ coincides with the Laplace estimator. On the other hand, in the high-confidence regime where $\delta < e^{-d}$, it smooths the empirical distribution more strongly than the Laplace

estimator, with a confidence-dependent level of smoothing—the higher the desired confidence level, the stronger the smoothing.

We now comment on the quantitative upper bound of Theorem 3. First, note that the condition $\delta > e^{-n/6}$ is not restrictive, since for $\delta = e^{-n/6}$ and $n \ge d$ the upper bound (20) is of order log d; but a deterministic upper bound of log d may be achieved by letting $\hat{P}_n = (1/d, \ldots, 1/d)$ be the uniform distribution. In fact, for $\delta \le e^{-n/6}$ the estimator $\hat{P}_{n,\delta}$ also satisfies a deterministic bound $\mathrm{KL}(P, \hat{P}_{n,\delta}) \le \log(7d)$, as $\hat{p}_j \ge 1/(7d)$ for $j = 1, \ldots, d$.

Second, observe that the upper bound (20) provides an improvement over the best possible guarantee for confidence-independent estimators, as characterized by Theorems 1 and 2. This amounts to saying that, regardless of $\delta \in (0, e^{-2})$ and $d \ge 2$, one has

$$\frac{d + \log(d)\log(1/\delta)}{n} \lesssim \frac{d + \log(1/\delta)\log\log(1/\delta)}{n} \,. \tag{21}$$

To see why (21) holds, consider the following two cases. If $\log(1/\delta) \leq d/\log d$, then the left-hand side of (21) is of order d/n and the bound holds. On the other hand, if $\log(1/\delta) \geq d/\log d \geq \sqrt{d}$, then $\log d \leq \log \log(1/\delta)$ and the bound (21) also holds.

Hence, Theorems 2 and 3 together imply an advantage of confidence-dependent estimators over confidence-independent ones. We note that such an advantage has been previously observed for a different problem, namely mean estimation under heavy-tailed noise [DLLO16, Cat12]. It is notable that it also manifests itself in the basic problem of estimation of discrete distributions, in the absence of robustness constraints.

We refer to Section 6.5 for the proof of Theorem 3, which shares a common structure with that of Theorem 1, although some details differ. Again, the core of the analysis is to control the contribution to the relative entropy of classes whose frequency is underestimated, which is technically achieved through sharp moment estimates on the corresponding terms.

3.2 Lower bound for confidence-dependent estimators

While the upper bound of Theorem 3 for the estimator $\hat{P}_{n,\delta}$ circumvents the lower bound of Theorem 2 for confidence-independent estimators, it still exceeds the ideal asymptotic tail bound (3) by a factor of log *d* in the deviation term. Theorem 4 below shows that this extra log *d* factor cannot be avoided, even for confidence-dependent estimators:

Theorem 4. Let $n \ge d \ge 5000$ and $\delta \in (e^{-n}, e^{-1})$. For any estimator $\Phi = \Phi_{\delta} : [d]^n \to \mathcal{P}_d$, there exists a distribution $P \in \mathcal{P}_d$ such that, letting $\widehat{P}_n = \widehat{P}_{n,\delta} = \Phi_{\delta}(X_1, \ldots, X_n)$ we have

$$\mathbb{P}_P\left(\mathrm{KL}(P,\hat{P}_n) \geqslant \frac{d + \log(d)\log(1/\delta)}{5000\,n}\right) \geqslant \delta\,.$$
(22)

Together, Theorems 3 and 4 characterize, up to universal constant factors, the minimax highprobability risk (in other words, the "sample complexity") for estimation of discrete distributions in relative entropy. An interesting feature of this lower bound is that it exceeds the asymptotic rate (3) by a log d term; this establishes a separation between asymptotic guarantees and uniform non-asymptotic guarantees.

The proof of Theorem 4 relies on the following lemma, which is proved in Section 7.2.

Lemma 1. Let $n \ge d \ge 2$ and $\delta \in (e^{-n}, e^{-1})$. There exists a set $\mathcal{F} = \mathcal{F}_{n,d,\delta} \subset \mathcal{P}_d$ of d distributions with support size at most 2 such that the following holds. For any estimator $\widehat{P}_n = \Phi(X_1, \ldots, X_n)$, there exists a distribution $P \in \mathcal{F}$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \geqslant \frac{\log(d)\log(1/\delta)}{14\,n}\right) \geqslant \delta\,.$$
(23)

The idea behind Lemma 1 is quite simple, and can be summarized as follows: let P be either a Dirac mass at 1, namely $P = \delta_1 = (1, 0, ..., 0)$, or a mixture of the form $(1 - \frac{\log(1/\delta)}{n})\delta_1 + \frac{\log(1/\delta)}{n}\delta_j$ for some j = 2, ..., d. Then, regardless of which of these distributions P is, with probability at least of order δ , only the first class is observed $(X_1 = \cdots = X_n = 1)$. In this case, it is impossible to tell which distribution P is. In order to avoid incurring a large error in the first case where $P = \delta_1$, one must assign a large probability to the first class, and thus a low probability to all remaining classes. However, this entails a large error in the second case. Specifically, the extra $\log(d)$ factor comes from the fact that in the second case where P puts some mass on a class $j \in \{2, ..., d\}$, no information on j is available if only the first class is observed. Hence, all the remaining mass must be shared among the d-1 remaining classes j = 2, ..., d, effectively dividing by d-1 the per-class probability and thus inflating the relative entropy.

We note in passing that Lemma 1 has broader implications to the theory of aggregation and density estimation, beyond the present context of estimation of discrete distributions. Indeed, it is known from [YB99, Cat04] (building on an idea of Barron [Bar87]), that for any finite model/class \mathcal{F} of distributions, given an i.i.d. sample of size n from an unknown distribution $P \in \mathcal{F}$, there exists an estimator \hat{P}_n such that $\mathbb{E}_P[\mathrm{KL}(P, \hat{P}_n)] \leq \log(|\mathcal{F}|)/n$ for all $P \in \mathcal{P}$. This naturally raises the question of whether a corresponding ideal high-probability guarantee, of the form $\mathbb{P}_P(\mathrm{KL}(P, \hat{P}_n) \geq C \{\log |\mathcal{F}| + \log(1/\delta)\}/n) \leq \delta$ for some absolute constant C, can also be achieved, possibly by another estimator. Lemma 1 shows that this is not the case: since $|\mathcal{F}_{n,d,\delta}| = d$, the best tail bound one may hope for is of order $\log(|\mathcal{F}|) \log(1/\delta)/n$.

4 Adaptation to the effective support size

The results of Section 3 settle the question of the minimax-optimal high-probability guarantees. In addition, the results of Section 2 show that the classical Laplace estimator is rather close to being optimal, in the same sense.

While these results attest to the soundness of the Laplace rule—and of its confidencedependent modification—, one should keep in mind that in several applications such as natural language processing, its use has been supplanted by that of more advanced methods, such as Kneser-Ney smoothing [KN95, CG99].

This apparent contradiction between theory and practice comes from the fact that we have so far focused on minimax guarantees that hold uniformly over all distributions. By itself, this uniformity is a strength, as it requires no restrictive assumptions on the true distribution. In addition, it is not obviously detrimental, given that the limiting risk (3) of the MLE does not depend on the distribution $P \in \mathcal{P}_d$, as long as the latter assigns positive probability to all classes.

Nevertheless, this last restriction points towards a possible improvement: if only s < d classes have positive probability, then by (3), the limiting high-probability risk of the MLE scales as $\{s + \log(1/\delta)\}/n$, which can be much smaller than d/n if $s \ll d$. This suggests that, more generally, the complexity of distribution estimation should be governed by some notion of support size, such as the number of positive or (for a fixed sample size) large enough probabilities.

Indeed, while the minimax rate of d/n cannot be improved for worst-case distributions that are approximately uniform over $\{1, \ldots, d\}$, distributions that arise in practice often exhibit a non-uniform structure, with a small number of frequent classes and a large number of less frequent classes. Under such a configuration, it may be possible to estimate the distribution Peven when the sample size n is smaller than the total number d of classes.

In this section, we consider high-probability upper and lower bounds for estimation of discrete distributions, which depend on suitable notions of "support size" of the distribution. Section 4.1 contains minimax lower bounds for estimation of sparse distributions, while Section 4.2 is devoted to high-probability upper bounds for suitable adaptive estimators.

4.1 Minimax lower bounds for sparse distributions

We start by establishing minimax lower bounds on the best possible estimation guarantee for "sparse" distributions. Since we are interested in lower bounds, we consider a small class of sparse distributions P, namely distributions with support size at most $s \leq d$. Naturally, such lower bounds transfer to larger classes, that is, to less stringent notions of sparsity.

For any $P \in \mathcal{P}_d$, we let $\operatorname{supp}(P) = \{1 \leq j \leq d : p_j > 0\}$ denote the support of P. In addition, for $1 \leq s \leq d$ we define the class $\mathcal{P}_{s,d} = \{P \in \mathcal{P}_d : |\operatorname{supp}(P)| \leq s\}$ of probability distributions on $\{1, \ldots, d\}$ supported on at most s elements. We call such distributions s-sparse.

Our main lower bound for the class of s-sparse distributions is the following:

Proposition 1. Let $n, d \ge 2$. For any $1 \le s \le \min(n, d/55)$ and any estimator $\widehat{P}_n = \Phi(X_1, \ldots, X_n)$, there exists a distribution $P \in \mathcal{P}_{s,d}$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P, \widehat{P}_n) \ge \frac{s\log(ed/s)}{300\,n}\right) \ge 1 - 3\exp\left(-\frac{s}{35}\right). \tag{24}$$

We refer to Section 7.3 for the proof of Proposition 1. In particular, letting $s = \lfloor d/55 \rfloor$ in Proposition 1 shows the following: for any $n \ge d \ge 110$ and any estimator $\hat{P}_n = \Phi(X_1, \ldots, X_n)$, there exists a distribution $P \in \mathcal{P}_d$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \ge \frac{d}{4600\,n}\right) \ge 1 - 3\exp\left(-\frac{d}{3000}\right). \tag{25}$$

Two aspects of the lower bound of Proposition 1 merit further discussion.

First, the minimax lower bound is of order $s \log(ed/s)/n$, which exceeds by a $\log(ed/s)$ factor the rate of estimation of a distribution P with known support of size s. This extra factor, which is standard in the context of estimation under sparsity, can be seen as the price of estimation of the support of P: indeed, one has $s \log(ed/s) \approx \log \binom{d}{s}$, where $\binom{d}{s}$ is the number of possible supports of size s. We note that the rate of $s \log(ed/s)/n$ coincides with the minimax rate of estimation of a sparse vector under Gaussian noise (e.g., [Wai19, p. 156]).

However, despite the similarity in rates, there are qualitative differences between the Gaussian setting and the multinomial setting that we consider. Indeed, it follows from the results of Agrawal [Agr22] (see Lemma 4 below) that when $P \in \mathcal{P}_{s,d}$, the empirical distribution achieves an upper bound in squared Hellinger distance of order s/n with probability at least $1 - e^{-s}$. This rate no longer features the extra $\log(ed/s)$ factor; by contrast, the extra $\log(ed/s)$ factor does appear in the Gaussian model, even when the error is measured in squared Hellinger distance. Roughly speaking, this difference stems from the fact that in the multinomial model, unlike in the Gaussian model, the amount of "noise" in a coordinate $j = 1, \ldots, d$ decays when the magnitude of the corresponding coefficient parameter p_j goes to 0.

In particular, the optimal rate of estimation of sparse discrete distributions in Kullback-Leibler divergence exceeds the optimal rate under squared Hellinger divergence by a $\log(ed/s)$ factor. This gap comes from the fact that the empirical distribution only identifies a subset of the support of the true distribution, and the price of missing part of the support of P is higher in Kullback-Leibler divergence than in squared Hellinger distance.

A second important feature of the lower bound of Proposition 1 (and of its consequence (25) in the non-sparse case) is that it holds with high probability. Such a statement is stronger than minimax lower bounds that are usually found in the literature, which hold either in expectation or with constant probability. For estimation of non-sparse discrete distributions, lower bounds in expectation are proven in [HJW15, KOPS15], while [CKT24] obtains a lower bound with constant probability. To appreciate the difference, consider the d/n lower bound in the non-sparse case. Then, a d/n lower bound in expectation is in principle compatible with the following

behavior: a risk of 1 with probability d/n, and a risk of $0 \ll d/n$ with high probability 1-d/n. In contrast, a stronger lower bound with constant probability rules out such a behavior: it asserts that a lower bound of order d/n must hold with probability bounded away from 0, say 10% or 80%. However, even this lower bound does not rule out the possibility that the estimator achieves an error significantly smaller than d/n (say, of 0) with nontrivial probability—respectively, with probability 90% or 20%. A high-probability lower bound such as (25) excludes such a behavior: it asserts that a lower bound of order d/n holds with overwhelming probability, effectively ruling out the possibility that the estimator "gets lucky". In addition, the probability with which the lower bound (25) holds, which behaves as $1 - e^{-d/c}$, is best possible, as shown by the convergence of the properly scaled risk of the MLE to a χ^2_{d-1} distribution as $n \to \infty$.

Proposition 1 is, to our knowledge, the first minimax lower bound that holds with probability exponentially close to 1. We believe that such lower bounds could be of broader interest in statistical estimation, beyond the present context of estimation of discrete distributions.

We note that the classical method for proving lower bounds in statistical estimation (put forward by Ibragimov and Has'minskii [IH81], see [Tsy09, Chapter 2] and [Wai19, Chapter 15]), which consists in a reduction from estimation to testing, followed by a lower bound for testing though Fano's inequality, actually provides a high-probability lower bound. However, the probability estimate that the commonly used version of this inequality (see e.g. [Wai19, Proposition 5.12 p. 502]) leads to is weaker than (25). Specifically, it leads to a lower bound that holds with probability 1 - c/d for some constant c, in contrast with the lower bound (25) with probability $1 - e^{-d/c}$. As it turns out, in the non-sparse case, one could also obtain a probability of $1 - e^{-d/c}$ by instead using the sharp version of Fano's inequality (e.g., [Tsy09, Lemma 2.10]) and further simplifying the resulting bound. However, the reduction from estimation to testing is not straightforward in the sparse case, essentially because the Kullback-Leibler divergence does not behave like a distance over sparse distributions with distinct supports.

The proof of Proposition 1, which can be found in Section 7.3, is not based on a reduction to testing, but instead on the probabilistic method—in other words, on a Bayesian lower bound. Specifically, we consider suitable distributions with randomly chosen support of size s, and show that on average over the random draw of such a distribution, the probability that any estimator achieves a risk at least of order $s \log(ed/s)/n$ is overwhelmingly close to 1.

In addition to the high-probability lower bound of Proposition 1, we now provide a *low-probability* lower bound. Low-probability lower bounds (such as Theorems 2 and 4 in previous sections) are more common in the literature; their interest stems from the fact that they are precisely the converse of high-probability upper bounds, and can thus attest to the optimality of such upper bounds. Specifically, combining Proposition 1 with Lemma 1 above—which applies to our present setting as it involves 2-sparse distributions—leads to the following:

Corollary 1. Let $n, d \ge 110$, $s \in \{2, \ldots, \min(n, d/55)\}$ and $\delta \in (e^{-n}, e^{-2})$. For any estimator $\Phi = \Phi_{s,\delta} : [d]^n \to \mathcal{P}_d$, there exists a distribution $P \in \mathcal{P}_{s,d}$ such that, letting $\widehat{P}_n = \Phi(X_1, \ldots, X_n)$, we have

$$\mathbb{P}_P\left(\mathrm{KL}(P,\hat{P}_n) \geqslant \frac{s\log(ed/s) + \log(d)\log(1/\delta)}{320\,n}\right) \geqslant \delta\,.$$
(26)

We refer to Section 7.4 for the proof of this result. Corollary 1 recovers the minimax lower bound of Theorem 4 up to constants, by setting $s = \lfloor d/55 \rfloor$.

4.2 Upper bounds for adaptive estimators

Having obtained in Corollary 1 a lower bound on the best possible high-probability guarantee, we now turn to upper bounds. Of particular interest are estimators that achieve such guarantees without prior knowledge of the support size s of P; we call such estimators *adaptive*.

4.2.1 Effective sparsity parameters

Before stating such guarantees, it is worth discussing what an ideal upper bound might look like. Of course, a natural objective would be to obtain minimax high-probability upper bounds under sparsity constraints that match the lower bound of Corollary 1. However, it should be noted that the notion of sparsity (support size) we considered in the previous section is rather restrictive: it excludes distributions with full support, but for which only a small number of classes have significant probability. Intuitively, classes with positive but very small probability should have a limited effect on the estimation error, at least for moderate sample sizes.

These considerations call for identifying notions of "effective support size" or "effective sparsity" that capture the hardness of estimation more accurately than the number of nonzero entries of the distribution. One should expect such notions of effective support size to depend on the sample size n, given that the large-sample asymptotic behavior (3) of the MLE is governed by the number of nonzero entries.

Perhaps the most natural notion of "effective support size of P under sample size n" is the typical number of distinct classes that would appear on an i.i.d. sample of size n from P. As we will explain briefly, this leads to the following definition:

Definition 1. For any distribution $P = (p_1, \ldots, p_d) \in \mathcal{P}_d$ and $n \ge 1$, the effective support size of P at sample size n is the quantity $s_n(P) \in [1, \min(n, d)]$ defined by

$$s_n(P) = \sum_{j=1}^{d} \min(np_j, 1) = \left| \left\{ 1 \le j \le d : p_j \ge 1/n \right\} \right| + \sum_{j: p_j < 1/n} np_j.$$
(27)

It is clear from the first expression that $s_n(P)$ increases with n while $s_n(P)/n$ decreases with n, that $s_1(P) = 1$, that $s_n(P) \leq s$ if $P \in \mathcal{P}_{s,d}$, and that $s_n(P)$ converges as $n \to \infty$ to $s_{\infty}(P) = |\{1 \leq j \leq d : p_j > 0\}|$. Now, denote by D_n the number of distinct classes in the sample X_1, \ldots, X_n , namely

$$D_n = \sum_{j=1}^d \mathbf{1}(N_j \ge 1) = \left| \left\{ X_i : 1 \le i \le n \right\} \right|.$$
(28)

The following fact relates $\mathbb{E}_P[D_n]$ to $s_n(P)$.

Fact 1. For any $P \in \mathcal{P}_d$ and $n \ge 1$, one has $(1 - e^{-1})s_n(P) \le \mathbb{E}_P[D_n] \le s_n(P)$.

Proof. The upper bound comes from the fact that $\mathbb{E}_P[D_n] = \sum_{j=1}^d \mathbb{P}(N_j = 1)$ and that $\mathbb{P}(N_j = 1) = \mathbb{P}(\bigcup_{i=1}^n \{X_i = j\}) \leq \min(np_j, 1)$ by a union bound. For the lower bound, we start with the expression $\mathbb{E}_P[D_n] = \sum_{j=1}^d \{1 - (1 - p_j)^n\}$. Now, if $p_j \geq 1/n$, then $(1 - p_j)^n \leq (1 - 1/n)^n \leq e^{-1}$, hence $1 - (1 - p_j)^n \geq 1 - e^{-1}$. On the other hand, if $p_j \leq 1/n$, then $(1 - p_j)^n \leq e^{-np_j} \leq 1 - (1 - e^{-1})np_j$ by convexity of exp, so that $1 - (1 - p_j)^n \geq (1 - e^{-1})np_j$.

One reason why the quantity $s_n(P)$ is a natural "effective sparsity" parameter is that it controls the error of the empirical distribution in squared Hellinger distance, as shown by Lemma 4 below. To see how it arises here, recall that for an estimator $\hat{P}_n = (\hat{p}_1, \ldots, \hat{P}_d)$, one has

$$\operatorname{KL}(P, \widehat{P}_n) = \sum_{j=1}^d D(p_j, \widehat{p}_j).$$

Now, consider the situation in which for each class $j = 1, \ldots, d$, we are given two options:

- estimate p_j based on the binomial N_j , for instance using the empirical frequency or the binary Laplace estimate $(N_j + 1)/(n + 2)$. By (5), the latter gives $\mathbb{E}[D(p_j, \hat{p}_j)] \leq 1/n$;
- alternatively, resort to an "oracle" that returns an approximation \tilde{p}_j of p_j guaranteed to be of the same order of magnitude: $p_j/2 \leq \tilde{p}_j \leq 2p_j$. This ensures that $D(p_j, \tilde{p}_j) \leq p_j$.

Using the best of the two options for each coordinate leads to an error in relative entropy of order at most $\sum_{j=1}^{d} \min(p_j, 1/n) = s_n(P)/n$. However, one cannot expect a bound of $s_n(P)/n$ to be achievable without the oracle, as this would violate the lower bound of Corollary 1. This is because a class j with small probability $p_j \ll 1/n$ typically does not appear in the sample; in this case, one has no information about the order of magnitude of p_j except that $p_j \lesssim 1/n$. This obstruction corresponds to the difficulty of estimating the support of P discussed in Section 4.1, which leads to an extra $\log(ed/s)$ factor.

In light of this discussion, a natural objective would be to aim for a high-probability upper bound matching the lower bound (26), but with the support size s possibly replaced by the effective support size $s_n(P)$.

An in-expectation version of such a result was established by Falahatgar, Ohannessian, Orlitsky and Pichapati [FOOP17, Theorem 1], who showed that the "absolute discounting" estimator \hat{P}_n satisfies the following guarantee: for any $P \in \mathcal{P}_d$, denoting $s_n = s_n(P)$ one has

$$\mathbb{E}_{P}[\mathrm{KL}(P,\widehat{P}_{n})] \leqslant c \cdot \frac{s_{n} \log(ed/s_{n})}{n}$$
(29)

for some constant c that only depends on the tuning parameter of the estimator.

Before establishing a high-probability extension of this result, we first show that an improvement of the bound (29) is possible even in expectation. To see why, consider a distribution $P \in \mathcal{P}_d$ with support size $s \ll d$, and consider the asymptotic regime where $n \to \infty$. In this case, it follows from (3) that as $n \to \infty$, the MLE \overline{P}_n achieves a risk of order s/n with high probability. On the other hand, again as $n \to \infty$, one has $s_n \to s$, hence the bound (29) scales as $s \log(ed/s)/n$, which exceeds the risk of the MLE by a $\log(ed/s)$ factor.

To understand this discrepancy, recall that the $\log(ed/s)$ factor comes from the fact that the support of P is unknown. Meanwhile, $s_n(P)$ measures the typical number of distinct classes that appear in the sample. But since classes that appear in the sample are known, they do not contribute to the uncertainty about the support of P. In contrast, the only classes that contribute to the uncertainty on the support are those that are missing from the sample².

Hence, one should expect that the effective support size $s_n(P)$ does not suffice to describe the achievable error rate in relative entropy. Instead, the logarithmic term owing to the lack of knowledge of the support should be tied to a different sparsity parameter, which accounts for the contribution of classes that are missing from the sample. We now define such a parameter:

Definition 2. For any distribution $P = (p_1, \ldots, p_d) \in \mathcal{P}_d$ and real number $n \ge 1$, the *effective* missing support size of P at sample size n is the quantity $s_n^{\circ}(P) \in (0, \min(n, d)]$ defined by

$$s_n^{\circ}(P) = \sum_{j=1}^d \min\left(e^{1-np_j}, np_j\right) = \sum_{j: p_j \ge 1/n} e^{1-np_j} + \sum_{j: p_j < 1/n} np_j.$$
(30)

The reason why $s_n^{\circ}(P)$ accounts for the contribution of missing classes is that $s_n^{\circ}(P)/n$ is closely related to the expected "missing mass" (total probability of classes that do not appear in the sample), as shown in Lemma 2 below. We defer to Section 5 for more discussion on the

²To be accurate, the situation is slightly more subtle: classes that do appear, but with an empirical frequency significantly smaller than their theoretical one, also contribute to the inflation of the relative entropy.

behavior of the missing mass. In addition, again by Lemma 2, $s_n^{\circ}(P)$ is also closely related to the expected number of new classes that would appear on an independent sample X_{n+1}, \ldots, X_{2n} of size n (but not in X_1, \ldots, X_n), namely $\mathbb{E}_P[D_{2n} - D_n]$; this justifies the terminology "effective missing support size" of Definition 2. This suggests that $s_n^{\circ}(P)$ (roughly the number of classes that first appear after $\asymp n$ observations) can be seen as a "local" counterpart to the "global" effective sparsity parameter $s_n(P)$ (number of classes that first appear after $\lesssim n$ observations).

It is clear from the second expression in (30) that $s_n^{\circ}(P) \leq s_n(P)$. In addition, from the first expression, the quantity $s_n^{\circ}(P)/n$ decreases with n. On the other hand, contrary to $s_n(P)$, the quantity $s_n^{\circ}(P)$ does not increase with n; in fact, one has $s_n^{\circ}(P) \to 0$ as $n \to \infty$ for any $P \in \mathcal{P}_d$.

To appreciate the difference between the two effective support size parameters, one may use the following rough approximations: $s_n(P)$ usually amounts to the number of classes j such that $p_j \gtrsim 1/n$, while $s_n^{\circ}(P)$ often roughly corresponds to the number of classes j with $p_j \approx 1/n$.

4.2.2 Upper bound in expectation for sparse distributions

As discussed above, the contribution to the estimation error of classes that appear in the sample (with an empirical frequency of the same order as their true frequency) should ideally scale as $s_n(P)/n$, without additional logarithmic factors. On the other hand, the contribution to the estimation error of classes that do not appear in the sample (or appear with an empirical frequency significantly smaller than their true frequency) is governed by the parameter $s_n^{\circ} = s_n^{\circ}(P)$. To anticipate their contribution, consider the situation where s_n° indeed scales as the number of classes j such that $p_j \approx 1/n$. Such classes have a constant probability of not appearing in the sample, in which case their identity is unknown. Thus, there are typically about s_n° missing classes that have a probability of order 1/n. However, assuming that the total number of missing classes, of order s_n°/n , must be shared among roughly d classes. This means each of the roughly s_n° classes j with $p_j \approx 1/n$ is assigned a probability $\hat{p}_j \approx (s_n^{\circ}(P)/n)/d \lesssim 1/n$, leading to a total contribution to the estimation error of order $s_n^{\circ} \times \frac{1}{n} \log(\frac{1/n}{s_n^{\circ}/(dn)}) \approx \frac{s_n^{\circ}}{n} \log(d/s_n^{\circ})$.

The next result (proved in Section 9) shows that a matching upper bound can indeed be achieved by a suitable adaptive estimator:

Proposition 2. Let $\widehat{P}_n^{\text{ad}} = (\widehat{p}_1, \ldots, \widehat{p}_d)$ denote the estimator defined by, for $j = 1, \ldots, d$,

$$\widehat{p}_j = \frac{N_j + \widehat{\lambda}}{n + \widehat{\lambda} d} \quad \text{with} \quad \widehat{\lambda} = \frac{D_n}{d},$$
(31)

where $D_n = \sum_{j=1}^d \mathbf{1}(N_j \ge 1)$ is the number of distinct classes among X_1, \ldots, X_n . Then, for any $n, d \ge 2$ and distribution $P \in \mathcal{P}_d$, letting $s_n = s_n(P)$ and $s_{n/2}^\circ = s_{n/2}^\circ(P)$ one has

$$\mathbb{E}_P\left[\mathrm{KL}(P, \widehat{P}_n^{\mathrm{ad}})\right] \leqslant \frac{2.4s_n + 2s_{n/2}^\circ \log(ed/s_{n/2}^\circ)}{n+1}.$$
(32)

The estimator \widehat{P}_n^{ad} defined by (31) can be viewed as an adaptive modification of the Laplace estimator (4), where the regularization parameter $\widehat{\lambda}$ is chosen in a data-dependent manner, so as to adapt to the "shape" of the distribution P. In the case where $n \gtrsim d$ and the distribution is "dense", namely puts significant probability $p_j \gtrsim 1/d$ to all classes, then typically D_n is of order d and thus $\widehat{\lambda} \simeq 1$, hence the estimator behaves similarly to the Laplace estimator. On the other hand, when the distribution is highly sparse, then typically $D_n \ll d$ and thus the regularization parameter $\widehat{\lambda} \ll 1$ is much smaller than for the Laplace estimator—it may be as low as 1/d, in the extreme case where only one class appears in the sample. The specific choice of tuning parameter $\widehat{\lambda} = D_n/d$ is motivated by the decomposition of the risk in Lemma 3 below.

To the best of our knowledge, the estimator $\widehat{P}_n^{\rm ad}$ defined by (31) is new. However, it is related to existing estimators from the literature. First, it bears some relation with absolute discounting [NEK94, KN95], a core component in the Kneser-Ney smoothing method which has long been favored in natural language processing [KN95, CG99] and whose performance has been analyzed in [OD12, FOOP17]; see also [Teh06] for a justification of this method in the case of polynomially decaying class frequencies. Specifically, absolute discounting consists in removing a constant $\eta \in (0,1)$ to the count $N_j \ge 1$ of all classes $j = 1, \ldots, d$ that appear in the sample, and sharing the freed mass among missing classes. When $D_n \leq d/2$, absolute discounting and the estimator $\widehat{P}_n^{\text{ad}}$ both assign a probability $\widehat{p}_j \simeq D_n/(dn)$ to missing classes j; on the other hand, the corrections of these two estimators for classes j with $N_j \gg 1$ differ. Second, in an online setting where observations accrue one by one, Hutter [Hut13] proposed a sequential prediction method using a regularization $\lambda_t \simeq D_t / \{d \log[et/D_t]\}$ after t observations, and shows that it satisfies regret guarantees in sparse situations. This method is related to the estimator $\hat{P}_n^{\rm ad}$, the main difference being a logarithmic factor in the regularization parameter. We also note that guarantees in the sequential setting necessarily feature an additional $\log n$ factor compared to the fixed-sample size setting, due to the contribution of small sample sizes.

Since $s_{n/2}^{\circ} \leq s_{n/2} \leq s_n$, Proposition 2 matches the upper bound (29) from [FOOP17] for the absolute discounting estimator. In addition, it improves this bound when $s_{n/2}^{\circ} \ll s_n$.

Remark 1 (Comparison between s_n and s_n°). In order to quantitatively appreciate similarities and differences between the two effective sample sizes, it helps to consider the following stylized situations. Let $P = (p_1, \ldots, p_d)$, and assume (up to relabeling classes) that $p_1 \ge \ldots \ge p_d$.

- (a) Polynomial decay: first, assume that $p_j \simeq j^{-\alpha}$ for some constant $\alpha > 1$, and that $d \gg n^{1/\alpha}$. In this case, one has $s_n \simeq s_n^{\circ} \simeq n^{1/\alpha}$ up to constants that may depend on α . Hence, the two effective sample size parameters are equivalent, and so are the bounds (29) and (32).
- (b) Geometric decay: assume now that $c_1 e^{-C_1 j} \leq p_j \leq C_2 e^{-c_2 j}$ for every $j = 1, \ldots, d$ for some constants $c_1, c_2, C_1, C_2 > 0$. If $d \geq C \log n$ for some large enough constant C (depending on the previous constants c_k, C_k), then $s_n \asymp \log n$ while $s_n^{\circ} \asymp 1$. In particular, if d scales polynomially in n, then the bound (29) scales as $\log^2(n)/n$, while the bound (32) scales as $\log(n)/n$.
- (c) Sparse distributions: finally, assume that P is supported on $s \leq d$ classes, and that the frequencies of these classes are lower-bounded; namely, $p_j \geq c/s$ for $j = 1, \ldots, s$ for some $c \in (0, 1)$, while $p_j = 0$ for $s < j \leq d$. Then, for $n \geq 2s \log(es)/c$, one has $s_n = s$ while $s_n^{\circ} \leq s_{n/2}^{\circ} \leq ese^{-cn/(2s)} \leq 1$, hence $s_n^{\circ} \ll s_n$ if $s \gg 1$. Thus if $d^{\varepsilon} \leq s \leq d^{1-\varepsilon}$ for some $\varepsilon \in (0, 1)$, ignoring the dependence on c, ε , the bound (29) scales as $s \log(d)/n$, while the bound (32) scales as s/n, removing a log d factor.

We note that the bound of Proposition 2 holds in expectation, while the focus of the present work is on high-probability guarantees. High-probability extensions will be proved in what follows, but we have nonetheless included Proposition 2 for two reasons. First, this bound features small numerical constants. Second and more importantly, the proof of Proposition 2 is selfcontained and significantly simpler than that of the high-probability bounds. Specifically, this proof relies on a combination of exchangeability and leave-one-out arguments. Such arguments already appear in the proof of the bound (5) for the Laplace estimator (see [Cat97, MG22]), although the proof of Proposition 2 requires somewhat more careful counting and conditioning.

4.2.3 High-probability upper bounds for sparse distributions

We now turn to the strongest positive guarantees in this work, namely high-probability upper bounds for the estimator \hat{P}_n^{ad} and its confidence-dependent version $\hat{P}_{n,\delta}^{\text{ad}}$, defined below. These results can be seen as sparsity-adaptive extensions of Theorems 1 and 3, respectively.

Theorem 5. Let $n \ge 12, d \ge 3$, $\delta \in (e^{-n/6}, e^{-2})$ and $P \in \mathcal{P}_d$, and denote $s_n = s_n(P)$ and $s_n^\circ = s_n^\circ(P)$. The add- $\widehat{\lambda}$ estimator $\widehat{P}_n^{\mathrm{ad}} = (\widehat{p}_1, \ldots, \widehat{p}_d)$ given by

$$\widehat{p}_j = \frac{N_j + \widehat{\lambda}}{n + \widehat{\lambda} d} \tag{33}$$

with $\widehat{\lambda} = D_n/d$, where $\widehat{D}_n = \sum_{j=1}^d \mathbf{1}(N_j \ge 1)$ is the number of distinct classes that appear in the sample X_1, \ldots, X_n , satisfies the following guarantee: with probability at least $1 - 14\delta$ under P,

$$\mathrm{KL}(P, \widehat{P}_{n}^{\mathrm{ad}}) \leq 121000 \, \frac{s_{n} + s_{n/112}^{\circ} \log(ed/s_{n}) + \max\{\log d, \log\log(1/\delta)\}\log(1/\delta)}{n} \,. \tag{34}$$

In addition, letting $\widehat{\lambda}_{\delta} = \max\{D_n, \log(1/\delta)\}/d$, the add- $\widehat{\lambda}_{\delta}$ estimator $\widehat{P}_{n,\delta}^{ad}$ satisfies the following guarantee: with probability at least $1 - 14\delta$ under P,

$$\mathrm{KL}(P, \widehat{P}_{n,\delta}^{\mathrm{ad}}) \leqslant 121000 \, \frac{s_n + s_{n/112}^\circ \log(ed/s_n) + \log(d) \log(1/\delta)}{n} \,. \tag{35}$$

The proof of Theorem 5 is provided in Section 6.6. This proof relies on a combination of lemmata already used in the analysis of the Laplace estimator, together with additional results on the "underestimated mass" described in Section 5 below. We now comment on Theorem 5.

First, since $s_{n/112}^{\circ} \leq s_{n/112} \leq s_n \leq d$, using that $s \mapsto s \log(ed/s)$ is increasing on [0, d] and recalling inequality (21), Theorem 5 shows that the confidence-independent estimator \hat{P}_n^{ad} achieves the best possible uniform high-probability guarantee over \mathcal{P}_d among confidence-independent estimators, while the confidence-dependent estimator $\hat{P}_{n,\delta}^{\mathrm{ad}}$ achieves the minimax high-probability guarantee over \mathcal{P}_d . These results match the lower bounds of previous sections and the guarantees for the Laplace estimator and its confident-dependent modification.

Second, for any $s \in \{2, \ldots, d\}$ and $P \in \mathcal{P}_{s,d}$ one has $s_{n/112}^{\circ} \leq s_n \leq s$; hence, Theorem 5 implies that the estimator $\widehat{P}_{n,\delta}^{\mathrm{ad}}$ achieves an upper bound of order $\{s \log(ed/s) + \log(d) \log(1/\delta)\}/n$ over $\mathcal{P}_{s,d}$. This matches the minimax lower bound over $\mathcal{P}_{s,d}$ of Corollary 1, and thus completes the characterization of the minimax high-probability rate over this class. In addition, the estimator $\widehat{P}_{n,\delta}^{\mathrm{ad}}$ achieves this rate simultaneously for all sparsity levels $s \in \{2, \ldots, d\}$.

Finally, the guarantees of Theorem 5 mainly dependent on the intrinsic and distributiondependent parameters $s = s_n(P)$ and $s^\circ = s_{n/112}^\circ(P)$, with only a (necessary) logarithmic dependence on the total number d of classes. They may therefore be viewed as essentially "non-parametric". Note that since $s^\circ \leq s$, the term $s^\circ \log(ed/s)$ in (34) and (35) may be bounded by $s^\circ \log(ed/s^\circ)$; this matches the corresponding term in the in-expectation bound of Proposition 2, up to constant factors in the error bound and sample size. In fact, the complexity term in Theorem 5 may appear to be of a smaller order of magnitude than that of Proposition 2 when $s^\circ \ll s$. This is not the case, since one also has $s + s^\circ \log(ed/s^\circ) \lesssim s + s^\circ \log(ed/s)$ (indeed, either $s^\circ \log(ed/s^\circ) \leq s$ and the bound holds, or otherwise $\frac{ed/s^\circ}{\log(ed/s^\circ)} < ed/s$ and thus $\log(ed/s^\circ) \lesssim \log(ed/s)$ hence the bound also holds). Consequences for the sample complexity. While we have stated our results in terms of error rates, one may also formulate them in terms of the *sample complexity*, which is the sample size n required to achieve an error in relative entropy of ε with probability at least $1 - \delta$.

For instance, the results of Section 3 (Theorems 3 and 4) imply that for $d \ge 5000$, $0 < \delta < e^{-2}$ and $0 < \varepsilon < 1$, a sample size of

$$n\gtrsim N_{\max}(d,\varepsilon,\delta)=\frac{d+\log(d)\log(1/\delta)}{\varepsilon}$$

is both necessary and sufficient for the existence of an estimator $\widehat{P}_n = \Phi(X_1, \ldots, X_n)$ such that $\mathbb{P}_P(\mathrm{KL}(P, \widehat{P}_n) \leq \varepsilon) \geq 1 - \delta$ for every $P \in \mathcal{P}_d$.

It is also instructive to formulate the distribution-dependent upper bound (35) for the adaptive and confidence-dependent estimator $\widehat{P}_{n,\delta}^{\mathrm{ad}}$ in this way. For any dimension $d \ge 2$, accuracy $\varepsilon \in (0,1)$, failure probability $\delta \in (0, e^{-2})$ and distribution $P \in \mathcal{P}_d$, define the following "critical sample sizes":

$$N_{\text{obs}}(P,\varepsilon) = \inf\left\{n \ge 1 : \frac{s_n(P)}{n} \le \varepsilon\right\},\tag{36}$$

$$N_{\mathsf{miss}}(P, d, \varepsilon) = \inf\left\{n \ge 1 : \frac{s_n^{\circ}(P)}{n} \leqslant \frac{\varepsilon}{\log(ed/\varepsilon n)}\right\},\tag{37}$$

$$N_{\mathsf{dev}}(d,\varepsilon,\delta) = \frac{\log(d)\log(1/\delta)}{\varepsilon}.$$
(38)

Note that for $n \ge N_{obs}(P,\varepsilon)$ (resp. $n \ge N_{miss}(P,d,\varepsilon)$), one has $s_n(P)/n \le \varepsilon$ (resp. $s_n^{\circ}(P)/n \le \varepsilon/\log(ed/\varepsilon n)$) since $s_n(P)/n$ (resp. $s_n^{\circ}(P)/n$ and $\log(ed/\varepsilon n)$) is non-increasing in n, as noted above. The sample sizes (36) and (38) serve to control the first and third term in the error bound (35). In addition, up to constant factors, the second sample size allows one to control the second term in the bound (35). Indeed, bounding $s_n \ge s_{n/112}^{\circ}$, the second term is at most of order $s_{n/112}^{\circ}\log(ed/s_{n/112}^{\circ})/n$. In addition, up to replacing n by 112n, for this term to be bounded by $C\varepsilon$ for some absolute constant C > 1, it suffices that (recalling that $s_n^{\circ} \le d$)

$$\frac{s_n^\circ \log(ed/s_n^\circ)}{n} \lesssim \varepsilon \quad \text{i.e.} \quad \frac{d/s_n^\circ}{\log(ed/s_n^\circ)} \gtrsim \frac{d}{\varepsilon n} \quad \text{i.e.} \quad \frac{d}{s_n^\circ} \gtrsim \frac{d}{\varepsilon n} \log\left(\frac{ed}{\varepsilon n}\right),$$

which amounts to the condition in (37).

The upper bound (35) from Theorem 5, together with the previous discussion, implies the following: for some absolute constants $c_1, c_2 \ge 1$ (one may take $c_1 = 112$), if

$$n \ge c_1 \max\left\{ N_{\mathsf{obs}}(P,\varepsilon), N_{\mathsf{miss}}(P,d,\varepsilon), N_{\mathsf{dev}}(d,\varepsilon,\delta) \right\},\tag{39}$$

then the estimator $\widehat{P}_{n,\delta}^{\mathrm{ad}}$ satisfies

$$\mathbb{P}_P\big(\mathrm{KL}(P,\widehat{P}_{n,\delta}^{\mathrm{ad}}) \geqslant c_2 \,\varepsilon\big) \leqslant \delta \,. \tag{40}$$

5 High-probability bound on the missing mass

In this section, we present a result that plays an important role in the proof of Theorem 5, namely in the high-probability analysis of adaptive estimators, but which may also be of independent interest.

Specifically, the following quantities appear naturally in our analysis and in other contexts:

Definition 3 (Missing and underestimated masses). Given a distribution P and an i.i.d. sample X_1, \ldots, X_n from P, we define the missing mass $M_n = \sum_{j=1}^d p_j \mathbf{1}(N_j = 0)$ as the total mass under P of classes that do not appear in the sample.

We also define the *underestimated mass* $U_n = \sum_{j=1}^d p_j \mathbf{1}(N_j \leq np_j/4)$ as the total mass of classes whose empirical frequency underestimates their true probability by at least a factor of 4.

The missing and underestimated masses depend on both the sample X_1, \ldots, X_n and on the true distribution; as such, they are not "observable" from the data.

It is clear from the definition that $M_n \leq U_n$. As it happens, the quantity that plays a role in our analysis is the underestimated mass U_n , and thus our main goal is to provide a highprobability upper bound on U_n . On the other hand, the missing mass M_n is a classical quantity, which has been studied in Statistics since the work of Good [Goo53], and our bound on U_n will also yield a new high-probability upper bound on M_n .

Specifically, our aim is to address the following question:

For a given $\varepsilon > 0$ and $\delta \in (0, e^{-1})$ and a distribution $P \in \mathcal{P}_d$, how large must the sample size n be to ensure that the underestimated mass U_n (and thus the missing mass M_n) is smaller than ε with probability at least $1 - \delta$?

Before stating our main result, we first clarify that the behavior of the expected missing mass is essentially governed by the behavior of the parameter $s_n^{\circ}(P)$ (Definition 2), a fact we alluded to in Section 4.

Lemma 2. For $P \in \mathcal{P}_d$, let $s_n^{\bullet}(P) = \sum_{j=1}^d (np_j)e^{-np_j}$ for $n \in (0, +\infty)$, and recall the definitions of $s_n^{\circ}(P)$ and M_n (Definitions 2 and 3). For any integer $n \ge 1$, the following holds:

1. $\mathbb{E}_{P}[M_{n}] = \sum_{j=1}^{d} p_{j}(1-p_{j})^{n};$ 2. $s_{2n}^{\bullet}(P)/(2n) - e^{-0.3n} \leq \mathbb{E}_{P}[M_{n}] \leq s_{n}^{\bullet}(P)/n;$ 3. $e^{-1}s_{n}^{\circ}(P) \leq s_{n}^{\bullet}(P) \leq 2s_{n/2}^{\circ}(P).$

4. For
$$n \ge 3$$
, letting $s_n^{\diamond}(P) = \mathbb{E}_P[D_{2n} - D_n]$, one has $s_{2n}^{\diamond}(P)/12 - e^{-n} \le s_n^{\diamond}(P) \le s_n^{\diamond}(P)$.

In short, the quantities $n \mathbb{E}_P[M_n], s_n^{\circ}(P), s_n^{\circ}(P)$ are essentially equivalent, up to constant factors in their values and in the sample size n, and possibly additive exponentially small terms.

Proof of Lemma 2. For the first identity, write $\mathbb{E}_P[M_n] = \sum_{j=1}^d p_j \mathbb{P}_P(N_j = 0) = \sum_{j=1}^d p_j (1 - p_j)^n$. The upper bound $\mathbb{E}_P[M_n] \leq s_n^{\bullet}(P)/n$ comes from the fact that $(1 - p_j)^n \leq e^{-np_j}$. For the lower bound, given $\lambda \in \mathbb{R}^+$ let $N(\lambda) \sim \mathsf{Poisson}(\lambda)$, and set $N_j^{(\lambda)} = \sum_{1 \leq i \leq N(\lambda)} \mathbf{1}(X_i = 0)$.

For the lower bound, given $\lambda \in \mathbb{R}^{-1}$ let $N(\lambda)$'s Poisson(λ), and set $N_j = \sum_{1 \leq i \leq N(\lambda)} \mathbf{I}(X_i = j)$ and $M^{(\lambda)} = M_{N(\lambda)} = \sum_{j=1}^d p_j \mathbf{1}(N_j^{(\lambda)} = 0)$. It is a classical fact (e.g., this follows from [MU17, Theorem 5.6 p. 100] and the convolution property of Poisson distribution) that $N_j^{(\lambda)} \sim \mathsf{Poisson}(\lambda p_j)$. Thus $\mathbb{E}[M^{(\lambda)}] = \sum_{j=1}^d p_j \mathbb{P}(N_j^{(\lambda)} = 0) = \sum_{j=1}^d p_j e^{-\lambda p_j} = s_{\lambda}^{\bullet}(P)/\lambda$. Clearly, if $N(2n) \geq n$ then $M^{(2n)} \leq M_n$. On the other hand, if N(2n) < n we may write $M_n \geq 0 \geq M^{(2n)} - 1$. Hence,

$$n \mathbb{E}_P[M_n] \ge n \mathbb{E}_P[M^{(2n)} - \mathbf{1}(N(2n) < n)] = n \times \frac{s_{2n}^{\bullet}(P)}{2n} - n\mathbb{P}(N(2n) < n),$$

and the desired claim follows from the bound $\mathbb{P}(N(2n) < n) \leq \exp(-D(n, 2n)) \leq e^{-(1-\log 2)n} \leq e^{-0.3n}$ by Lemma 17.

For the third point, we apply the bounds $(np_j)e^{-np_j} \leq np_j$ and $(np_j)e^{-np_j} = 2(np_j/2)e^{-np_j} \leq 2e^{-1+np_j/2}e^{-np_j} \leq e^{1-np_j/2}$ to get the upper bound. To get the lower bound, we use that $np_j \geq 1$ when $p_j \geq 1/n$, and $e^{-np_j} \geq e^{-1}$ when $p_j < 1/n$.

Finally, for the fourth point, write

$$s_n^{\diamond}(P) = \mathbb{E}_P[D_{2n} - D_n] = \sum_{j=1}^d \left\{ (1 - p_j)^n - (1 - p_j)^{2n} \right\} = \sum_{j=1}^d (1 - p_j)^n \left\{ 1 - (1 - p_j)^n \right\}.$$

The upper bound $s_n^{\diamond}(P) \leq s_n^{\diamond}(P)$ follows from the bounds $(1-p_j)^n \leq e^{-np_j}$ and $1-(1-p_j)^n \leq \min(1, np_j)$. For the lower bound, consider first indices j for which $np_j \leq 1$. In this case, one has $1-(1-p_j)^n \geq (1-e^{-1})np_j$ (see the proof of Fact 1) while $(1-p_j)^n \geq (1-1/n)^n \geq (2/3)^3$, thus $(1-p_j)^n \{1-(1-p_j)^n\} \geq (2/3)^3(1-e^{-1})np_j \geq np_j/6 \geq \min(2np_j, e^{1-2np_j})/12$. When $np_j > 1$, one has $1-(1-p_j)^n > 1-(1-1/n)^n \geq 1-e^{-1}$, while $(1-p_j)^n \geq e^{-2np_j} - e^{-n}\mathbf{1}(p_j > 3/4)$ as $1-p > e^{-2p}$ for $p \in [0,3/4]$; thus $(1-p_j)^n \{1-(1-p_j)^n\} \geq (1-e^{-1})e^{-2np_j} - e^{-n}\mathbf{1}(p_j > 3/4) \geq \min(2np_j, e^{1-2np_j})/12 - e^{-n}\mathbf{1}(p_j > 3/4)$. Combining the previous inequalities and using that there is at most one index j such that $p_j > 3/4$ concludes the proof. \Box

In particular, the inequalities of Lemma 2 imply that:

$$\frac{s_{2n}^{\circ}(P)}{2en} - e^{-0.3n} \leqslant \mathbb{E}_P[M_n] \leqslant \frac{2s_{n/2}^{\circ}(P)}{n}.$$
(41)

(In addition, the $e^{-0.3n}$ term can often be ignored: for instance, it is dominated by the first term for large *n* whenever there is a class *j* with $p_j \leq 0.1$, which occurs whenever at least 10 classes have nonzero probability.) Hence, ignoring the $e^{-0.3n}$ term, the behavior of the typical value $\mathbb{E}_P[M_n]$ of the missing mass is essentially equivalent to that of $s_n^{\circ}(P)/n$.

Theorem 6 below provides a deviation upper bound on the underestimated and missing masses, involving the same complexity parameter as the in-expectation estimates (41):

Theorem 6. For any $n, d \ge 2$, $\delta \in (0, e^{-1})$ and any $P \in \mathcal{P}_d$, with probability at least $1 - 8\delta$ under P one has

$$M_n \leqslant U_n \leqslant \frac{336 \, s_{n/112}^\circ(P) + 2500e \log(1/\delta)}{n} \,.$$
(42)

The proof of Theorem 6 is provided in Section 8. Roughly speaking, the proof relies on a careful control of the contribution to U_n of classes p_j of each order of magnitude, followed by a combination of these per-scale contributions.

As we argue now, Theorem 6 constitutes an almost optimal high-probability upper bound on the missing mass, up to constant factors in the mass and sample size. Indeed, combining (41) and Theorem 6 shows that for every $\delta \in (0, e^{-1})$, with probability at least $1 - \delta$,

$$M_n \leqslant U_n \leqslant 3e \mathbb{E}_P[M_{n/224}] + 2600e \frac{\log(1/\delta)}{n}.$$

$$\tag{43}$$

(Specifically, from (41) we obtain an additive term in $3e^{1-0.15n} \leq 20e \log(1/\delta)/n$.) In addition, as will be discussed below, the deviation term in $\log(1/\delta)/n$ is necessary, except in situations where the probabilities in P exhibit a significant gap; and in this case, the bound of Theorem 6 admits a simple strengthening that addresses its sub-optimality.

Critical sample size. Thanks to Theorem 6, we can address the question of the critical sample size for which the missing mass is small with high probability. Specifically, for $P \in \mathcal{P}_d$ and $\varepsilon \in (0, 1)$, define

$$N^{\circ}(P,\varepsilon) = \inf\left\{n \ge 1 : \frac{s_n^{\circ}(P)}{n} \le \varepsilon\right\}.$$
(44)

Note that $s_n^{\circ}(P)/n \leq \varepsilon$ for $n \geq N^{\circ}(P,\varepsilon)$ as $s_n^{\circ}(P)/n$ is non-increasing in n. Then, it follows from Theorem 6 that for any $\delta \in (0, e^{-1})$, if

$$n \ge 2500 \max\left\{N^{\circ}(P,\varepsilon), \frac{\log(1/\delta)}{\varepsilon}\right\},$$
(45)

then $\mathbb{P}_P(M_n \ge 6\varepsilon) \le \delta$. Below, we show that the estimate (45) is optimal in many situations, and that a simple tightening is optimal in the general case. We finally compare our condition with those derived from previous upper bounds on the missing mass in the literature.

Necessity of the complexity parameter. We first argue that the condition $n \gtrsim N^{\circ}(P, \varepsilon)$ is almost necessary for the missing mass to be small in expectation. Indeed, the latter condition amounts to

$$n \ge N_{\exp}(P,\varepsilon) = \inf\left\{n \ge 1 : \mathbb{E}_P[M_n] \le \varepsilon\right\}.$$
(46)

But by inequality (41), one has for some absolute constant c > 1:

$$N^{\circ}(P, c \varepsilon) \leqslant c N_{\exp}(P, \varepsilon) + c \log\left(\frac{1}{\varepsilon}\right),$$
(47)

hence whenever $\log(1/\varepsilon) \ll N^{\circ}(P,\varepsilon)$ (which as discussed above occurs in most situations of interest) the condition $n \gtrsim N^{\circ}(P,\varepsilon)$ is necessary. Besides, whenever the second term $\log(1/\delta)/\varepsilon \ge \log(1/\varepsilon)$ in (45) is necessary, so is $\log(1/\varepsilon)$ and thus $N^{\circ}(P,\varepsilon)$.

In fact, it follows from inequality (43) that $\mathbb{P}_P(U_n \ge 4e\varepsilon) \le \delta$ whenever

$$n \ge 2600 \max\left\{ N_{\exp}(P,\varepsilon), \frac{\log(1/\delta)}{\varepsilon} \right\}.$$
 (48)

Deviation term and refinement. We now address the question of whether the condition $n \geq \log(1/\delta)/\varepsilon$ is necessary. As it turns out, this is often but not always the case; however, a simple strengthening of this result does provide a necessary condition. In order to state this refinement, define the set of all sums of class probabilities in P:

$$\mathcal{S}(P) = \left\{ p_J = \sum_{j \in J} p_j : J \subset [d] \right\} \subset [0, 1].$$

$$\tag{49}$$

It is clear that $U_n, M_n \in \mathcal{S}(P)$. Hence, in order to ensure that $M_n < \varepsilon$, it suffices that $M_n < \overline{\varepsilon}$, where we let

$$\overline{\varepsilon} = \overline{\varepsilon}(P,\varepsilon) = \inf \left(\mathcal{S}(P) \cap [\varepsilon, 1] \right).$$
(50)

(Equivalently, if $M_n \leq t$ then $M_n \leq \underline{b}(P, t) = \sup (\mathcal{S}(P) \cap [0, t])$, which strengthens the bound of Theorem 6.) Hence, condition (45) ensures that $\mathbb{P}_P(M_n \geq 6\varepsilon) \leq \delta$ whenever

$$n \ge 15000 \max\left\{ N^{\circ}(P,\varepsilon), \frac{\log(1/\delta)}{\overline{\varepsilon}(P,\varepsilon)} \right\}.$$
(51)

Indeed, if $\overline{\varepsilon} \leq 6\varepsilon$, apply (45) and bound $1/\varepsilon \leq 6/\overline{\varepsilon}$. On the other hand, if $\overline{\varepsilon} > 6\varepsilon$, apply (45) to $\overline{\varepsilon}/6 > \varepsilon$, so that with probability $1 - \delta$ one has $M_n < 6 \times \overline{\varepsilon}/6 = \overline{\varepsilon}$, hence $M_n < \varepsilon$, and use that $N^{\circ}(P, \overline{\varepsilon}/6) \leq N^{\circ}(P, \varepsilon)$.

We now show that the condition $n \gtrsim \log(1/\delta)/\overline{\varepsilon}$ is necessary to ensure this high-probability bound, whenever $\overline{\varepsilon}$ is bounded away from 1. For concreteness, assume that $\overline{\varepsilon} \leq 3/4$, so that $1-\overline{\varepsilon} > e^{-2\overline{\varepsilon}}$. Let $J \subset [d]$ such that $p_J = \sum_{j \in J} p_j = \overline{\varepsilon}$, and let $N_J = \sum_{j \in J} N_j = \sum_{i=1}^n \mathbf{1}(X_i \in J)$. Clearly, if $N_J = 0$ then $M_n \ge p_J = \overline{\varepsilon} \ge \varepsilon$. Hence, $\mathbb{P}_P(M_n \ge \varepsilon) \ge \mathbb{P}_P(N_J = 0) = (1 - p_J)^n > e^{-2\overline{\varepsilon}n} \ge \delta$ whenever $n \le \log(1/\delta)/(2\overline{\varepsilon})$.

Summarizing the previous discussion, we obtain up to universal constants the critical sample size after which the missing mass is small, in expectation and with high probability (the sufficient condition being deduced from (48) in the same way that (51) is deduced from (45)).

Corollary 2. Let $\varepsilon \in (0, 3/4)$, $\delta \in (0, e^{-1})$, $n \ge 1$ and $P \in \mathcal{P}_d$.

1. When $\overline{\varepsilon}(P,\varepsilon) \leq 3/4$, if $\mathbb{E}_P[M_n] \leq \varepsilon$ and $\mathbb{P}_P(M_n \geq \varepsilon) \leq \delta$, then

$$n \geqslant \frac{1}{2} \max \left\{ N_{\exp}(\varepsilon), \frac{\log(1/\delta)}{\overline{\varepsilon}(P,\varepsilon)} \right\}$$

2. Conversely, if

$$n \ge 11000e \max\left\{ N_{\exp}(P,\varepsilon), \frac{\log(1/\delta)}{\overline{\varepsilon}(P,\varepsilon)} \right\},\tag{52}$$

then $\mathbb{E}_P[M_n] \leq \varepsilon$ and $\mathbb{P}_P(M_n \geq 4e\varepsilon) \leq \delta$.

Since the optimal deviation term scales as $\log(1/\delta)/\overline{\varepsilon}$, the deviation term $\log(1/\delta)/\varepsilon$ is optimal unless $\overline{\varepsilon} \gg \varepsilon$. Distributions for which this occurs can be characterized as follows:

Fact 2. Let $\varepsilon \in (0, 1/2)$ and $\overline{\varepsilon} \ge 2\varepsilon$. Let $P = (p_1, \ldots, p_d) \in \mathcal{P}_d$, with $p_1 \ge \ldots \ge p_d$ without loss of generality. The following properties are equivalent:

- (i) $\overline{\varepsilon}(P,\varepsilon) \ge \overline{\varepsilon};$
- (ii) there exists $j^* \in \{1, \ldots, d\}$ such that $p_{j^*} \ge \overline{\varepsilon}$ while $\sum_{j^* < j \le d} p_j < \varepsilon$.

In this case, one has $\overline{\varepsilon}(P,\varepsilon) = p_{j^*}$.

Proof. We first prove the implication $(i) \Rightarrow (ii)$. Let $j^* = \max\{1 \le j \le d : p_j \ge \varepsilon\} \in \{0, \ldots, d\}$, with the convention that $j^* = 0$ if the set is empty.

We first show that $\sum_{j^* < j \leq d} p_j < \varepsilon$. If $j^* = d$, the sum equals 0 and the property holds; we thus assume that $j^* \leq d-1$. In this case, let j' denote the largest integer in $\{j^*+1,\ldots,d\}$ such that $\sum_{j^* < j \leq j'} p_j < \varepsilon$ (by definition of j^* , this inequality holds for $j' = j^*+1$, hence $j' \geq j^*+1$). We need to show that j' = d. We proceed by contradiction and assume that $j' \leq d-1$. In this case, one has $\sum_{j^* < j \leq j'+1} p_j < \varepsilon$ while $\sum_{j^* < j \leq j'+1} p_j \geq \varepsilon$. But since $\sum_{j^* < j \leq j'+1} p_j \in \mathcal{S}(P)$, this also implies that $\sum_{j^* < j \leq j'+1} p_j \geq \overline{\varepsilon} \geq 2\varepsilon$, hence $p_{j'+1} = \sum_{j^* < j \leq j'+1} p_j - \sum_{j^* < j \leq j'} p_j > \varepsilon$. But this contradicts the fact that $p_{j'+1} \leq p_{j^*+1} < \varepsilon$, proving the claim of this paragraph.

Now since $\sum_{0 < j^* \leq d} p_j = 1$, the inequality $\sum_{j^* < j \leq d} p_j < \varepsilon < 1$ implies that $j^* \geq 1$. By definition of j^* , one has $p_{j^*} \geq \varepsilon$; but since $p_{j^*} \in \mathcal{S}(P)$, this implies that $p_{j^*} \geq \overline{\varepsilon}(P, \varepsilon) \geq \overline{\varepsilon}$.

We now conclude with the implication $(ii) \Rightarrow (i)$. Let $J \subset [d]$. If $J \subset \{j^* + 1, \ldots, d\}$, then $\sum_{j \in J} p_j \leq \sum_{j^* < j \leq d} p_j < \varepsilon$. Otherwise, J contains an element $j' \in \{1, \ldots, j^*\}$, hence $\sum_{j \in J} p_j \geq p_{j'} \geq p_{j^*} \geq \overline{\varepsilon}$. Since the inequality $\sum_{j \in J} p_j \geq p_{j^*}$ is an equality for $J = \{j^*\}$, one has $\overline{\varepsilon}(P, \varepsilon) = p_{j^*} \geq \overline{\varepsilon}$.

It follows from Fact 2 that if $\overline{\varepsilon}(P,\varepsilon) \gg \varepsilon$, then there exists $j^* \in \{1,\ldots,d\}$ such that $p_{j^*} \gg \varepsilon$ while $\sum_{j^* < j \leq d} p_j < \varepsilon$. Then, either $j^* = d$ and $p_d \gg \varepsilon$, or $j^* \leq d-1$ and $\sum_{j=j^*+1}^d p_j < \varepsilon \ll p_{j^*}$, so in particular $p_{j^*+1} \ll p_{j^*}$.

We note that the latter property occurs neither for polynomial decay of probabilities (as defined in Remark 1) nor for exponential decay, for which $p_{j+1} \gtrsim p_j$ for every j < d. However, this property does occur in the case of sparse distributions (third example of Remark 1), for which the refinement in terms of $\overline{\varepsilon}(P, \varepsilon)$ brings an improvement.

Previous results. The behavior of the missing mass and the question of its estimation have been studied in a rich literature, which can be traced to the work of Good [Goo53] introducing the Good-Turing estimate. In what follows, we focus our discussion on existing high-probability bounds on the missing mass, and on conditions they imply on the sample size for the missing mass to be small. We additionally refer to the discussion in [BHBO17] (and references therein) for background and references on the distinct question of estimation of the missing mass.

First, a deviation bound due to McAllester and Schapire [MS00], with constants later improved by McAllester and Ortiz [MO03] (see also Berend and Kontorovitch [BK13] for an alternative approach), shows that the missing mass exhibits sub-Gaussian tails. Specifically, the following distribution-free deviation bound holds [MO03, Theorem 16]: for any $d \ge 2$, distribution $P \in \mathcal{P}_d$ and $\delta \in (0, 1)$, one has

$$\mathbb{P}_P\left(M_n \ge \mathbb{E}_P[M_n] + \sqrt{\frac{\log(1/\delta)}{n}}\right) \le \delta.$$
(53)

This implies that for any $\varepsilon \in (0, 1)$ and $\delta \in (0, e^{-1})$, if

$$n \ge \max\left\{N_{\exp}(P,\varepsilon), \frac{\log(1/\delta)}{\varepsilon^2}\right\},$$
(54)

then $\mathbb{P}_P(M_n \ge 2\varepsilon) \le \delta$. While significant and nontrivial, condition (54) exhibits a suboptimal dependence on ε compared to (48), although it involves smaller numerical constants.

To the best of our knowledge, the previous best deviation bound on the missing mass is due to Ben-Hamou, Boucheron and Ohannessian [BHBO17], tightening previous multiplicative concentration bounds by Ohannessian and Dahleh [OD12]. Specifically, the following tail bound follows from [BHBO17, Theorem 3.9]:

$$\mathbb{P}_P\left(M_n \ge \mathbb{E}_P[M_n] + \frac{\sqrt{2d_n^+ \log(1/\delta)}}{n} + \frac{\log(1/\delta)}{n}\right) \le \delta,$$
(55)

where
$$d_n^+ = d_n^+(P) = \sum_{j=1}^d \left\{ 1 - (np_j + 1)e^{-np_j} \right\} \approx \sum_{j=1}^d \min\left[1, (np_j)^2 \right].$$
 (56)

(Note that $d_n^+(P) \simeq s_n(P)$ whenever $\sum_{j: p_j < 1/n} (np_j) \lesssim \sum_{j: p_j \ge 1/n} 1$.) Since $d_n^+ \leqslant \sum_{j=1}^d (1 - e^{-np_j}) \leqslant \sum_{j=1}^d (np_j) = n$, this bound recovers (53) up to constants in the regime of interest $\delta \in (e^{-n}, 1)$. Furthermore, as shown in [BHBO17], inequality (55) strictly improves (53) in most situations, and in fact provides sharp results in several interesting cases.

However, there are situations in which the second term in (55) dominates and leads to a suboptimal bound. For instance, consider the case where $p_j \approx 1/d$ for $1 \leq j \leq d/2$, while $p_j \approx 2^{-(j-d/2)}/d$ for $d/2 < j \leq d$, and the regime of constant probability $\delta \in (0,1)$. For $d \log d \ll n \ll 2^d$, the typical missing mass is at most of order $\mathbb{E}_P[M_n] \approx 1/n$. However, in this regime one has $d_n^+ \approx d$, hence the bound (55) is of order \sqrt{d}/n , which is suboptimal by a \sqrt{d} factor. This \sqrt{d} factor can be removed by resorting to Theorem 6.

6 Proof of high-probability upper bounds

In this section, we provide the proof of high-probability upper bounds for estimation in this paper, namely Theorems 1, 3 and 5 (the proof of the last result relies in part on Theorem 6 on the underestimated mass, which is proved in Section 8).

Specifically, we start with lemmata that are used in the proof of all high-probability upper bounds, before concluding with the proof of each specific result.

6.1 Risk decomposition

All estimators we consider are "add- λ " smoothing rules, where λ may be confidence-dependent and/or data-dependent. We therefore start our analysis with the following deterministic risk decomposition for add- λ estimators. Below, we let $\overline{p}_j = N_j/n$ be empirical frequency of class $j = 1, \ldots, d$.

Lemma 3. Consider the distribution $\widehat{P}_n = (\widehat{p}_1, \dots, \widehat{p}_d)$ given by

$$\widehat{p}_j = \frac{N_j + \lambda}{n + \lambda d}, \quad j = 1, \dots, d, \qquad (57)$$

for some $\lambda \in (0, n/d]$ that may depend on X_1, \ldots, X_n . Then, we have

$$\operatorname{KL}(P,\widehat{P}_n) \leqslant 6 \sum_{j=1}^d \left(\sqrt{\overline{p}_j} - \sqrt{p_j}\right)^2 + \frac{7\lambda d}{n} + \sum_{j: p_j \ge 4\lambda/n} p_j \log\left(\frac{2np_j}{\lambda}\right) \mathbf{1}\left(N_j \leqslant \frac{np_j}{4}\right).$$
(58)

The decomposition (58) features three terms.

The first term, which does not depend on the estimator (that is, on λ), corresponds to the squared Hellinger distance between the empirical distribution $\overline{P}_n = (\overline{p}_1, \ldots, \overline{p}_d)$ and the true distribution P. It constitutes a natural "hard limit" for our estimation guarantees in relative entropy. This term is controlled in Section 6.2.

The second term accounts for the "bias" due to the use of regularization, which diminishes the probability assigned to high-frequency classes. This term increases with the smoothing parameter λ . Its control is immediate when λ is data-independent, and relatively straightforward when λ is data-dependent.

The third and final term accounts for the contribution of classes whose frequency is significantly underestimated, which inflates the relative entropy over ideal asymptotic rates. The effect of underestimation of true frequencies is mitigated by the use of smoothing, and indeed this term decreases with the smoothing parameter λ . The control of this term is at the core of the analysis, and is carried out in Section 6.3.

Proof of Lemma 3. By Lemma 20, for any $p, q \in \mathbb{R}^+$, if $q \ge p/8$ then $D(p,q) \le \phi(8)(\sqrt{p}-\sqrt{q})^2 \le 3(\sqrt{p}-\sqrt{q})^2$. On the other hand, if $q \le p/8$ then $D(p,q) = p\log(p/q) - p + q \le p\log(p/q)$. Hence, for any $p, q \in \mathbb{R}^+$, we have

$$D(p,q) \leq 3(\sqrt{p} - \sqrt{q})^2 + p \log\left(\frac{p}{q}\right) \mathbf{1}(q \leq p/8).$$
(59)

It follows from this inequality that

$$\operatorname{KL}(P,\widehat{P}_n) = \sum_{j=1}^d D(p_j,\widehat{p}_j) \leqslant 3\sum_{j=1}^d \left(\sqrt{\widehat{p}_j} - \sqrt{p_j}\right)^2 + \sum_{j=1}^d p_j \log\left(\frac{p_j}{\widehat{p}_j}\right) \mathbf{1}\left(\widehat{p}_j \leqslant \frac{p_j}{8}\right).$$
(60)

We now bound the two terms of the right-hand side of (60), starting with the first one. First, since $\hat{p}_j = (N_j + \lambda)/(n + \lambda d) = (\bar{p}_j + \lambda/n)/(1 + \lambda d/n)$, we have $\bar{p}_j/(1 + \lambda d/n) \leq \hat{p}_j \leq \bar{p}_j + \lambda/n$. We therefore have

$$\left(\sqrt{\widehat{p}_{j}} - \sqrt{p_{j}}\right)^{2} \leq 2\left(\sqrt{\widehat{p}_{j}} - \sqrt{\overline{p}_{j}}\right)^{2} + 2\left(\sqrt{\overline{p}_{j}} - \sqrt{p_{j}}\right)^{2}$$

$$\leq 2 \max\left\{\left(\sqrt{\overline{p}_{j}} - \sqrt{\overline{p}_{j}}\left(1 + \frac{\lambda d}{n}\right)^{-1/2}\right)^{2}, \left(\sqrt{\overline{p}_{j}} + \frac{\lambda}{n} - \sqrt{\overline{p}_{j}}\right)^{2}\right\} + 2\left(\sqrt{\overline{p}_{j}} - \sqrt{\overline{p}_{j}}\right)^{2}$$

$$\leq 2\left(1 - \left(1 + \frac{\lambda d}{n}\right)^{-1/2}\right)^{2}\overline{p}_{j} + 2\left(\sqrt{\overline{p}_{j}} + \frac{\lambda}{n} - \sqrt{\overline{p}_{j}}\right)^{2} + 2\left(\sqrt{\overline{p}_{j}} - \sqrt{\overline{p}_{j}}\right)^{2}.$$
(61)

Using that, for any $u \in \mathbb{R}^+$, one has $1-(1+u)^{-1/2} = u/(1+u+\sqrt{1+u}) \leq u/(2\sqrt{u}+\sqrt{u}) = \sqrt{u}/3$, we can bound the first term in (61) as

$$\left(1 - \left(1 + \frac{\lambda d}{n}\right)^{-1/2}\right)^2 \overline{p}_j \leqslant \left(\frac{\sqrt{\lambda d/n}}{3}\right)^2 \overline{p}_j \leqslant \frac{\lambda d}{9n} \cdot \overline{p}_j$$

Likewise, we can bound the second term as follows:

$$\left(\sqrt{\overline{p}_j + \lambda/n} - \sqrt{\overline{p}_j}\right)^2 = \left(\frac{(\overline{p}_j + \lambda/n) - \overline{p}_j}{\sqrt{\overline{p}_j + \lambda/n} + \sqrt{\overline{p}_j}}\right)^2 \leqslant \left(\frac{\lambda/n}{\sqrt{\lambda/n}}\right)^2 = \frac{\lambda}{n}$$

Plugging these two bounds into (61), summing over j = 1, ..., d and using that $\sum_{j=1}^{d} \overline{p}_j = 1$, we deduce that

$$\sum_{j=1}^{d} \left(\sqrt{\widehat{p}_{j}} - \sqrt{p_{j}}\right)^{2} \leqslant \frac{2\lambda d}{9n} \cdot \sum_{j=1}^{d} \overline{p}_{j} + \frac{2\lambda d}{n} + 2\sum_{j=1}^{d} \left(\sqrt{\overline{p}_{j}} - \sqrt{p_{j}}\right)^{2}$$
$$= \frac{20\lambda d}{9n} + 2\sum_{j=1}^{d} \left(\sqrt{\overline{p}_{j}} - \sqrt{p_{j}}\right)^{2}.$$
(62)

We now turn to bounding the second term in the decomposition (60). First, since $\lambda \leq n/d$, we have $n + \lambda d \leq 2n$, and thus $\hat{p}_j \geq \max\{N_j, \lambda\}/(2n)$. This implies that

$$\sum_{j=1}^{d} p_j \log\left(\frac{p_j}{\widehat{p}_j}\right) \mathbf{1}\left(\widehat{p}_j \leqslant \frac{p_j}{8}\right) \leqslant \sum_{j=1}^{d} p_j \log\left(\frac{p_j}{\lambda/(2n)}\right) \mathbf{1}\left(\frac{\max\{N_j,\lambda\}}{2n} \leqslant \frac{p_j}{8}\right)$$
$$= \sum_{j: p_j \geqslant 4\lambda/n} p_j \log\left(\frac{2np_j}{\lambda}\right) \mathbf{1}\left(N_j \leqslant \frac{np_j}{4}\right). \tag{63}$$

Plugging upper bounds (62) and (63) into the decomposition (61) concludes the proof. \Box

6.2 Upper bound in Hellinger distance

In this section, we proceed with the control of the first term in the decomposition of Lemma 3. Specifically, Lemma 4 below provides a high-probability bound on the squared Hellinger distance between the empirical distribution \overline{P}_n and the true distribution P. This bound is established by combining the the analysis of the reverse-relative entropy from [Agr22] (for large probabilities) with concentration inequalities deduced from negative association [JDP83] (for small probabilities).

Lemma 4. Let $n, d \ge 2$ and $\delta \in (0, 1)$. For any $P \in \mathcal{P}_d$, letting $s_n(P) = \sum_{j=1}^d \min(1, np_j)$ and denoting by $\overline{P}_n = (\overline{p}_1, \ldots, \overline{p}_d)$ the empirical distribution of X_1, \ldots, X_n , one has

$$\mathbb{P}_P\left(\sum_{j=1}^d \left(\sqrt{\overline{p}_j} - \sqrt{\overline{p}_j}\right)^2 \geqslant \frac{4s_n(P) + 7\log(1/\delta)}{n}\right) \leqslant 2\delta.$$
(64)

We note in passing that this bound is of the right order of magnitude, as one may check that $\mathbb{E}_P[\sum_{j=1}^d (\sqrt{\overline{p_j}} - \sqrt{\overline{p_j}})^2] \approx s_n(P)/n.$

Proof of Lemma 4. Let $J^+ = \{1 \leq j \leq d : p_j \geq 1/n\}$ and $J^- = \{1, \ldots, d\} \setminus J^+$. Also, let $s^+ = |J^+|$ and $s^- = n \sum_{j \in J^-} p_j$, so that $s_n(P) = s^+ + s^-$. From the inequalities $(\sqrt{p} - \sqrt{q})^2 \leq D(p,q)$ (Lemma 20) and $(\sqrt{p} - \sqrt{q})^2 \leq p + q$, it follows that

$$\sum_{j=1}^{a} \left(\sqrt{\overline{p}_j} - \sqrt{\overline{p}_j} \right)^2 \leqslant \sum_{j \in J^+} D(\overline{p}_j, p_j) + \sum_{j \in J^-} \left(p_j + \overline{p}_j \right).$$
(65)

First term. We start by controlling the first term in (65), following [Agr22]. Specifically, let $H^+ = \sum_{j \in J^+} D(\overline{p}_j, p_j)$. By proceeding as in [Agr22, Corollary 1.7], except that in [Agr22, Proposition 2.4] we sum only over indices $j \in J^+$ (rather than over $1 \leq j \leq d$), we obtain the following inequality: for any $t \in (0, n/2)$,

$$\log \mathbb{E}_P\left[e^{t(H^+ - \mathbb{E}_P[H^+])}\right] \leqslant \frac{4s^+ t^2/n^2}{1 - 2t/n}.$$

In other words, $H^+ - \mathbb{E}_P[H^+]$ is sub-gamma [BLM13, §2.4] with variance factor $8s^+/n^2$ and shape parameter 2/n. Hence, by [BLM13, p. 29], one has for every $\delta \in (0, 1)$,

$$\mathbb{P}\left(H^{+} - \mathbb{E}_{P}[H^{+}] \geqslant \frac{4\sqrt{s^{+}\log(1/\delta)}}{n} + \frac{2\log(1/\delta)}{n}\right) \leqslant \delta.$$
(66)

In addition, recalling the inequality $h(t) \leq (t-1)^2$ for any $t \in \mathbb{R}^+$ (Lemma 14), we have for all $p, q \in \mathbb{R}^+, q > 0$:

$$D(p,q) = qh\left(\frac{p}{q}\right) \leqslant q\left(\frac{p}{q}-1\right)^2 = \frac{(p-q)^2}{q}$$

This implies that

$$\mathbb{E}_P[H^+] = \sum_{j \in J^+} \mathbb{E}_P[D(\overline{p}_j, p_j)] \leqslant \sum_{j \in J^+} \frac{\mathbb{E}_P[(\overline{p}_j - p_j)^2]}{p_j} = \sum_{j \in J^+} \frac{p_j(1 - p_j)/n}{p_j} \leqslant \frac{s^+}{n}$$

Plugging the inequality into (66), we deduce that with probability at least $1 - \delta$, we have

$$H^{+} < \frac{s^{+} + 4\sqrt{s^{+}\log(1/\delta)} + 2\log(1/\delta)}{n} \leqslant \frac{(1+2\sqrt{2})s^{+} + (2+2\sqrt{2})\log(1/\delta)}{n}.$$
 (67)

Second term. We now turn to the second term in (65). Note that $\sum_{j \in J^-} \overline{p}_j = n^{-1} \sum_{j \in J^-} N_j$. In addition, by [JDP83, 3.1(a)], the variables $(N_j)_{1 \leq j \leq d}$ are negatively associated, in the sense of [JDP83, Definition 2.1]. This implies in particular that, for every $t \in \mathbb{R}^+$,

$$\mathbb{E}\left[\exp\left(t\sum_{j\in J^{-}}N_{j}\right)\right] \leqslant \prod_{j\in J^{-}}\mathbb{E}\left[e^{tN_{j}}\right] = \prod_{j\in J^{-}}\left(1-p_{j}+p_{j}e^{t}\right)^{n}$$
$$\leqslant \prod_{j\in J^{-}}\exp\left\{np_{j}(e^{t}-1)\right\} = \exp\left\{n\sum_{j\in J^{-}}p_{j}(e^{t}-1)\right\} = \exp\left\{s^{-}(e^{t}-1)\right\}.$$

Rearranging and using that $e^t - 1 - t = \sum_{k \ge 2} \frac{t^k}{k!} \leqslant \sum_{k \ge 2} \frac{t^2}{2} = \frac{t^2}{2(1-t)}$ for any $t \in [0, 1)$, we obtain that for any $0 \leqslant t < 1$:

$$\log \mathbb{E}\left[\exp\left(t\left\{\sum_{j\in J^{-}}N_{j}-s^{-}\right\}\right)\right] \leqslant s^{-}(e^{t}-1-t) \leqslant \frac{s^{-}t^{2}}{2(1-t)}$$

Using again the sub-gamma tail bound [BLM13, p. 29], we conclude that, with probability at least $1 - \delta$,

$$\sum_{j \in J^-} N_j < s^- + \sqrt{2s^- \log(1/\delta)} + \log(1/\delta) \leqslant 2(s^- + \log(1/\delta)),$$

namely

$$\sum_{j \in J^{-}} \left(p_j + \overline{p}_j \right) = n^{-1} \left(s^- + \sum_{j \in J^-} N_j \right) < \frac{3s^- + 2\log(1/\delta)}{n} \,. \tag{68}$$

Conclusion. For $\delta \in (0, 1/2)$, with probability at least $1 - 2\delta$ both (67) and (68) hold. In this case, inequality (65) together with the fact that $s^+ + s^- = s_n(P)$ and the bound $2\sqrt{2} \leq 3$ imply that

$$\sum_{j=1}^d \left(\sqrt{\overline{p}_j} - \sqrt{p_j}\right)^2 < \frac{4s_n(P) + 7\log(1/\delta)}{n}$$

This concludes the proof.

6.3 Control of the contribution of underestimated frequencies

The control of the second term, namely the bias due to the use of regularization, is immediate when the parameter λ is not data-dependent; we also postpone its analysis in the context of data-dependent regularization to the proof of Theorem 5 below.

We now turn to the control of the key term in the decomposition of Lemma 3, namely the third term accounting for the contribution of classes whose true frequency is significantly underestimated in the sample. Specifically, in this section we establish the following control on the residual, which is arguably the core of the analysis:

Lemma 5. Let $P \in \mathcal{P}_d$. For any $\lambda \ge 1$, let $d_{\lambda} = |\{1 \le j \le d : p_j \ge 4\lambda/n\}|$ and

$$R_{\lambda} = \sum_{j: p_j \geqslant 4\lambda/n} p_j \log\left(\frac{2np_j}{\lambda}\right) \mathbf{1}\left(N_j \leqslant \frac{np_j}{4}\right).$$

For any $\delta \in (e^{-n/6}, e^{-2})$, one has

$$\mathbb{P}_P\left(R_1 \ge \frac{62000\,d_1 + 106000\log(1/\delta)\log\log(1/\delta)}{n}\right) \le 2\delta\,.\tag{69}$$

In addition, for any $\delta \in (e^{-n/6}, e^{-d}]$, if $\lambda = \log(1/\delta)/d$, then

$$\mathbb{P}_P\left(R_\lambda \geqslant \frac{74000\log(d)\log(1/\delta)}{n}\right) \leqslant 2\delta.$$
(70)

Roughly speaking, inequality (69) will be used in the analysis of confidence-independent estimators such as the Laplace estimator (or of confidence-dependent estimators in the low confidence regime), while (70) will be used in the analysis of the confidence-dependent estimators in the high-confidence regime.

Observe that the residual R_{λ} of Lemma 5 is a sum of dependent random variables, since the counts $(N_j)_{1 \leq j \leq d}$ are dependent. Up to a standard technique of Poisson sampling, one may reduce its control to that of a "Poissonized" sum involving independent summands, stated next:

Lemma 6. Let $\widetilde{N}_1, \ldots, \widetilde{N}_d$ be independent random variables, with $\widetilde{N}_j \sim \mathsf{Poisson}(\lambda_j/2)$ for $j = 1, \ldots, d$. For any $\lambda \ge 1$, let $d_\lambda = |\{1 \le j \le d : \lambda_j \ge 4\lambda\}|$ and

$$\widetilde{R}_{\lambda} = \sum_{j: \lambda_j \ge 4\lambda} \lambda_j \log\left(\frac{2\lambda_j}{\lambda}\right) \mathbf{1}\left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4}\right).$$

For any $\delta \in (0, e^{-2})$, one has

$$\mathbb{P}\left(\widetilde{R}_1 \ge 62000 \, d_1 + 106000 \log(1/\delta) \log\log(1/\delta)\right) \le \delta \,. \tag{71}$$

In addition, for any $\delta \in (0, e^{-d}]$, if $\lambda = \log(1/\delta)/d$, then

$$\mathbb{P}\Big(\widetilde{R}_{\lambda} \ge 74000 \log(d) \log(1/\delta)\Big) \le \delta.$$
(72)

Proof of Lemma 5 from Lemma 6. We resort to the technique of Poisson sampling. Specifically, let $(X_i)_{i \ge n+1}$ be an i.i.d. sequence of random variables with distribution P, independent from X_1, \ldots, X_n . In addition, let $N \sim \mathsf{Poisson}(n/2)$ be independent from the sequence $(X_i)_{i \ge 1}$. For $j = 1, \ldots, d$, define

$$\widetilde{N}_j = \sum_{1 \leqslant i \leqslant N} \mathbf{1}(X_i = j) \,. \tag{73}$$

It is a classical fact about Poisson random variables (which follows, e.g., from the combination of [MU17, Theorem 5.6 p. 100] and [Dur10, Exercise 2.1.11 p. 55]) that $\tilde{N}_1, \ldots, \tilde{N}_d$ are independent random variables, with $\tilde{N}_j \sim \mathsf{Poisson}(np_j/2)$ for $j = 1, \ldots, d$. Consider now the event $E = \{N \leq n\}$; by the Poisson deviation bound (Lemma 17), one has

$$\mathbb{P}(E) = 1 - \mathbb{P}(N > n) \ge 1 - e^{-D(n,n/2)} \ge 1 - e^{-n/6} \ge 1 - \delta.$$
(74)

In addition, under E one has $\widetilde{N}_j \leq N_j$ for $j = 1, \ldots, d$, hence letting $\lambda_j = np_j$,

$$R_{\lambda} = \sum_{j: p_j \geqslant 4\lambda/n} p_j \log\left(\frac{2np_j}{\lambda}\right) \mathbf{1}\left(N_j \leqslant \frac{np_j}{4}\right) \leqslant \frac{1}{n} \sum_{j: \lambda_j \geqslant 4\lambda} \lambda_j \log\left(\frac{2\lambda_j}{\lambda}\right) \mathbf{1}\left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4}\right).$$
(75)

The right-hand side of (75) is controlled in Lemma 6, which concludes the proof.

We now turn to the proof of Lemma 6. It should be noted that \widetilde{R}_{λ} is a nonnegative weighted sum of independent Bernoulli variables, with varying coefficient and parameters. Perhaps a natural approach to control such a sum is to apply Bennett's inequality [BLM13, Theorem 2.9 p. 35]. Unfortunately, this would lead to a highly suboptimal tail bound. Roughly speaking, the reason why Bennett's inequality fails to capture the right tail behavior of this sum is that it is highly inhomogeneous, in the sense that the coefficients of the independent Bernoulli variables may be of different orders of magnitude. In order to handle this structure, we instead evaluate the upper envelope of the tails of the individual summands, and then resort to a sharp estimate from Latała [Lat97] on moments of sums of independent random variables.

We start with the control on the tails of the individual summands that comprise \widetilde{R}_{λ} :

Lemma 7. For any $\lambda \ge 1$, let $W^{(\lambda)}$ be a random variable such that $\mathbb{P}(W^{(\lambda)} = 0) = 1 - e^{-2\lambda/7}$, and $\mathbb{P}(W^{(\lambda)} \ge t \log(t/\lambda)) = e^{-t/14}$ for any $t \ge 4\lambda$. Then, for any $j = 1, \ldots, d$ such that $\lambda_j \ge 4\lambda$, the random variable

$$V_j^{(\lambda)} = \lambda_j \log\left(\frac{\lambda_j}{\lambda}\right) \mathbf{1}\left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4}\right)$$

is stochastically dominated by $W^{(\lambda)}$, in the sense that $\mathbb{P}(V_j^{(\lambda)} \ge w) \le \mathbb{P}(W^{(\lambda)} \ge w)$ for any $w \in \mathbb{R}$.

In particular, for $\lambda = 1$, Lemma 7 asserts that for any $\delta < e^{-2/7}$, the quantile of order $1 - \delta$ of $V_j^{(1)}$ is smaller than $14 \log(1/\delta) \log(14 \log(1/\delta)) \approx \log(1/\delta) \log \log(1/\delta)$, regardless of the value of $\lambda_j \ge 4$. By independence of the variables $V_j^{(\lambda)}$ and by Lemma 18, we deduce from Lemma 7 that the random variable \tilde{R}_{λ} is stochastically dominated by $\sum_{j=1}^{d_{\lambda}} W_j^{(\lambda)}$, where $W_1^{(\lambda)}, W_2^{(\lambda)}, \ldots$ are i.i.d. random variables with the same distribution as $W^{(\lambda)}$.

Proof of Lemma 7. First, since $\widetilde{N}_j \sim \mathsf{Poisson}(\lambda_j/2)$, by the Poisson deviation bound (Lemma 17) we have

$$\mathbb{P}\left(\widetilde{N}_{j} \leqslant \frac{\lambda_{j}}{4}\right) \leqslant \exp\left(-D\left(\frac{\lambda_{j}}{4}, \frac{\lambda_{j}}{2}\right)\right) = \exp\left(-(1 - \log 2)\lambda_{j}/4\right) \leqslant e^{-\lambda_{j}/14}.$$
 (76)

We need to show that $\mathbb{P}(V_j^{(\lambda)} \ge w) \le \mathbb{P}(W^{(\lambda)} \ge w)$ for any $w \in \mathbb{R}$. For $w \le 0$, both probabilities are equal to 1. For $0 < w \le 4\log(4)\lambda$, this is a consequence of (76) as $\mathbb{P}(V_j^{(\lambda)} \ge w) \le \mathbb{P}(\widetilde{N}_j \le \lambda_j/4) \le e^{-\lambda_j/14} \le e^{-4\lambda/14} = e^{-2\lambda/7} = \mathbb{P}(W^{(\lambda)} \ge w)$.

For $w > 4 \log(4)\lambda$, using that the map $t \mapsto t \log(t/\lambda)$ is an increasing bijection from $(4\lambda, +\infty)$ to $(4 \log(4)\lambda, +\infty)$, we may write $w = t \log(t/\lambda)$ for some $t > 4\lambda$. There are now two cases. First, if $\lambda_j \ge t$, then by (76) $\mathbb{P}(V_j^{(\lambda)} \ge w) \le \mathbb{P}(\widetilde{N}_j \le \lambda_j/4) \le e^{-\lambda_j/14} \le e^{-t/14} = \mathbb{P}(W^{(\lambda)} \ge w)$. On the other hand, if $\lambda_j < t$, then $V_j^{(\lambda)} \le \lambda_j \log(\lambda_j/\lambda) < t \log(t/\lambda)$, thus $\mathbb{P}(V_j^{(\lambda)} \ge w) = 0$ and the inequality also holds.

We now aim to obtain a high-probability bound on the sum $\sum_{j=1}^{d_{\lambda}} W_j^{(\lambda)}$, which is a sum of i.i.d. nonnegative random variables. Unfortunately, since $W^{(\lambda)}$ has super-exponential tails, its moment generating functions is infinite at any positive value, thus one cannot rely on the Chernoff approach to obtain such a tail bound. We will instead work with its moments, using a moment estimate of Latała [Lat97]. Specifically, Lemma 8 below is a consequence of (part of) [Lat97, Corollary 1] obtained by tracking the numerical constants in this result. We recall that, for any $p \in \mathbb{R}^+$ such that $p \ge 1$ and any real random variable Z, we let $||Z||_p = \mathbb{E}[|Z|^p]^{1/p} \in \mathbb{R}^+ \cup \{+\infty\}$.

Lemma 8 ([Lat97], Corollary 1). Let Z, Z_1, \ldots, Z_m be *i.i.d.* nonnegative random variables. Then, for any $p \in [1, +\infty)$,

$$\left\|\sum_{i=1}^{m} Z_i\right\|_p \leq 2e^2 \sup\left\{\frac{p}{s}\left(\frac{m}{p}\right)^{1/s} \|Z\|_s : \max\left(1, \frac{p}{m}\right) \leq s \leq p\right\}.$$

In order to apply Lemma 8, we need to control the L^p norm of $W^{(\lambda)}$. This is achieved in the following lemma.

Lemma 9. For any $p \in [1, +\infty)$, one has

$$\left\|W^{(\lambda)}\right\|_{p} \leq 215p \log\left(\max\left\{e, \frac{50p}{\lambda}\right\}\right).$$
(77)

Proof of Lemma 9. We have

$$\begin{split} \left\|W^{(\lambda)}\right\|_{p}^{p} &= \mathbb{E}\left[(W^{(\lambda)})^{p}\right] = \int_{0}^{\infty} \mathbb{P}\left((W^{(\lambda)})^{p} \ge u\right) \mathrm{d}u \\ &= \int_{0}^{(4\log(4)\lambda)^{p}} e^{-2\lambda/7} \mathrm{d}u + \int_{4\log(4)\lambda}^{\infty} \mathbb{P}\left(W^{(\lambda)} \ge w\right) pw^{p-1} \mathrm{d}w \\ &= (4\log(4)\lambda)^{p} e^{-2\lambda/7} + \int_{4\lambda}^{\infty} \mathbb{P}\left(W^{(\lambda)} \ge t\log\left(\frac{t}{\lambda}\right)\right) p\left(t\log\left(\frac{t}{\lambda}\right)\right)^{p-1} \log\left(\frac{et}{\lambda}\right) \mathrm{d}t \\ &\leqslant (4\log(4)\lambda)^{p} e^{-2\lambda/7} + \frac{p}{2\lambda} \int_{4\lambda}^{\infty} e^{-t/14} \left(t\log\left(\frac{t}{\lambda}\right)\right)^{p} \mathrm{d}t \,, \end{split}$$
(78)

where we used that $\log(et/\lambda) \leq 2\log(t/\lambda) \leq t\log(t/\lambda)/(2\lambda)$. Now, for any $t \geq 4\lambda$, let

$$\phi_{\lambda}(t) = \log\left\{e^{-t/28}\left(t\log\left(\frac{t}{\lambda}\right)\right)^{p}\right\} = p\log(t) + p\log\log\left(\frac{t}{\lambda}\right) - \frac{t}{28}$$

Since

$$\phi_{\lambda}'(t) = \frac{p}{t} + \frac{p}{t\log(t/\lambda)} - \frac{1}{28} \leqslant \left(1 + \frac{1}{\log 4}\right)\frac{p}{t} - \frac{1}{28}$$

we deduce that $\phi'_{\lambda}(t) < 0$ for any $t \ge 50p$, thus ϕ_{λ} decreases over $[\max\{50p, 4\lambda\}, +\infty)$.

Small *p*. We first consider the case where $50p \leq 4\lambda$. In this case, we get

$$\sup_{t \ge 4\lambda} \left\{ e^{-t/28} \left(t \log\left(\frac{t}{\lambda}\right) \right)^p \right\} = e^{-\lambda/7} \left(4 \log(4)\lambda \right)^p$$

and therefore

$$\frac{p}{2\lambda} \int_{4\lambda}^{\infty} e^{-t/14} \left(t \log\left(\frac{t}{\lambda}\right) \right)^p \mathrm{d}t \leqslant \frac{1}{25} e^{-\lambda/7} \left(4 \log(4)\lambda \right)^p \int_{4\lambda}^{\infty} e^{-t/28} \mathrm{d}t = \frac{28}{25} \left(4 \log(4)\lambda \right)^p e^{-2\lambda/7}$$

Plugging this inequality into (78), we deduce that

$$\left\|W^{(\lambda)}\right\|_{p}^{p} \leq 2.2 \left(4\log(4)\lambda\right)^{p} e^{-2\lambda/7}$$

We now apply the bound $(ex/p)^p \leq e^x$ for $x \in \mathbb{R}^+$ to $x = \lambda/7$, which gives

$$\left\|W^{(\lambda)}\right\|_{p}^{p} \leq 2.2 \left(4\log(4)\right)^{p} \left(\frac{7p}{e}\right)^{p} e^{-\lambda/7} \leq 2.2 (15p)^{p} e^{-12.5p/7} \leq 2.2 (2.4p)^{p},$$

so that $||W^{(\lambda)}||_p \leq (2.2)^{1/p} \cdot 2.4p \leq 6p.$

Large *p*. We now turn to the case where $4\lambda \leq 50p$. In this case, we get

$$\sup_{t \ge 4\lambda} \left\{ e^{-t/28} \left(t \log\left(\frac{t}{\lambda}\right) \right)^p \right\} = \sup_{4\lambda \leqslant t \leqslant 50p} \left\{ e^{-t/28} \left(t \log\left(\frac{t}{\lambda}\right) \right)^p \right\}$$
$$\leqslant \sup_{t \in \mathbb{R}^+} \left\{ e^{-t/28} t^p \right\} \left(\log\left(\frac{50p}{\lambda}\right) \right)^p = e^{-p} (28p)^p \left(\log\left(\frac{50p}{\lambda}\right) \right)^p.$$

This implies that

$$\begin{split} \frac{p}{2\lambda} \int_{4\lambda}^{\infty} e^{-t/14} \Big(t \log\left(\frac{t}{\lambda}\right) \Big)^p \mathrm{d}t &\leqslant \frac{p}{2} \Big(28e^{-1}p \log\left(\frac{50p}{\lambda}\right) \Big)^p \int_{4\lambda}^{\infty} e^{-t/28} \mathrm{d}t \\ &\leqslant 14p \Big(28e^{-1}p \log\left(\frac{50p}{\lambda}\right) \Big)^p \,. \end{split}$$

Plugging this inequality into (78) and using the bound $(4 \log(4)\lambda)^p e^{-2\lambda/7} \leq (2.4p)^p$ shown above, we get

$$\begin{split} \left\|W^{(\lambda)}\right\|_{p} &\leqslant \left[(2.4p)^{p} + 14p\left(28e^{-1}p\log\left(\frac{50p}{\lambda}\right)\right)^{p}\right]^{1/p} \\ &\leqslant 2.4p + 14p^{1/p} \cdot 28e^{-1}p\log\left(\frac{50p}{\lambda}\right) \\ &\leqslant 2.4p + 210p\log\left(\frac{50p}{\lambda}\right) \leqslant 215\log\left(\frac{50p}{\lambda}\right). \end{split}$$

Hence, the desired bound holds for all $p \ge 1$.

Combining the moment estimate of Lemma 9 with Lemma 8, we obtain the following control on the moments of the term that dominates residuals when $\lambda = 1$.

Lemma 10. For any $p \ge 1$, one has

$$\left\|\sum_{j=1}^{d_1} W_j^{(1)}\right\|_p \le 15000d_1 + 4600p\log(50p)\,.$$
⁽⁷⁹⁾

Proof of Lemma 10. We start with the case where $p \ge d_1$. In this case, plugging Lemma 9 into Lemma 8 and bounding $d_1/p \le 1$ gives:

$$\begin{split} \left\| \sum_{j=1}^{d_1} W_j^{(1)} \right\|_p &\leq 2e^2 \sup\left\{ \frac{p}{s} \left(\frac{d_1}{p} \right)^{1/s} 215s \log\left(\max\left\{ e, 50s \right\} \right) : \max\left(1, \frac{p}{d_1} \right) \leqslant s \leqslant p \right\} \\ &\leq 2e^2 \sup_{1 \leqslant s \leqslant p} \left\{ 215p \log(50s) \right\} = 430e^2 p \log\left(50p \right). \end{split}$$

We now turn to the case where $p \leq d_1$. In this case, Lemmas 8 and 9 imply that

$$\left\|\sum_{j=1}^{a_1} W_j^{(1)}\right\|_p \leqslant 430e^2 \sup_{1\leqslant s\leqslant p} \left\{ p\left(\frac{d_1}{p}\right)^{1/s} \log(50s) \right\}.$$

We bound the supremum over $1 \leq s \leq p$ as follows. If $1 \leq s \leq 2$, then (using that $d_1/p \geq 1$)

$$p\left(\frac{d_1}{p}\right)^{1/s}\log(50s) \leqslant p\left(\frac{d_1}{p}\right)^{1/s}\log(100) \leqslant p\left(\frac{d_1}{p}\right)\log(100) = \log(100)d_1.$$

If $s \ge 2$, we consider two cases. If $s \le \sqrt{d_1/p}$, then

$$p\left(\frac{d_1}{p}\right)^{1/s}\log(50s) \leqslant p\left(\frac{d_1}{p}\right)^{1/2}\log\left(50\sqrt{\frac{d_1}{p}}\right) = d_1\frac{\log(50) + \log(\sqrt{d_1/p})}{\sqrt{d_1/p}} \leqslant (\log 50 + e^{-1})d_1.$$

Finally, if $\sqrt{d_1/p} \leqslant s \leqslant p$, then

,

$$p\left(\frac{d_1}{p}\right)^{1/s}\log(50s) \leqslant ps^{1/s}\log(50p) \leqslant e^{1/e}p\log(50p)$$

Putting things together, for any value of p one has

$$\left\|\sum_{j=1}^{d_1} W_j^{(1)}\right\|_p \leq 430e^2 \max\left\{\log(100)d_1, e^{1/e}p\log(50p)\right\},\$$

which proves (79) after bounding the maximum by a sum and simplifying constants. \Box

Next, we obtain similarly a control of suitable large moments of the sum.

Lemma 11. Assume that $p \ge d$ and that $\lambda = p/d$. Then

$$\left\|\sum_{j=1}^{d_{\lambda}} W_{j}^{(\lambda)}\right\|_{p} \leq 3200 \log(50d)p.$$

$$(80)$$

Proof of Lemma 11. Since $p \ge d$, we have in particular that $d_{\lambda}/p \le 1$. Plugging Lemma 9 into Lemma 8 therefore gives:

$$\begin{split} \left\|\sum_{j=1}^{a_{\lambda}} W_{j}^{(\lambda)}\right\|_{p} &\leq 2e^{2} \sup\left\{\frac{p}{s} \left(\frac{d_{\lambda}}{p}\right)^{1/s} 215s \log\left(\max\left\{e, \frac{50s}{\lambda}\right\}\right) : \max\left(1, \frac{p}{d_{\lambda}}\right) \leqslant s \leqslant p\right\} \\ &\leq 2e^{2} \sup_{1\leqslant s\leqslant p} \left\{215p \log\left(\max\left\{e, \frac{50s}{\lambda}\right\}\right)\right\} \\ &= 430e^{2}p \log \max\left\{e, \frac{50p}{\lambda}\right\} = 430e^{2}p \log\left(50d\right). \end{split}$$

The conclusion follows by bounding $430e^2 \leq 3200$.

With these results in place, we can conclude the proof of Lemma 6—and thus, of Lemma 5.

Proof of Lemma 6. Denote by \preccurlyeq the stochastic domination relation between real random variables (Definition 4). As noted above, by independence of $\widetilde{N}_1, \ldots, \widetilde{N}_d$ it follows from Lemmas 7 and 18 that

$$\widetilde{R}_{\lambda} \leqslant \frac{3}{2} \sum_{j: \lambda_j \geqslant 4\lambda} \lambda_j \log\left(\frac{\lambda_j}{\lambda}\right) \mathbf{1}\left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4}\right) \preccurlyeq \frac{3}{2} \sum_{j=1}^{a_{\lambda}} W_j^{(\lambda)}.$$
(81)

(Above, we used that $\log(2\lambda_j/\lambda) \leq \frac{3}{2}\log(\lambda_j/\lambda)$ for $\lambda_j/\lambda \geq 4$.) It therefore suffices to establish tail bounds on the right-hand side of (81). In both cases $\lambda = 1$ and $\lambda = \log(1/\delta)/d$, we deduce such tail bounds from the moment bounds of Lemmas 10 and 11, respectively, together with the following inequality: for any real random variable Z and $p \geq 1$,

$$\mathbb{P}(|Z| \ge e ||Z||_p) = \mathbb{P}(|Z|^p \ge e^p \mathbb{E}[|Z|^p]) \le e^{-p},$$

applied to $p = \log(1/\delta)$. Inequalities (71) and (72) are obtained by further bounding constants, using in particular that $\log(50) \leq \log_2(50) \min\{\log(1/\delta), \log d\}$ as $\delta < e^{-2}$ and $d \ge 2$.

In the following sections, we proceed with the proofs of Theorem 1, 3 and 5—the first two directly obtained by combining the results above.

6.4 Proof of Theorem 1

We apply the decomposition of Lemma 3 with $\lambda = 1$. The first term in this decomposition is bounded through the Hellinger bound of Lemma 4, together with the inequality $s_n(P) \leq d$. The second term is equal to 7d/n. Finally, the third term is bounded through the bound (69) on R_1 from Lemma 5. Putting things together and using a union bound, we obtain, for any $\delta \in (e^{-n/6}, e^{-2})$,

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \ge 6 \times \frac{4d + 7\log(1/\delta)}{n} + \frac{7d}{n} + \frac{62000d + 106000\log(1/\delta)\log\log(1/\delta)}{n}\right) \le 4\delta.$$

Further bounding constants gives the bound of Theorem 1.

6.5 Proof of Theorem 3

We consider two cases. If $\log(1/\delta) \leq d$, then $\lambda_{\delta} = 1$ and $\widehat{P}_{n,\delta}$ coincides with the Laplace estimator. Then, Theorem 1 ensures that, with probability at least 4δ ,

$$\operatorname{KL}(P, \widehat{P}_{n,\delta}) \leqslant 110000 \frac{d + \log(1/\delta) \log \log(1/\delta)}{n} \leqslant 110000 \frac{d + \log(1/\delta) \log d}{n} \,. \tag{82}$$

On the other hand, if $\log(1/\delta) > d$, namely $\delta \in (e^{-n/6}, e^{-d})$, then $\lambda_{\delta} = \log(1/\delta)/d > 1$. We then proceed similarly to the proof of Theorem 1, with only two changes: the second term in the decomposition of Lemma 3 now equals $7\lambda_{\delta}d/n = 7\log(1/\delta)/n$, while the third term is now controlled using the bound (70) on $R_{\lambda_{\delta}}$ from Lemma 5. This gives, with probability at least $1 - 4\delta$,

$$\mathrm{KL}(P, \widehat{P}_{n, \delta}) < 6 \times \frac{4d + 7\log(1/\delta)}{n} + \frac{7\log(1/\delta)}{n} + \frac{74000\log(d)\log(1/\delta)}{n}$$

which also implies the desired tail bound.

6.6 Proof of Theorem 5

We are now in position to complete the proof of Theorem 5, up to Theorem 6 which we prove in Section 8 below.

In what follows, we fix $P \in \mathcal{P}_d$ and let $s_n = s_n(P)$ and $s_n^\circ = s_n^\circ(P)$. For now, let $\widetilde{\lambda}$ be either $\widehat{\lambda}$ or $\widehat{\lambda}_{\delta}$, and let \widetilde{P} be the add- $\widetilde{\lambda}$ estimator. First, the decomposition of Lemma 3 writes:

$$\operatorname{KL}(P,\widetilde{P}) \leqslant 6 \sum_{j=1}^{d} \left(\sqrt{\overline{p}_j} - \sqrt{\overline{p}_j} \right)^2 + \frac{7\widetilde{\lambda}d}{n} + \sum_{j: p_j \ge 4\widetilde{\lambda}/n} p_j \log\left(\frac{2np_j}{\widetilde{\lambda}}\right) \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right).$$
(83)

First term. The first term in the decomposition (83) is bounded through Lemma 4: with probability at least $1 - 2\delta$,

$$6\sum_{j=1}^{d} \left(\sqrt{\overline{p}_{j}} - \sqrt{p_{j}}\right)^{2} < \frac{24s_{n} + 42\log(1/\delta)}{n} \,. \tag{84}$$

Second term. For the second term, note that $\lambda \leq \hat{\lambda}_{\delta} \leq \max\{D_n, \log(1/\delta)\}/d$, and that by inequality (125) from Lemma 19, with probability at least $1 - \delta$,

$$D_n \leqslant 2\mathbb{E}_P[D_n] + 2\log(1/\delta)$$
.

Combining these inequalities and recalling that $\mathbb{E}_P[D_n] \leq s_n$ (Fact 1), we get: with probability $1 - \delta$,

$$\frac{7\widetilde{\lambda}d}{n} \leqslant \frac{14s_n + 14\log(1/\delta)}{n} \,. \tag{85}$$

Third term. We now turn to the control of the third term, which requires the most effort.

We first deal with the case where $\lambda > 1$, which directly reduces to Lemma 5. Indeed, since $\hat{\lambda} = D_n/d \leq 1$, one must have $\tilde{\lambda} = \tilde{\lambda}_{\delta} = \log(1/\delta)/d$, and $\delta \leq e^{-d}$. Hence, inequality (70) from Lemma 5 gives: with probability $1 - 2\delta$,

$$\sum_{j: p_j \ge 4\widetilde{\lambda}/n} p_j \log\left(\frac{2np_j}{\widetilde{\lambda}}\right) \mathbf{1}\left(N_j \le \frac{np_j}{4}\right) = R_{\lambda_\delta} \le \frac{74000 \log(d) \log(1/\delta)}{n} \,. \tag{86}$$

From now on, we assume that $\widetilde{\lambda} \leq 1$. In this case, we further decompose the third term (denoted $R_{\widetilde{\lambda}}$) as follows:

$$R_{\widetilde{\lambda}} = \sum_{j: p_j \ge 4\widetilde{\lambda}/n} p_j \log(2np_j) \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right) + \sum_{j: p_j \ge 4\widetilde{\lambda}/n} p_j \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right) \cdot \log\left(\frac{1}{\widetilde{\lambda}}\right).$$
(87)

We then bound the first term above as

$$\sum_{j: p_j \ge 4\tilde{\lambda}/n} p_j \log(2np_j) \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right)$$

$$= \sum_{j: 4\tilde{\lambda}/n \leqslant p_j < 4/n} p_j \log(2np_j) \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right) + \sum_{j: p_j \ge 4/n} p_j \log(2np_j) \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right)$$

$$\leqslant \log(8) \sum_{j: 4\tilde{\lambda}/n \leqslant p_j < 4/n} p_j \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right) + R_1 \leqslant \log(8) \sum_{j: p_j < 4/n} p_j + R_1$$

$$\leqslant \frac{4 \log(8) s_n}{n} + R_1.$$

Hence, bounding R_1 via the bound (69) of Lemma 5 and using that $d_1 = |\{j : p_j \ge 4/n\}| \le s_n$, we obtain: with probability $1 - 2\delta$,

$$\sum_{j: p_j \ge 4\tilde{\lambda}/n} p_j \log(2np_j) \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right) \leqslant \frac{4 \log(8) s_n}{n} + \frac{62000 s_n + 106000 \log(1/\delta) \log \log(1/\delta)}{n} \\ \leqslant \frac{63000 s_n + 106000 \log(1/\delta) \log \log(1/\delta)}{n} \,.$$

It remains to upper bound the second term in (87). Since $\lambda \geq \hat{\lambda} = D_n/d$, and recalling the definition of the underestimated mass U_n (Definition 3), we have

$$\sum_{j: p_j \ge 4\widetilde{\lambda}/n} p_j \mathbf{1}\left(N_j \leqslant \frac{np_j}{4}\right) \cdot \log\left(\frac{1}{\widetilde{\lambda}}\right) \leqslant \sum_{j=1}^d p_j \mathbf{1}\left(N_j \leqslant \frac{np_j}{4}\right) \cdot \log\left(\frac{d}{D_n}\right) = U_n \log\left(\frac{d}{D_n}\right)$$

Now, Theorem 6 shows that, with probability at least $1 - 8\delta$, one has

$$U_n \leqslant \frac{336 \, s_{n/112}^{\circ}(P) + 2500e \log(1/\delta)}{n}$$

Thus, under the same event (using that $D_n \ge 1$),

$$U_n \log\left(\frac{d}{D_n}\right) \leqslant \frac{336 \, s_{n/112}^\circ(P) \log(d/D_n) + 2500e \log(d) \log(1/\delta)}{n} \,. \tag{88}$$

We now consider two cases, depending on whether $\log(1/\delta)$ is larger or smaller than s_n . First, by deviation lower bound (124) from Lemma 19, letting $s'_n = \mathbb{E}_P[D_n]$ one has

$$\mathbb{P}\left(D_n \leqslant \frac{s'_n}{2}\right) \leqslant \exp\left\{-D\left(\frac{s'_n}{2}, s'_n\right)\right\} = \exp\left\{-\frac{1-\log 2}{2}s'_n\right\}.$$

But since $s'_n \ge (1 - e^{-1})s_n$ (Fact 1), after bounding constants we deduce that

$$\mathbb{P}(D_n \leqslant 0.3s_n) \leqslant e^{-s_n/21}$$

Thus, if $\log(1/\delta) \leq s_n/21$, then with probability at least $1 - \delta$ one has

$$D_n \ge 0.3s_n$$
,

which combined with (88) gives, with probability $1 - 9\delta$,

$$U_n \log\left(\frac{d}{D_n}\right) \leqslant \frac{336 \, s_{n/112}^{\circ}(P) \log(2ed/s_n) + 2500e \log(d) \log(1/\delta)}{n}$$

.

On the other hand, if $\log(1/\delta) > s_n/21$, then by lower-bounding $D_n \ge 1$ in (88) and using that

$$s_{n/112}^{\circ} \leqslant s_{n/112} \leqslant s_n \leqslant 21 \log(1/\delta)$$
,

we get with probability at least $1 - \delta$ that

$$U_n \log\left(\frac{d}{D_n}\right) \leqslant \frac{336 \times 21 \log(1/\delta) \times \log(d) + 2500e \log(d) \log(1/\delta)}{n} \leqslant \frac{14000 \log(d) \log(1/\delta)}{n}.$$

Summarizing, we get that regardless of δ , we have with probability at least $1 - 9\delta$ that

$$\sum_{j: p_j \ge 4\widetilde{\lambda}/n} p_j \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right) \cdot \log\left(\frac{1}{\widetilde{\lambda}}\right) \leqslant \frac{336 \, s_{n/112}^\circ(P) \log(2ed/s_n) + 14000 \log(d) \log(1/\delta)}{n}$$

Putting this inequality into the decomposition (87) of the third term, we get that whenever $\widetilde{\lambda} \leq 1$, we have with probability at least $1 - 11\delta$,

$$R_{\tilde{\lambda}} \leqslant \frac{63000 \, s_n + 106000 \log(1/\delta) \log \log(1/\delta)}{n} + \frac{336 \, s_{n/112}^{\circ}(P) \log(2ed/s_n) + 14000 \log(d) \log(1/\delta)}{n} \\ \leqslant \frac{64000 \, s_n + 336 \, s_{n/112}^{\circ}(P) \log(ed/s_n) + 120000 \max\{\log d, \log \log(1/\delta)\} \log(1/\delta)}{n} \,.$$
(89)

Conclusion. We first establish the bound (34) for the estimator $\widehat{P}_n^{\text{ad}}$. In this case, one has $\widehat{\lambda} = D_n/d \leq 1$, thus the bound (89) applies. Injecting this inequality, together with the bounds (84) and (85) on the first two terms, into the decomposition (83), we obtain: with probability at least at $1 - 14\delta$,

$$\begin{aligned} \operatorname{KL}(P, \widehat{P}_n^{\operatorname{ad}}) &< \frac{24s_n + 42\log(1/\delta)}{n} + \frac{14s_n + 14\log(1/\delta)}{n} + \\ &+ \frac{64000\,s_n + 336\,s_{n/112}^\circ(P)\log(ed/s_n) + 120000\max\{\log d, \log\log(1/\delta)\}\log(1/\delta)}{n}, \end{aligned}$$

which implies the claimed bound.

We now conclude with the estimator $\widehat{P}_{n,\delta}^{\mathrm{ad}}$. Then, either $\log(1/\delta) \leq d$, in which case $\widehat{\lambda}_{\delta} \leq 1$ and again the bound on the third term for $\widetilde{\lambda} \leq 1$ applies, so that the same bound as for $\widehat{P}_n^{\mathrm{ad}}$ holds. In addition, one has $\max\{\log d, \log \log(1/\delta)\} = \log d$ in this case. On the other hand, if $\log(1/\delta) > d$, then $\widehat{\lambda}_{\delta} > 1$, thus the third term is bounded using (86). Combining with the bounds (84) and (85) on the first two terms gives the desired inequality.

7 Proofs of lower bounds

In this section, we provide the proofs of the lower bounds stated in previous sections, specifically the tail (low-probability) lower bounds of Theorem 2, Lemma 1, Theorem 4 and Corollary 1, as well as the high-probability minimax lower bound of Proposition 1.

7.1 Proof of Theorem 2

We first prove Theorem 2 on confidence-independent estimators. The proof rests on the following lemma.

Lemma 12. Let $n \ge d \ge 2$ and $\kappa \ge 1$, and $\widehat{P}_n = \Phi(X_1, \ldots, X_n)$ be an estimator as in Theorem 2. Then, for any $\delta \in (e^{-n}, e^{-16\kappa^2})$, there exists a distribution $P \in \mathcal{P}_d$ such that

$$\mathbb{P}_{P}\left(\mathrm{KL}(P,\widehat{P}) \geqslant \frac{\log(1/\delta)\log\log(1/\delta)}{10n}\right) \geqslant \delta.$$
(90)

Proof of Lemma 12. Let $Q = \Phi(1, \ldots, 1) = (q_1, \ldots, q_d) \in \mathcal{P}_d$ be the value of the estimator when only the first class is observed. Clearly, if $P = \delta_1$, then $\hat{P}_n = Q$ almost surely, and thus condition (17) writes:

$$\operatorname{KL}(\delta_1, Q) = \log(1/q_1) \leqslant \frac{\kappa d}{n}.$$

Since $1 - q_1 \leq -\log(1 - (1 - q_1)) = \log(1/q_1)$, this implies that

$$\sum_{j=2}^d q_j = 1 - q_1 \leqslant \frac{\kappa d}{n} \,,$$

and thus there exists $j \in \{2, \ldots, d\}$ such that $q_j \leq (\kappa d/n)/(d-1) \leq 2\kappa/n$. Now, for $\delta \in$ $(e^{-n}, e^{-16\kappa^2})$, consider the distribution $P = P_{\delta,n} = (1-\rho)\delta_1 + \rho\delta_j$, where $\rho = 1 - \delta^{1/n}$. Then, the event $E = \{X_1 = \dots = X_n = 1\}$ is such that

$$\mathbb{P}_P(E) = (1-\rho)^n = \delta.$$

In addition, under E one has $\widehat{P} = Q$, thus

$$\operatorname{KL}(P, \widehat{P}_n) = \operatorname{KL}(P, Q) \ge D(\rho, q_j) = q_j h\left(\frac{\rho}{q_j}\right).$$

By convexity of the exponential function, one has $1 - e^{-x} \ge (1 - e^{-1})x$ for $x \in [0, 1]$; since $\frac{\log(1/\delta)}{n} \le 1$, this implies that $\rho = 1 - \exp\left(-\frac{\log(1/\delta)}{n}\right) \ge (1 - e^{-1})\frac{\log(1/\delta)}{n}$. Since $\delta \le e^{-16\kappa^2}$, we 1.10.9 therefore have

$$\frac{\rho}{q_j} \ge \frac{(1 - e^{-1})16\kappa^2/n}{2\kappa/n} = 8(1 - e^{-1})\kappa \ge e.$$

Hence, by Lemma 14 we have

$$q_{j}h\left(\frac{\rho}{q_{j}}\right) \ge q_{j} \times e^{-1}\frac{\rho}{q_{j}}\log\left(\frac{\rho}{q_{j}}\right) \ge e^{-1}(1-e^{-1})\frac{\log(1/\delta)}{n}\log\left(\frac{(1-e^{-1})\log(1/\delta)/n}{2\kappa/n}\right)$$
$$\ge \frac{\log(1/\delta)}{5n}\log\left(\frac{\log(1/\delta)}{4\kappa}\right) \ge \frac{\log(1/\delta)\log\log(1/\delta)}{10n}.$$

This concludes the proof of Lemma 12.

We can now conclude the proof of Theorem 2.

Proof of Theorem 2. The statement follows from a combination of the lower bounds of Lemma 12 and of the consequence (25) of Proposition 1. Specifically, one the one hand, since $d \ge 3300 \ge$ $3000 \log\left(\frac{3}{1-e^{-16}}\right)$, it follows from (25) that there exists $P \in \mathcal{P}_d$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \ge \frac{d}{4600n}\right) \ge 1 - 3\exp\left(-\frac{d}{3000}\right) \ge e^{-16} \ge \delta.$$

On the other hand, by Lemma 12, there exists $P \in \mathcal{P}_d$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \geqslant \frac{\log(1/\delta)\log\log(1/\delta)}{10n}\right) \geqslant \delta.$$

By taking the best of these two lower bounds (depending on d, δ), we deduce that there exists $P \in \mathcal{P}_d$ such that, with probability at least δ ,

$$\begin{split} \operatorname{KL}(P,\widehat{P}_n) &\ge \max\left\{\frac{d}{4600n}, \frac{\log(1/\delta)\log\log(1/\delta)}{10n}\right\} \\ &\ge \frac{99}{100} \cdot \frac{d}{4600n} + \frac{1}{100} \cdot \frac{\log(1/\delta)\log\log(1/\delta)}{10n} \ge \frac{d + \log(1/\delta)\log\log(1/\delta)}{5000n} \,, \end{split}$$
 blishes the claim. \Box

which establishes the claim.

7.2 Proof of Lemma 1 and Theorem 4

In this section, we establish Lemma 1 and then deduce Theorem 4.

Proof of Lemma 1. Fix n, d, δ as in Lemma 1. We define the class $\mathcal{F} = \mathcal{F}_{n,d,\delta}$ as

$$\mathcal{F} = \left\{ P^{(j)} = \delta^{1/n} \delta_1 + (1 - \delta^{1/n}) \delta_j : 1 \le j \le d \right\} = \left\{ \delta_1 \right\} \cup \left\{ P^{(j)} : 2 \le j \le d \right\}.$$

Let $\Phi : [d]^n \to \mathcal{P}_d$ be an estimator, and let $Q = (q_1, \ldots, q_d) = \Phi(1, \ldots, 1)$ denote the value of this estimator when only the first class is observed. Let $\alpha \in \mathbb{R}^+$ such that $1 - q_1 = \alpha d/n$.

First, assume that $\alpha \ge \log(1/\delta)/(7\sqrt{d})$. In this case, if $P = P^{(1)} = \delta_1$, then $\hat{P}_n = Q$ almost surely, hence

$$\operatorname{KL}(P, \widehat{P}_n) = \operatorname{KL}(\delta_1, Q) = \log(1/q_1) \ge 1 - q_1 = \frac{\alpha d}{n}.$$

Since $\alpha \ge \log(1/\delta)/(7\sqrt{d})$, we deduce that

$$\operatorname{KL}(P, \widehat{P}_n) \ge \frac{\sqrt{d}\log(1/\delta)}{7n} \ge \frac{\log(d)\log(1/\delta)}{14n} \,. \tag{91}$$

Now, assume that $\alpha < \log(1/\delta)/(7\sqrt{d})$. Since $\sum_{j=2}^{d} q_j = 1 - q_1 = \alpha d/n$, there exists $2 \leq j \leq d$ such that

$$q_j \leqslant \frac{\alpha d}{n(d-1)} \leqslant \frac{2\alpha}{n} < \frac{2\log(1/\delta)}{7n\sqrt{d}}.$$

Let $P = P^{(j)}$, so that letting $E = \{X_1 = 1, \dots, X_n = 1\}$, we have $\mathbb{P}_{P^{(j)}}(E) = (\delta^{1/n})^n = \delta$. Under E, one has $\widehat{P}_n = Q$, thus denoting $\rho = 1 - \delta^{1/n}$ we have

$$\operatorname{KL}(P,\widehat{P}_n) = \operatorname{KL}(P^{(j)},Q) \ge D(\rho,q_j) = q_j h\left(\frac{\rho}{q_j}\right).$$

By convexity of the exponential function, one has $1 - e^{-x} \ge (1 - e^{-1})x$ for $x \in [0, 1]$; since $\frac{\log(1/\delta)}{n} \le 1$, this implies that $\rho = 1 - \exp\left(-\frac{\log(1/\delta)}{n}\right) \ge (1 - e^{-1})\frac{\log(1/\delta)}{n}$. We therefore have

$$\frac{\rho}{q_j} \ge \frac{(1-e^{-1})\log(\delta^{-1})/n}{2\log(\delta^{-1})/(7n\sqrt{d})} = \frac{7(1-e^{-1})\sqrt{d}}{2} \ge 2\sqrt{d} \ge e.$$

Hence, by Lemma 14 we have

$$\begin{aligned} \operatorname{KL}(P, \widehat{P}_n) &\ge q_j h\left(\frac{\rho}{q_j}\right) \ge q_j \times e^{-1} \frac{\rho}{q_j} \log\left(\frac{\rho}{q_j}\right) \ge e^{-1} (1 - e^{-1}) \frac{\log(1/\delta)}{n} \log(2\sqrt{d}) \\ &\ge \frac{\log(1/\delta) \log(\sqrt{d})}{5n} = \frac{\log(d) \log(1/\delta)}{10n} \,. \end{aligned}$$

This concludes the proof of Lemma 1.

Proof of Theorem 4. We again use the consequence (25) of Proposition 1, this time combined with Lemma 1. Specifically, since $d \ge 5000 \ge 3000 \log(\frac{3}{1-e^{-1}})$, it follows from (25) that there exists $P \in \mathcal{P}_d$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \ge \frac{d}{4600n}\right) \ge 1 - 3\exp\left(-\frac{d}{3000}\right) \ge e^{-1} \ge \delta.$$

On the other hand, by Lemma 1, there exists $P \in \mathcal{P}_d$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \geqslant \frac{\log(d)\log(1/\delta)}{14\,n}\right) \geqslant \delta\,.$$

As before, taking the best of these two lower bounds shows that there exists $P \in \mathcal{P}_d$ under which, with probability at least δ ,

$$\mathrm{KL}(P, \widehat{P}_n) \ge \frac{99}{100} \cdot \frac{d}{4600n} + \frac{1}{100} \cdot \frac{\log(d)\log(1/\delta)}{14n} \ge \frac{d + \log(d)\log(1/\delta)}{5000n}.$$

7.3 Proof of Proposition 1

In this section, we turn to the proof of the high-probability lower bound of Proposition 1.

Note that if $s \leq 35$, then $1 - 3e^{-s/35} < 0$ and the inequality is trivial. From now on, we therefore assume that $s \ge 36$.

Random support. Fix n, d, s and an estimator $\hat{P}_n = \Phi(X_1, \ldots, X_n)$. Let \mathcal{S} denote the class of subsets σ of $\{2, \ldots, d\}$ with cardinality $|\sigma| = s - 1$. For any $\sigma \in \mathcal{S}$, we let $P_{\sigma} = (p_1^{\sigma}, \ldots, p_d^{\sigma})$ be the element of $\mathcal{P}_{s,d}$ defined by

$$P_{\sigma} = \left(1 - \frac{s-1}{2en}\right)\delta_1 + \frac{1}{2en}\sum_{j\in\sigma}\delta_j.$$
(92)

In addition, we let π denote the uniform distribution on S, which induces a "prior" distribution on $\mathcal{P}_{s,d}$. We can then define a joint distribution for $(\sigma, (X_1, \ldots, X_n))$ on $S \times [d]^n$ as follows: $\sigma \sim \pi$, and conditionally on σ , the variables X_1, \ldots, X_n are i.i.d. with distribution P_{σ} . We denote by πP the marginal distribution of X_1, \ldots, X_n under this distribution. As before, for $1 \leq j \leq d$ we denote by $N_j = \sum_{i=1}^n \mathbf{1}(X_i = j)$ the number of occurrences of the class j.

We start by establishing an upper bound on the number $D_n = \sum_{j=1}^d \mathbf{1}(N_j \ge 1)$ of distinct classes. For any $\sigma \in \mathcal{S}$, using that $\mathbb{E}[D_n|\sigma] = \mathbb{E}_{P_{\sigma}}[D_n] \le 1 + (s-1)n/(2en) = 1 + (s-1)/(2e)$ (where we used that $\mathbb{P}_{P_{\sigma}}(N_j \ge 1) \le \sum_{i=1}^n \mathbb{P}_{P_{\sigma}}(X_i = j)$) and applying Lemma 19, we obtain that

$$\mathbb{P}\left(D_n \ge e + \frac{s-1}{2}\right) = \mathbb{P}\left(D_n \ge e \cdot \left[1 + \frac{s-1}{2e}\right]\right) \le \exp\left(-1 - \frac{s-1}{2e}\right) \le \exp\left(-\frac{s}{2e}\right).$$

In what follows, we define the event $E = \{D_n < e + (s-1)/2, N_j \ge 1\}$, so that $\mathbb{P}(E^c) \le \exp(-s/(2e)) + (\frac{s-1}{2en})^n \le \exp(-s/(2e)) + (2e)^{-n}$. We also let $\hat{\sigma} = \{2 \le j \le d : N_j \ge 1\}$, so that $|\hat{\sigma}| = D_n - 1$ under E.

We now proceed to the lower bound on the estimation error. For any t > 0, we have

$$\mathbb{E}_{\sigma \sim \pi}[\mathbb{P}(\mathrm{KL}(P_{\sigma}, \widehat{P}_n) \ge t | \sigma)] = \mathbb{E}[\mathbb{P}(\mathrm{KL}(P_{\sigma}, \widehat{P}_n) \ge t | X_1, \dots, X_n)].$$
(93)

We thus aim at establishing a lower bound of the following form, for some constant C > 0:

$$\mathbb{P}\left(\mathrm{KL}(P_{\sigma}, \widehat{P}_n) \geqslant \frac{s \log(ed/s)}{Cn} \middle| X_1, \dots, X_n\right) \geqslant e^{-s/C}$$

Below, we reason conditionally on X_1, \ldots, X_n . First, note that the (posterior) conditional distribution of σ given X_1, \ldots, X_n , denoted $\hat{\pi}$, is the uniform distribution on the set

$$\widehat{\mathcal{S}} = \{ \sigma \in \mathcal{S} : \widehat{\sigma} \subset \sigma \} \,.$$

(Indeed, for any $\sigma \in S \setminus \widehat{S}$, the distribution $P_{\sigma}^{\otimes n}$ puts a mass of 0 to the sequence (X_1, \ldots, X_n) ; while all measures $P_{\sigma}^{\otimes n}$ with $\sigma \in \widehat{S}$ put the same positive mass to such a sequence.) In other words, $\sigma = \widehat{\sigma} \cup \widetilde{\sigma}$, where (conditionally on X_1, \ldots, X_n) $\widetilde{\sigma}$ is uniformly distributed over all subsets of $\{2, \ldots, d\} \setminus \widehat{\sigma}$ with $s - D_n$ elements. Write $\widehat{P}_n = (\widehat{p}_1, \ldots, \widehat{p}_d)$, and let $\widehat{\alpha} = n(1 - \widehat{p}_1)$. **Large** $\hat{\alpha}$. Assume first that $\hat{\alpha} \ge \sqrt{sd}/20$. In this case, applying Lemma 16 with $J = \{2, \ldots, d\}$ gives, for any $\sigma \in S$,

$$\mathrm{KL}(P_{\sigma}, \widehat{P}_n) \ge D\left(\sum_{j=2}^d p_j^{\sigma}, \sum_{j=2}^d \widehat{p}_j\right) = D\left(\frac{s-1}{2en}, \frac{\widehat{\alpha}}{n}\right) = \frac{\widehat{\alpha}}{n} \cdot h\left(\frac{s-1}{2e\widehat{\alpha}}\right).$$

Now, if $\widehat{\alpha} \ge \sqrt{sd}/20$, we have $\frac{s-1}{2e\widehat{\alpha}} \le 10e^{-1}\sqrt{s/d} \le 1/2$ as $d/s \ge 55$, thus $h(\frac{s-1}{2e\widehat{\alpha}}) \ge h(1/2) = (1-\log 2)/2 \ge 1/7$. The previous inequality then writes, for any $\sigma \in \widehat{S}$ (and thus with probability 1 over $\sigma \sim \widehat{\pi}$),

$$\operatorname{KL}(P_{\sigma}, \widehat{P}_{n}) \geqslant \frac{1}{7} \frac{\sqrt{sd}}{20n} = \frac{s}{140n} \sqrt{\frac{d}{s}} \geqslant \frac{s}{140n} \log\left(e\sqrt{\frac{d}{s}}\right) \geqslant \frac{s\log(ed/s)}{280n} \,. \tag{94}$$

Small $\hat{\alpha}$. Assume from now on that $\hat{\alpha} \leq \sqrt{sd}/20$. We start by noting that

$$\operatorname{KL}(P_{\sigma}, \widehat{P}_{n}) = \sum_{j=1}^{d} D(p_{j}^{\sigma}, \widehat{p}_{j}) \geqslant \sum_{2 \leqslant j \leqslant d, \, j \notin \widehat{\sigma}} D\left(p_{j}^{\sigma}, \widehat{p}_{j}\right) \geqslant \sum_{2 \leqslant j \leqslant d, \, j \notin \widehat{\sigma}} D\left(\frac{1}{2en}, \widehat{p}_{j}\right) \mathbf{1}(j \in \widetilde{\sigma}) \,. \tag{95}$$

Observe that in the right-hand side of (95), conditionally on X_1, \ldots, X_n , the only randomness comes from the presence of $\tilde{\sigma}$. Now since $\sum_{j=2}^{d} \hat{p}_j = \hat{\alpha}/n \leq \sqrt{sd}/(20n)$, we have:

$$\left|\left\{2 \leqslant j \leqslant d : \widehat{p}_j \geqslant \frac{\sqrt{s/d}}{10n}\right\}\right| = \left|\left\{2 \leqslant j \leqslant d : \widehat{p}_j \geqslant \frac{\sqrt{sd}/(20n)}{d/2}\right\}\right| \leqslant d/2$$

Recall that, under the event E, one has $|\hat{\sigma}| = D_n - 1 \leq e - 1 + (s - 1)/2$. It follows that, letting

$$\widehat{J} = \left\{ 2 \leqslant j \leqslant d : \widehat{p}_j < \frac{\sqrt{s/d}}{10n}, \ N_j = 0 \right\},\$$

we have, recalling that $d \ge 55s \ge 1980$,

$$\begin{aligned} |\widehat{J}| \geqslant (d-1) - \frac{d}{2} - |\widehat{\sigma}| \geqslant \frac{d}{2} - 1 - (e-1) - \frac{s-1}{2} &= \frac{d-s - (2e-1)}{2} \\ \geqslant \frac{d}{2} \left(1 - \frac{1}{55} - \frac{2e-1}{1980} \right) \geqslant 0.48d \,. \end{aligned}$$
(96)

In addition, since $D(\frac{1}{2en}, \cdot)$ decreases on $(0, \frac{1}{2en})$, we have for every $j \in \widehat{J}$:

$$D\left(\frac{1}{2en}, \hat{p}_j\right) \ge D\left(\frac{1}{2en}, \frac{\sqrt{s/d}}{10n}\right) = \frac{\sqrt{s/d}}{10n} \cdot h\left(5e^{-1}\sqrt{d/s}\right)$$
$$\ge \frac{\sqrt{s/d}}{10n} \cdot e^{-1}5e^{-1}\sqrt{d/s}\log\left(5e^{-1}\sqrt{d/s}\right) \ge \frac{1}{4e^2n}\log\left(\frac{ed}{s}\right)$$

where we used that $5e^{-1}\sqrt{d/s} \ge e$ as $d \ge 55s$, and that $h(t) \ge e^{-1}t \log t$ for $t \ge e$ (Lemma 14). Plugging this lower bound into (95), we obtain

$$\operatorname{KL}(P_{\sigma}, \widehat{P}_{n}) \geqslant \frac{1}{4e^{2}n} \log\left(\frac{ed}{s}\right) \sum_{j \in \widehat{J}} \mathbf{1}(j \in \widetilde{\sigma}) = \frac{1}{4e^{2}n} \log\left(\frac{ed}{s}\right) |\widehat{J} \cap \widetilde{\sigma}|.$$

$$(97)$$

We will now show that $|\widehat{J} \cap \widetilde{\sigma}| \gtrsim s$ with high probability over the draw of $\widetilde{\sigma}$, conditionally on X_1, \ldots, X_n . Although one could in principle show this by purely combinatorial means, we will instead resort to the notion of negative association [JDP83], which provides a particularly convenient way to handle the dependence that arises here.

Denote $\widehat{\sigma}^c = \{2, \ldots, d\} \setminus \widehat{\sigma}$, so that $|\widehat{\sigma}^c| = d - D_n$ under E. Since $\widetilde{\sigma}$ is uniformly distributed on subsets of $\widehat{\sigma}^c$ with $s - D_n$ elements, conditionally on X_1, \ldots, X_n the vector $(\mathbf{1}(j \in \widetilde{\sigma}))_{j \in \widehat{\sigma}^c}$ is uniformly distributed over all permutations of the vector $(\mathbf{1}(j \in \widetilde{\sigma}_0))_{j \in \widehat{\sigma}^c}$ for some fixed $\widetilde{\sigma}_0 \subset \widehat{\sigma}^c$ with $s - D_n$ elements. In other words, this distribution is a permutation distribution, which by [JDP83, Theorem 2.11] is negatively associated, in the sense of [JDP83, Definition 2.1]. By the restriction property of negatively associated random variables [JDP83, Property P4], this implies that the variables $(\mathbf{1}(j \in \widetilde{\sigma}))_{j \in \widehat{J}}$ are negatively associated. This implies in particular that, for every $\lambda \in \mathbb{R}^+$,

$$\mathbb{E}\Big[\exp\big(-\lambda|\widehat{J}\cap\widetilde{\sigma}|\big)\big|X_1,\ldots,X_n\Big] = \mathbb{E}\bigg[\prod_{j\in\widehat{J}}\exp\big(-\lambda\mathbf{1}(j\in\widetilde{\sigma})\big)\Big|X_1,\ldots,X_n\bigg]$$
$$\leqslant \prod_{j\in\widehat{J}}\mathbb{E}\big[\exp\big(-\lambda\mathbf{1}(j\in\widetilde{\sigma})\big)\big|X_1,\ldots,X_n\big] = \left(\frac{s-D_n}{d-D_n}e^{-\lambda} + \frac{d-s}{d-D_n}\right)^{|\widehat{J}|}.$$

Denoting $\widehat{\eta} = (s - D_n)/(d - D_n) \in [0, 1]$ and recalling that $|\widehat{J}| \ge 0.48d$ by (96), we deduce that, for every $\lambda \in \mathbb{R}^+$ (using that $1 - \widehat{\eta} + \widehat{\eta}e^{-\lambda} \in [0, 1]$),

$$\log \mathbb{E}\left[\exp\left(-\lambda|\widehat{J}\cap\widetilde{\sigma}|\right)|X_1,\ldots,X_n\right] \leqslant 0.48d\log\left(1-\widehat{\eta}+\widehat{\eta}e^{-\lambda}\right) \leqslant 0.48d\,d\widehat{\eta}\,(e^{-\lambda}-1)\,.$$

Now under E and recalling that $s \ge 36$, we have

$$d\,\hat{\eta} = d \cdot \frac{s - D_n}{d - D_n} \ge s - \frac{s + 2e - 1}{2} = \frac{s - (2e - 1)}{2} \ge 0.43s$$

so that $\log \mathbb{E}[\exp(-\lambda | \widehat{J} \cap \widetilde{\sigma} |)] \leq 0.2s(e^{-\lambda} - 1)$. By a standard argument [BLM13, p. 23], this implies the following tail bound: for every $s' \in (0, s/10)$,

$$\mathbb{P}(|\widehat{J} \cap \widetilde{\sigma}| \leq s' | X_1, \dots, X_n) \leq \exp(-D(s', s/5)).$$

In particular,

$$\mathbb{P}\left(\left|\widehat{J}\cap\widetilde{\sigma}\right|\leqslant\frac{s}{10}\,\middle|\,X_1,\ldots,X_n\right)\leqslant\exp\left(-\frac{1-\log 2}{10}s\right)\leqslant e^{-s/35}\,.$$

Plugging this into the lower bound (97) shows that, under the event E,

$$\mathbb{P}\left(\mathrm{KL}(P_{\sigma}, \widehat{P}_{n}) \geqslant \frac{s}{40e^{2}n} \log\left(\frac{ed}{s}\right) \middle| X_{1}, \dots, X_{n}\right) \geqslant 1 - e^{-s/35}.$$
(98)

Note that this lower bound also holds in the case where $\hat{\alpha} \ge \sqrt{sd}/20$ due to (94).

Conclusion of the proof. Using the identity (93) together with the conditional lower bound (98) under E, as well as the bound $\mathbb{P}(E^c) \leq e^{-s/2e} + (2e)^{-n} \leq 2e^{-s/2e}$ established above, we get

$$\mathbb{E}_{\sigma \sim \pi} \left[\mathbb{P} \left(\mathrm{KL}(P_{\sigma}, \widehat{P}_{n}) \geq \frac{s \log(ed/s)}{40e^{2}n} \middle| \sigma \right) \right]$$

$$\geq \mathbb{E} \left[\mathbb{P} \left(\mathrm{KL}(P_{\sigma}, \widehat{P}_{n}) \geq \frac{s \log(ed/s)}{40e^{2}n} \middle| X_{1}, \dots, X_{n} \right) \mathbf{1}_{E} \right]$$

$$\geq \mathbb{P}(E) \cdot \left(1 - e^{-s/35} \right) \geq \left(1 - 2e^{-s/2e} \right) \left(1 - e^{-s/35} \right)$$

$$\geq 1 - 3e^{-s/35}. \tag{99}$$

Since $\max_{\sigma \in \mathcal{S}} \{\cdots\} \ge \mathbb{E}_{\sigma \sim \pi}[\cdots]$, it follows from (99) that there exists $\sigma \in \mathcal{S}$ such that

$$\mathbb{P}_{P_{\sigma}}\left(\mathrm{KL}(P_{\sigma}, \widehat{P}_n) \geqslant \frac{s \log(ed/s)}{40e^2 n}\right) \geqslant 1 - 3e^{-s/35}$$

The lower bound of Proposition 1 follows by further bounding $40e^2 \leq 300$.

7.4 Proof of Corollary 1

Fix $\Phi = \Phi_{s,\delta}$. By Lemma 1, there exists $P \in \mathcal{F} \subset \mathcal{P}_{2,d} \subset \mathcal{P}_{s,d}$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P, \widehat{P}_n) \geqslant \frac{\log(d)\log(1/\delta)}{14\,n}\right) \geqslant \delta\,.$$
(100)

On the other hand, by Proposition 1, there exists a distribution $P \in \mathcal{P}_{s,d}$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \geqslant \frac{s\log(ed/s)}{300\,n}\right) \geqslant 1 - 3e^{-s/35}$$

Now, if $s \ge 42$, then $1 - 3e^{-s/35} > e^{-2} \ge \delta$; on the other hand, if $2 \le s \le 41$ then (using that $d \ge 55s \ge 41$)

$$\frac{\log(d)\log(1/\delta)}{14n} \geqslant \frac{3\log(d)}{14n} \geqslant \frac{2\cdot 41\log(ed/41)}{14\cdot 41n} \geqslant \frac{2\cdot s\log(ed/s)}{574n} \geqslant \frac{s\log(ed/s)}{300n}$$

Hence, using the previous inequalities, we deduce that regardless of $s \ge 2$, there exists $P \in \mathcal{P}_{s,d}$ such that

$$\mathbb{P}_P\left(\mathrm{KL}(P,\widehat{P}_n) \geqslant \frac{s\log(ed/s)}{300n}\right) \geqslant \delta.$$
(101)

Taking the best of the two lower bounds (100) and (101), we obtain that under some $P \in \mathcal{P}_{s,d}$, with probability at least δ one has

$$\begin{split} \mathrm{KL}(P,\widehat{P}_n) &\geqslant \max\left\{\frac{s\log(ed/s)}{300n}, \frac{\log(d)\log(1/\delta)}{14n}\right\} \\ &\geqslant \frac{20}{21}\frac{s\log(ed/s)}{300n} + \frac{1}{21}\frac{\log(d)\log(1/\delta)}{14n} \geqslant \frac{s\log(ed/s) + \log(d)\log(1/\delta)}{320\,n} \end{split}$$

8 Proof of Theorem 6

We now provide the proof of Theorem 6. First, note that if $\delta \leq e^{-n/6}$, then the right-hand side of (42) is greater than 1, hence the inequality holds. We now assume that $\delta \in (e^{-n/6}, 1)$.

Poisson sampling. As before, we let $(X_i)_{i\geq 1}$ denote an i.i.d. sequence from P, and N be an independent random variable with distribution $\mathsf{Poisson}(n/2)$. In addition, we let $\widetilde{N}_j = \sum_{1 \leq i \leq N} \mathbf{1}(X_i = j)$. We work under the event $E = \{N \leq n\}$, such that $\mathbb{P}(E) \geq 1 - e^{-n/6} \geq 1 - \delta$ (inequality (74)), and under which $\widetilde{N}_j \leq N_j$ for $1 \leq j \leq d$. Thus, letting $\lambda_j = np_j$ we have under E that

$$U_n \leqslant \sum_{j: p_j < 1/n} p_j + \sum_{j: p_j \ge 1/n} p_j \mathbf{1} \left(N_j \leqslant \frac{np_j}{4} \right) \leqslant \sum_{j: p_j < 1/n} p_j + \frac{1}{n} \sum_{j: \lambda_j \ge 1} \lambda_j \mathbf{1} \left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4} \right).$$
(102)

Also, recall that $\widetilde{N}_1, \ldots, \widetilde{N}_d$ are independent, with $\widetilde{N}_j \sim \mathsf{Poisson}(\lambda_j/2)$ for $j = 1, \ldots, d$. We therefore need to control the last term in (102).

Multi-scale decomposition and domination. We collect terms in the sum as follows:

$$\sum_{j:\ \lambda_j \ge 1} \lambda_j \mathbf{1} \left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4} \right) = \sum_{k \ge 0} \sum_{j:\ 2^k \leqslant \lambda_j < 2^{k+1}} \lambda_j \mathbf{1} \left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4} \right)$$
$$\leqslant \sum_{k \ge 0} 2^{k+1} \sum_{j:\ 2^k \leqslant \lambda_j < 2^{k+1}} \mathbf{1} \left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4} \right).$$

Now since $\mathbb{P}(\widetilde{N}_j \leq \lambda_j/4) \leq e^{-\lambda_j/14}$, the (Bernoulli) random variable $\mathbf{1}(\widetilde{N}_j \leq \lambda_j/4)$ is stochastically dominated by $P'_j \sim \mathsf{Poisson}(e^{-\lambda_j/14})$. By choosing $(P'_j)_{1 \leq j \leq d}$ to be independent, we deduce from Lemma 18 that, for every $k \in \mathbb{N}$,

$$\sum_{j: 2^k \leqslant \lambda_j < 2^{k+1}} \mathbf{1} \Big(\widetilde{N}_j \leqslant \frac{\lambda_j}{4} \Big) \preccurlyeq P_k'' = \sum_{j: 2^k \leqslant \lambda_j < 2^{k+1}} P_j'.$$

Note that the random variables $(\widetilde{P}''_k)_{k\geq 0}$ are independent. In addition, P''_k follows a Poisson distribution with parameter $\sum_{j: 2^k \leqslant \lambda_j < 2^{k+1}} e^{-\lambda_j/14} \leqslant d'_k e^{-2^k/14}$ where $d'_k = |\{1 \leqslant j \leqslant d : 2^k \leqslant j < 2^{k+1}\}|$. Let $(P_k)_{k\geq 0}$ be independent random variables, with $P_k \sim \mathsf{Poisson}(d'_k e^{-2^k/14})$. The previous inequalities imply that

$$\sum_{j:\ \lambda_j \ge 1} \lambda_j \mathbf{1} \left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4} \right) \preccurlyeq \sum_{k \ge 0} 2^{k+1} P_k \,. \tag{103}$$

Control for individual scales. The following key lemma controls the tails of individual scales in the sum (103).

Lemma 13. For every $k \in \mathbb{N}$ and $t \ge 0$, one has

$$\mathbb{P}\left(2^k P_k \geqslant 56d'_k e^{-2^k/56} + 28et\right) \leqslant e^{-t}.$$
(104)

Proof. Since $P_k \sim \mathsf{Poisson}(d'_k e^{-2^k/14})$, the Poisson tail bound (Lemma 17) implies that for any $u \ge d'_k e^{-2^k/28}$,

$$\mathbb{P}(P_k \ge eu) \le \exp\left\{-D(eu, d'_k e^{-2^k/14})\right\}$$

Now since $eu/(d'_k e^{-2^k/14}) \ge e(d'_k e^{-2^k/28})/(d'_k e^{-2^k/14}) \ge e \cdot e^{2^k/28}$, and since $D(p,q) \ge e^{-1}p\log(p/q)$ when $p \ge eq$ (Lemma 15), the previous bound implies that

$$\mathbb{P}(P_k \ge eu) \le \exp\left\{-e^{-1} \cdot eu \log\left(e \cdot e^{2^k/28}\right)\right\} \le \exp\left\{-\frac{2^k u}{28}\right\}.$$

Hence, letting $u = 4t/2^k$, the inequality

$$\mathbb{P}(2^k P_k \ge 4et) = \mathbb{P}(P_k \ge eu) \le \exp\left\{-\frac{2^k u}{28}\right\} \le e^{-t/7}$$
(105)

holds for every $t \ge 2^k d'_k e^{-2^k/28}/4$. Since

$$\frac{2^k \cdot d'_k e^{-2^k/28}}{4} \leqslant \frac{56e^{-1}e^{2^k/56} \cdot d'_k e^{-2^k/28}}{4} = 14e^{-1}d'_k e^{-2^k/56} = t_0 \,,$$

the bound (105) holds for any $t \ge t_0$. Thus, for any $t \ge 0$,

$$\mathbb{P}(2^k P_k \ge 4e(t_0 + 7t)) \le e^{-(t_0 + 7t)/7} \le e^{-t},$$

which is precisely the claimed inequality (104).

Lemma 13 states that $2^k P_k$ is stochastically dominated by $56d'_k e^{-2^k/56} + 28eE_k$, where $E_k \sim \text{Exp}(1)$ is an exponential random variable. In addition, since $\lambda_j = np_j \leq n$ for $j = 1, \ldots, d$, we have $d'_k = 0$ (and thus $P_k = 0$) for $k > \log_2 n$. Thus, letting $(E_k)_{k \in \mathbb{N}}$ be independent exponential random variables, inequality (103) implies that

$$\sum_{j:\ \lambda_j \ge 1} \lambda_j \mathbf{1} \left(\widetilde{N}_j \leqslant \frac{\lambda_j}{4} \right) \preccurlyeq 112 \sum_{0 \leqslant k \leqslant \log_2 n} d'_k e^{-2^k/56} + 56e \sum_{0 \leqslant k \leqslant \log_2 n} E_k \,. \tag{106}$$

In addition, the first term in the r.h.s. of (106) is controlled by $cs_{n/112}^{\circ}$ for some constant c (see below). The second term may also be controlled using a deviation bound for sums of exponential random variables (that is, gamma random variables) [BLM13, p. 28], which gives, with probability $1 - \delta$,

$$\sum_{0 \leqslant k \leqslant \log_2 n} E_k \lesssim \log n + \log(1/\delta) \,.$$

While already non-trivial and in fact near-optimal, the resulting bound on the missing mass features an additional $\log n$ term, which is suboptimal in some (rather extreme) cases.

In order to address this sub-optimality, we need to account more carefully for the contribution of certain scales to the sum $\sum_{k\geq 0} 2^k P_k$.

Accounting for "rarely contributing" scales. We now define two types of "scale indices" $k \in \mathbb{N}$. Specifically, we let

$$A = \left\{ k \in \mathbb{N} : d'_k e^{-2^k/56} \ge e \right\}.$$
(107)

Intuitively speaking, since $\mathbb{E}[P_k] = d'_k e^{-2^k/14}$, the scale indices $k \in A$ are those for which P_k is often larger than 1, i.e. non-zero—with the technical caveat that we have changed the exponent in (107). On the contrary, indices $k \in \mathbb{N} \setminus A$ are those for which typically $P_k = 0$. However, given that there may be several such terms and that we account for low-probability events, some of these terms may be positive with small probability. Hence, they may contribute to the sum, especially if their coefficient 2^k is large. Thus, they should also be accounted for.

In what follows, we separately account for the contribution of indices $k \in A$ and $k \notin A$. First recall that, by Lemma 13, for any $k \in \mathbb{N}$ one has

$$2^k P_k \preccurlyeq 56d'_k e^{-2^k/56} + 28eE_k \tag{108}$$

where $(E_k)_{k \in \mathbb{N}}$ are i.i.d. exponential random variables.

We also recall (see [BLM13, p. 28]) that for any $\delta \in (0, 1)$ and finite subset $B \subset \mathbb{N}$, with probability $1 - \delta$ one has

$$\sum_{k \in B} E_k < |B| + \sqrt{2|B|\log(1/\delta)} + \log(1/\delta) \le 2|B| + 2\log(1/\delta).$$
(109)

When $k \in A$, one has $d'_k e^{-2^k/56} \ge e$, hence (108) gives

$$2^{k}P_{k} \preccurlyeq 56d'_{k}e^{-2^{k}/56} + 56e + 28e(E_{k} - 2) \leqslant 112d'_{k}e^{-2^{k}/56} + 28e(E_{k} - 2).$$
(110)

Now applying the tail bound (109) to B = A (which is finite since $|A| \leq \log_2 n$), we obtain with probability $1 - \delta$,

$$\sum_{k \in A} (E_k - 2) < 2|A| + 2\log(1/\delta) - 2|A| = 2\log(1/\delta)$$

which combined with (110) (and Lemma 18) gives

$$\mathbb{P}\bigg(\sum_{k\in A} 2^k P_k \ge 112 \sum_{k\in A} d'_k e^{-2^k/56} + 56e \log(1/\delta)\bigg) \le \delta.$$
(111)

We now turn to the case $k \notin A$, in which case $d'_k e^{-2^k/56} < e$. Using that if $P \sim \mathsf{Poisson}(\lambda)$ then $\mathbb{P}(P \neq 0) = 1 - e^{-\lambda} \leq \lambda$, this implies that

$$\mathbb{P}(P_k \neq 0) \leqslant d'_k e^{-2^k/14} = (d'_k e^{-2^k/56}) e^{-2^k 3/56} \leqslant e \cdot e^{-2^k/19}.$$

Now, let $k^* = \lceil 19 \log(1/\delta) \rceil \leq 20 \log(1/\delta)$ (as $\delta \leq e^{-1}$), so that $e^{-2^{k^*}/19} \leq \delta$. Using the previous inequality, we may bound

$$\mathbb{P}\left(\sum_{k\notin A,\,k\geqslant k^*} 2^k P_k \neq 0\right) \leqslant \sum_{k\notin A,\,k\geqslant k^*} \mathbb{P}(P_k \neq 0) \leqslant \sum_{k\geqslant k^*} e \cdot e^{-2^k/19}$$
$$= e \sum_{l\geqslant 0} e^{-2^l 2^{k^*}/19} \leqslant e \sum_{l\geqslant 0} \left(e^{-2^{k^*}/19}\right)^{l+1}$$
$$\leqslant e \sum_{l\geqslant 0} \delta^{l+1} = \frac{e\delta}{1-\delta} \leqslant \frac{e}{1-e^{-1}}\delta,$$

so that

$$\mathbb{P}\left(\sum_{k\notin A,\,k\geqslant k^*} 2^k P_k \neq 0\right) \leqslant 5\delta\,.$$
(112)

Finally, it remains to control indices $0 \leq k < k^*$ such that $k \notin A$. For this, we combine the domination (108) with the tail bound (109) to obtain:

$$\mathbb{P}\bigg(\sum_{k\notin A,\,k< k^*} 2^k P_k \geqslant 56 \sum_{k\notin A,\,k< k^*} d'_k e^{-2^k/56} + 56e\{k^* + \log(1/\delta)\}\bigg) \leqslant \delta.$$

Since $k^* \leq 20 \log(1/\delta)$, we conclude that

$$\mathbb{P}\bigg(\sum_{k\notin A,\,k< k^*} 2^k P_k \ge 56 \sum_{k< k^*,\,k\notin A} d'_k e^{-2^k/56} + 1176e \log(1/\delta)\bigg) \le 7\delta.$$
(113)

Combining inequalities (111), (112) and (113) through a union bound to control the sum of the three terms, we obtain

$$\mathbb{P}\bigg(\sum_{k\geq 0} 2^k P_k \geq 168 \sum_{k\geq 0} d'_k e^{-2^k/56} + 1232e \log(1/\delta)\bigg) \leqslant 7\delta.$$

Conclusion. Together with inequality (103), the previous bound implies that

$$\mathbb{P}\bigg(\sum_{j:\,\lambda_j \ge 1} \lambda_j \mathbf{1}\Big(\widetilde{N}_j \leqslant \frac{\lambda_j}{4}\Big) \ge 336 \sum_{k \ge 0} d'_k e^{-2^k/56} + 2464e \log(1/\delta)\bigg) \leqslant 7\delta.$$

Now, by definition of d'_k one has

$$\sum_{k \ge 0} d'_k e^{-2^k/56} = \sum_{k \ge 0} \sum_{j: \ 2^k \le np_j < 2^{k+1}} e^{-2^k/56}$$
$$\leqslant \sum_{k \ge 0} \sum_{j: \ 2^k \le np_j < 2^{k+1}} e^{-np_j/112} = \sum_{j: \ p_j \ge 1/n} e^{-np_j/112}$$

Plugging the previous inequalities into the decomposition (102) (which holds except for an event of probability at most δ), we get that with probability at least $1 - 8\delta$,

$$U_n < \frac{1}{n} \left[\sum_{j: p_j < 1/n} (np_j) + 336 \sum_{j: p_j \ge 1/n} e^{-np_j/112} + 2464e \log(1/\delta) \right]$$

$$\leq \frac{336 s_{n/112}^\circ(P) + 2500e \log(1/\delta)}{n}.$$

9 Proof of Proposition 2

Let X_{n+1} be a new observation from P, independent from X_1, \ldots, X_n . For $j = 1, \ldots, d$, denote by $\widetilde{N}_j = \sum_{i=1}^{n+1} \mathbf{1}(X_i = j)$ the count of class j and by $\widetilde{D} = \sum_{j=1}^{d} \mathbf{1}(\widetilde{N}_j \ge 1)$ the number of distinct classes in the extended sample (X_1, \ldots, X_{n+1}) . Also, let $\widetilde{N} = (\widetilde{N}_j)_{1 \le j \le d}$. Since the joint distribution of (X_1, \ldots, X_{n+1}) is exchangeable (and permuting indices does not change \widetilde{N}), for every $j = 1, \ldots, d$ and $i = 1, \ldots, n$ one has $\mathbb{P}(X_{n+1} = j | \widetilde{N}) = \mathbb{P}(X_i = j | \widetilde{N})$; hence,

$$\mathbb{P}(X_{n+1} = j|\widetilde{N}) = \frac{1}{n+1} \sum_{i=1}^{n} \mathbb{P}(X_i = j|\widetilde{N}) = \frac{1}{n+1} \mathbb{E}\left[\sum_{i=1}^{n+1} \mathbf{1}(X_i = j) \middle| \widetilde{N}\right]$$
$$= \frac{\mathbb{E}[\widetilde{N}_j|\widetilde{N}]}{n+1} = \frac{\widetilde{N}_j}{n+1}.$$
(114)

Now, observe that for any probability distribution $Q \in \mathcal{P}_d$, one has $\operatorname{KL}(P,Q) = L(Q) - L(P)$, where $L(Q) = \mathbb{E}[\ell(Q,X)]$ with $X \sim P$, and where ℓ stands for the logarithmic loss $\ell(Q,x) = -\log Q(\{x\})$. In addition, since $\widehat{P}_n^{\operatorname{ad}}$ is independent of X_{n+1} one has

$$\mathbb{E}[\ell(\widehat{P}_n^{\mathrm{ad}}, X_{n+1})] = \mathbb{E}\left[\mathbb{E}[\ell(\widehat{P}_n^{\mathrm{ad}}, X_{n+1})|\widehat{P}_n^{\mathrm{ad}}]\right] = \mathbb{E}[L(\widehat{P}_n^{\mathrm{ad}})]$$

On the other hand, let $\tilde{P} = (\tilde{p}_j)_{1 \leq j \leq d}$ denote the maximum likelihood distribution on the extended sample (X_1, \ldots, X_{n+1}) , defined by

$$\widetilde{P} = \operatorname*{arg\,min}_{P' \in \mathcal{P}_d} \left\{ \frac{1}{n+1} \sum_{i=1}^{n+1} \ell(P', X_i) \right\}.$$

It is a standard fact that \tilde{P} is the empirical distribution, namely $\tilde{p}_j = \tilde{N}_j/(n+1)$. In addition, by definition of \tilde{P} , one has

$$L(P) = \mathbb{E}\left[\frac{1}{n+1}\sum_{i=1}^{n+1}\ell(P,X_i)\right] \ge \mathbb{E}\left[\frac{1}{n+1}\sum_{i=1}^{n+1}\ell(\widetilde{P},X_i)\right] = \mathbb{E}\left[\ell(\widetilde{P},X_{n+1})\right],$$

where the last step used the fact that the distribution of the vector (X_1, \ldots, X_{n+1}) is invariant under permutation, and that \tilde{P} is also invariant under permutation. Putting the previous inequalities together gives:

$$\mathbb{E}\big[\mathrm{KL}(P,\widehat{P}_n^{\mathrm{ad}})\big] = \mathbb{E}\big[L(\widehat{P}_n^{\mathrm{ad}})\big] - L(P) \leqslant \mathbb{E}\big[\ell(\widehat{P}_n^{\mathrm{ad}}, X_{n+1}) - \ell(\widetilde{P}, X_{n+1})\big].$$
(115)

Recall that $\hat{p}_j = (N_j + D_n/d)/(n + D_n)$ while $\tilde{p}_j = \tilde{N}_j/(n + 1)$. It then follows from (115) that

$$\mathbb{E}\left[\mathrm{KL}(P,\widehat{P}_{n}^{\mathrm{ad}})\right] \leq \mathbb{E}\left[\sum_{j=1}^{d} \left\{\ell(\widehat{P}_{n}^{\mathrm{ad}},X_{n+1}) - \ell(\widetilde{P},X_{n+1})\right\} \mathbf{1}(X_{n+1}=j)\right]$$

$$= \mathbb{E}\left[\sum_{j=1}^{d} \log\left(\frac{\widetilde{N}_{j}/(n+1)}{(N_{j}+D_{n}/d)/(n+D_{n})}\right) \mathbf{1}(X_{n+1}=j)\right]$$

$$= \mathbb{E}\left[\sum_{j=1}^{d} \log\left(\frac{\widetilde{N}_{j}}{N_{j}+D_{n}/d}\right) \mathbf{1}(X_{n+1}=j) + \log\left(\frac{n+D_{n}}{n+1}\right)\right]$$

$$\leq \mathbb{E}\left[\sum_{j=1}^{d} \mathbb{E}\left[\log\left(\frac{\widetilde{N}_{j}}{\widetilde{N}_{j}-1+D_{n}/d}\right) \mathbf{1}(X_{n+1}=j)\right] \widetilde{N}\right] + \frac{\mathbb{E}[D_{n}]-1}{n+1}, \quad (116)$$

where in the last inequality we used that if $X_{n+1} = j$, then $N_j = \widetilde{N}_j - 1$, and that $\log[(n + i)]$ $D_n)/(n+1)] = \log[1 + (D_n - 1)/(n+1)] \leq (D_n - 1)/(n+1).$ Consider first the case where $\widetilde{N}_j = 1$. In this case and if $X_{n+1} = j$, then $D_n = \widetilde{D} - 1$ (since

 $X_{n+1} = j$ does not appear in the first *n* observations) and hence:

$$\mathbb{E}\left[\log\left(\frac{\widetilde{N}_{j}}{\widetilde{N}_{j}-1+D_{n}/d}\right)\mathbf{1}(X_{n+1}=j)\Big|\widetilde{N}\right] = \mathbb{E}\left[\log\left(\frac{1}{(\widetilde{D}-1)/d}\right)\mathbf{1}(X_{n+1}=j)\Big|\widetilde{N}\right]$$
$$= \log\left(\frac{d}{\widetilde{D}-1}\right) \cdot \mathbb{P}(X_{n+1}=j|\widetilde{N})$$
$$= \frac{1}{n+1}\log\left(\frac{d}{\widetilde{D}-1}\right)$$
(117)

where we used (114) (and $\widetilde{N}_j = 1$) in the last equality.

Consider now the case $\widetilde{N}_j \ge 2$. We then bound

$$\mathbb{E}\left[\log\left(\frac{\widetilde{N}_{j}}{\widetilde{N}_{j}-1+D_{n}/d}\right)\mathbf{1}(X_{n+1}=j)\middle|\widetilde{N}\right] \leqslant \mathbb{E}\left[\log\left(\frac{\widetilde{N}_{j}}{\widetilde{N}_{j}-1}\right)\mathbf{1}(X_{n+1}=j)\middle|\widetilde{N}\right] \\
= \log\left(\frac{\widetilde{N}_{j}}{\widetilde{N}_{j}-1}\right) \cdot \mathbb{P}(X_{n+1}=j|\widetilde{N}) \\
= \frac{\widetilde{N}_{j}}{n+1} \cdot \log\left(\frac{\widetilde{N}_{j}}{\widetilde{N}_{j}-1}\right) \\
\leqslant \frac{\log 4}{n+1};$$
(118)

in the last inequality, we used that the function $\phi: t \mapsto t \log[t/(t-1)]$, such that

$$\phi'(t) = \log[t/(t-1)] + t[1/t - 1/(t-1)] = \log[1 + 1/(t-1)] - 1/(t-1) \le 0$$

for t > 1, is decreasing, so that $\phi(\widetilde{N}_j) \leq \phi(2) = 2 \log 2 = \log 4$. Denote now by $\widetilde{C}_1 = \sum_{j=1}^d \mathbf{1}(\widetilde{N}_j = 1)$ and $\widetilde{C}_2 = \sum_{j=1}^d \mathbf{1}(\widetilde{N}_j \geq 2)$, so that $\widetilde{D} = \widetilde{C}_1 + \widetilde{C}_2$. Plugging the bounds (117) and (118) into inequality (116), we obtain (using that if $\tilde{N}_j = 0$, then

$$\begin{split} \mathbf{1}(X_{n+1} = j) &= 0): \\ \mathbb{E}\big[\mathrm{KL}(P, \widehat{P}_n^{\mathrm{ad}})\big] \\ &\leqslant \mathbb{E}\Big[\sum_{j=1}^d \mathbb{E}\bigg[\log\bigg(\frac{\widetilde{N}_j}{\widetilde{N}_j - 1 + D_n/d}\bigg)\mathbf{1}(X_{n+1} = j)\Big|\widetilde{N}\bigg]\big\{\mathbf{1}(\widetilde{N}_j = 1) + \mathbf{1}(\widetilde{N}_j \geqslant 2)\big\}\bigg] + \frac{\mathbb{E}[D_n] - 1}{n+1} \\ &\leqslant \mathbb{E}\bigg[\sum_{j=1}^d \frac{1}{n+1}\log\bigg(\frac{d}{\widetilde{D}-1}\bigg)\mathbf{1}(\widetilde{N}_j = 1) + \sum_{j=1}^d \frac{\log 4}{n+1}\mathbf{1}(\widetilde{N}_j \geqslant 2)\bigg] + \frac{\mathbb{E}[D_n] - 1}{n+1} \\ &\leqslant \frac{1}{n+1}\mathbb{E}\bigg[\widetilde{C}_1\log\bigg(\frac{d}{\widetilde{D}-1}\bigg) + \widetilde{C}_2\log 4 + D_n - 1\bigg] \\ &\leqslant \frac{1}{n+1}\mathbb{E}\bigg[\widetilde{C}_1\log\bigg(\frac{ed}{\widetilde{C}_1}\bigg) + \widetilde{C}_2\log 4 + D_n - 1\bigg]. \end{split}$$

(In the last bound, we have used that if $\widetilde{D} = 1$, then only one class appears n+1 > 1 times, hence $\widetilde{C}_1 = 0$ and the first term vanishes; on the other hand, if $\widetilde{D} \ge 2$, then $\widetilde{D} - 1 \ge \widetilde{D}/e \ge \widetilde{C}_1/e$.)

By concavity of the map $x \mapsto -x \log x$ on \mathbb{R}^*_+ , we deduce that

$$\mathbb{E}_P\left[\mathrm{KL}(P,\widehat{P}_n^{\mathrm{ad}})\right] \leqslant \frac{1}{n+1} \left\{ \mathbb{E}_P[\widetilde{C}_1] \log\left(\frac{ed}{\mathbb{E}_P[\widetilde{C}_1]}\right) + \mathbb{E}_P[\widetilde{C}_2] \log 4 + \mathbb{E}_P[D_n] - 1 \right\}.$$
(119)

Now, note that (using Lemma 2 for the last inequality)

$$\mathbb{E}_{P}[\widetilde{C}_{1}] = \sum_{j=1}^{d} \mathbb{P}(\widetilde{N}_{j} = 1) = \sum_{j=1}^{d} np_{j}(1-p_{j})^{n} \leqslant \sum_{j=1}^{d} np_{j}e^{-np_{j}} = s_{n}^{\bullet}(P) \leqslant 2s_{n/2}^{\circ}(P),$$

and in addition $\mathbb{E}_P[\widetilde{C}_2] \leq \mathbb{E}_P[D_n] \leq s_n = s_n(P)$. Thus

$$\mathbb{E}_{P}[\widetilde{C}_{1}] \log\left(\frac{ed}{\mathbb{E}_{P}[\widetilde{C}_{1}]}\right) + \mathbb{E}_{P}[\widetilde{C}_{2}] \log 4 + \mathbb{E}_{P}[D_{n}]$$

$$\leq 2s_{n/2}^{\circ} \log\left(\frac{ed}{s_{n/2}^{\circ}}\right) + \mathbb{E}_{P}[\widetilde{C}_{1} + \widetilde{C}_{2}] \log 4 + \mathbb{E}_{P}[D_{n}]$$

$$\leq 2s_{n/2}^{\circ} \log\left(\frac{ed}{s_{n/2}^{\circ}}\right) + (1 + \log 4)s_{n}.$$

Plugging this inequality into (119), and using that the map $x \mapsto x \log(ed/x)$ is increasing over [0, d], leads to the desired bound (32).

10 Technical lemmata

In this section, we gather simple technical results that are used at various places in the proofs.

Lemma 14. Define the function $h : \mathbb{R}^+ \to \mathbb{R}^+$ by $h(t) = t \log t - t + 1$ for t > 0, and h(0) = 1. Then h is continuous and convex on \mathbb{R}^+ . In addition, $h(t) \leq t \log t$ for $t \geq 1$, while $h(t) \geq e^{-1}t \log t$ for $t \geq e$. Finally, $h(t) \leq (t-1)^2$ for any $t \in \mathbb{R}^+$.

Proof. Continuity of h is straightforward, while convexity comes from the fact that h''(t) = 1/t > 0 for t > 0. The inequality $h(t) \leq t \log t$ for $t \geq 1$ is immediate. We now turn to the lower

bound $h(t) \ge e^{-1}t \log t$ for $t \ge e$. By convexity of the map $t \mapsto t \log t$ over \mathbb{R}^*_+ , the quantity $t \log(t)/(t-1)$ increases in t, thus for $t \ge e$ one has $t \log(t)/(t-1) \ge e/(e-1)$, hence

$$h(t) = t \log t - (t-1) \ge \left(1 - \frac{e-1}{e}\right) t \log t = e^{-1} t \log t.$$

We conclude with the proof of the bound $h(t) \leq (t-1)^2$. To this end, let $f(t) = h(t)/(t-1)^2$ for $t \in \mathbb{R}^+ \setminus \{1\}$, and f(1) = 1/2. It is easily verified that f is continuous on \mathbb{R}^+ and differentiable on $\mathbb{R}^+ \setminus \{1\}$. In addition, $f'(t) = g(t)/(t-1)^3$ where $g(t) = 2(t-1) - (t+1)\log t$ for t > 0. One has $g'(t) = -(t^{-1} - 1 - \log(t^{-1})) \leq 0$, and since g(1) = 0, this implies that $g \leq 0$ on (0, 1] while $g \geq 0$ on $[1, +\infty)$. Therefore $f'(t) = g(t)/(t-1)^3 \leq 0$ for any $t \in \mathbb{R}^+ \setminus \{1\}$ and thus f is decreasing on \mathbb{R}^+ . We conclude by noting that $\lim_{t\to 0^+} f(t) = 1$, so that $f \leq 1$ on \mathbb{R}^+_+ .

Lemma 15. For any $p \in \mathbb{R}^*_+$, the function $q \mapsto D(p,q)$ is strictly convex on \mathbb{R}^*_+ , and reaches its minimum (equal to 0) at q = p. Hence, it is decreasing on (0,p] and increasing on $[p, +\infty)$. In addition, $D(p,q) \leq p \log(p/q)$ when $q \leq p$, and $D(p,q) \geq e^{-1}p \log(p/q)$ when $q \leq p/e$.

Proof. The claims on convexity and monotonicity follow from the fact that $\frac{\partial D}{\partial q}(p,q) = -\frac{p}{q} + 1$ (which cancels out at q = p) and $\frac{\partial D}{\partial q}(p,q) = \frac{p}{q^2} > 0$. The inequalities on D follow from the expression D(p,q) = qh(p/q) and from Lemma 14.

Lemma 16. Let $P = (p_1, \ldots, p_d) \in \mathcal{P}_d$ and $Q = (q_1, \ldots, q_d) \in \mathcal{P}_d$. For any subset $J \subset [d]$, we have

$$\operatorname{KL}(P,Q) \ge D\left(\sum_{j\in J} p_j, \sum_{j\in J} q_j\right).$$
(120)

Proof. We may assume, without loss of generality, that $q_j > 0$ for any $j \in J$. Indeed, either there exists a $j \in J$ such that $q_j = 0$ and $p_j > 0$, in which case the left-hand side of (120) is $+\infty$ and the inequality holds; or, for any $j \in J$ such that $q_j = 0$, one has $p_j = 0$: but in this case replacing J by $J' = \{j \in J : q_j > 0\}$ does not affect the right-hand side of (120).

Next, if $J = \emptyset$, the right-hand side of (120) is 0, thus the inequality holds. We now also assume that $J \neq \emptyset$. Then, by non-negativity and convexity of the function h (Lemma 14), we have

$$\operatorname{KL}(P,Q) = \sum_{j=1}^{d} D(p_j,q_j) \ge \sum_{j \in J} q_j h\left(\frac{p_j}{q_j}\right) = \left(\sum_{j \in J} q_j\right) \sum_{j \in J} \frac{q_j}{\sum_{k \in J} q_k} h\left(\frac{p_j}{q_j}\right)$$
$$\ge \left(\sum_{j \in J} q_j\right) h\left(\sum_{j \in J} \frac{q_j}{\sum_{k \in J} q_k} \cdot \frac{p_j}{q_j}\right) = D\left(\sum_{j \in J} p_j, \sum_{j \in J} q_j\right).$$

We also recall the following standard Poisson tail bound (e.g., [BLM13, p. 23]):

Lemma 17. Let $\lambda \in \mathbb{R}^+$ and $N \sim \text{Poisson}(\lambda)$. For any $\mu \in \mathbb{R}^+$ such that $\mu \ge \lambda$, one has

$$\mathbb{P}(N \ge \mu) \le \exp(-D(\mu, \lambda)).$$
(121)

In addition, for any $\mu \in \mathbb{R}^+$ such that $\mu \leq \lambda$, one has

$$\mathbb{P}(N \leqslant \mu) \leqslant \exp(-D(\mu, \lambda)).$$
(122)

Definition 4. Let X, Y be real random variables. We say that X is stochastically dominated by Y, denoted $X \preccurlyeq Y$, if $\mathbb{P}(X \ge t) \leqslant \mathbb{P}(Y \ge t)$ for every $t \in \mathbb{R}$.

Lemma 18. Let X_1, \ldots, X_n and Y_1, \ldots, Y_n be independent real random variables, such that $X_i \preccurlyeq Y_i$ for $i = 1, \ldots, n$. Then $\sum_{i=1}^n X_i \preccurlyeq \sum_{i=1}^n Y_i$.

Proof. For i = 1, ..., n, let $F_i(t) = \mathbb{P}(X_i \leq t)$ be the cumulative distribution function (c.d.f.) of X_i , and $F_i^+(u) = \inf\{t \in \mathbb{R} : F_i(t) \geq u\}$ for $u \in (0, 1)$ be its right-continuous inverse. Likewise, for i = 1, ..., n, let G_i be the c.d.f. of Y_i and G_i^+ its right-continuous inverse. Since $X_i \leq Y_i$, we have $F_i \geq G_i$ and thus $F_i^+ \leq G_i^+$.

Now, let U_1, \ldots, U_n be independent random variables that are uniformly distributed on [0, 1]. Then $(F_i^+(U_i))_{1 \leq i \leq n}$ has the same distribution as $(X_i)_{1 \leq i \leq n}$, while $(G_i^+(U_i))_{1 \leq i \leq n}$ has the same distribution as $(Y_i)_{1 \leq i \leq n}$, thus for any $t \in \mathbb{R}$,

$$\mathbb{P}(X_1 + \dots + X_n \ge t) = \mathbb{P}(F_1^+(U_1) + \dots + F_n^+(U_n) \ge t)$$

$$\leq \mathbb{P}(G_1^+(U_1) + \dots + G_n^+(U_n) \ge t) = \mathbb{P}(Y_1 + \dots + Y_n \ge t).$$

Hence $X_1 + \dots + X_n \preccurlyeq Y_1 + \dots + Y_n$.

We will also use the following concentration inequality, due to Ben-Hamou, Boucheron, and Ohannessian [BHBO17], for the number $D_n = \sum_{j=1}^d \mathbf{1}(N_j \ge 1)$ of distinct classes in the sample.

Lemma 19. Let $s'_n = s'_n(P) = \mathbb{E}_P[D_n] = \sum_{j=1}^d \{1 - (1 - p_j)^n\}$. For any $s \in \mathbb{R}^+$, the following holds: if $s > s'_n$, then

$$\mathbb{P}_{P}(D_n \ge s) \le \exp(-D(s, s'_n));$$
(123)

in addition, if $s < s'_n$, then

$$\mathbb{P}_P(D_n \leqslant s) \leqslant \exp(-D(s, s'_n)).$$
(124)

Finally, for every $\delta \in (0, 1)$,

$$\mathbb{P}_P\Big(D_n \ge 2\big\{s'_n + \log(1/\delta)\big\}\Big) \le \delta.$$
(125)

Proof. Applying [BHBO17, Proposition 3.4] with r = 1 gives, for any $\theta \in \mathbb{R}$,

 $\log \mathbb{E}\left[e^{\theta(D_n - s'_n)}\right] \leqslant s'_n (e^{\theta} - \theta - 1).$

Applying the standard Chernoff method to this Poisson-type moment generating function gives the tail bounds (123) and (124) (see, e.g., [BLM13, p. 23]). To obtain the bound (125), further relax the m.g.f. bound above by noting that for $\theta \in (0, 1)$, one has $e^{\theta} - \theta - 1 = \sum_{k \ge 2} \frac{\theta^k}{k!} \le \sum_{k \ge 2} \frac{\theta^k}{2(1-\theta)}$, and apply the sub-gamma tail bound [BLM13, p. 29] to conclude that, with probability at least $1 - \delta$,

$$D_n < s'_n + \sqrt{2s'_n \log(1/\delta)} + \log(1/\delta) \le 2(s'_n + \log(1/\delta)).$$

This concludes the proof.

Finally, the following lemma was used in the proof of the decomposition of Lemma 3.

Lemma 20. Let $C \ge 4$, and define $\phi(t) = (t \log t - t + 1)/(\sqrt{t} - 1)^2$ for $t \in \mathbb{R}_+ \setminus \{1\}$, extended by continuity by $\phi(1) = 2$. For every $p, q \in \mathbb{R}^+$ such that $q \ge p/C$, one has

$$(\sqrt{p} - \sqrt{q})^2 \leqslant D(p,q) \leqslant \phi(C)(\sqrt{p} - \sqrt{q})^2 \leqslant 4\log(C)(\sqrt{p} - \sqrt{q})^2.$$
(126)

Proof. If q = 0, then by assumption p = 0 as well, hence all terms in (126) are equal to 0 and the inequalities hold; the same holds when p = q. We thus assume q > 0 and $p \neq q$. In this case, we have $D(p,q)/(\sqrt{p} - \sqrt{q})^2 = \phi(p/q)$. Now, a direct computation shows that the map $t \mapsto \phi(t^2)$ is continuously differentiable on \mathbb{R}^*_+ , with derivative $t \mapsto 2(t-1)^{-3}\psi(t)$ where $\psi(t) = t^2 - 2t\log t - 1$. Since $\psi'(t) = 2(t-1-\log t) \ge 0$ for any t > 0, the function ψ is non-decreasing, and since $\psi(1) = 0$ we deduce that $\psi \le 0$ on (0,1] and $\psi \ge 0$ on $[1, +\infty)$. Hence $\frac{d}{dt}\phi(t^2) = 2(t-1)^{-3}\psi(t) \ge 0$ for any t > 0, $t \ne 1$, thus by continuity of ϕ at 0 and 1, the function ϕ is non-decreasing on \mathbb{R}^+ .

Since $\phi(0) = 1$ and $0 \leq p/q \leq C$, we deduce that $(\sqrt{p} - \sqrt{q})^2 \leq D(p,q) \leq \phi(C)(\sqrt{p} - \sqrt{q})^2$. Finally, since $C \geq 4$ we have

$$\phi(C) = \frac{C \log C - C + 1}{C(1 - 1/\sqrt{C})^2} \leqslant \frac{C \log C}{C(1 - 1/\sqrt{4})^2} = 4 \log C \,,$$

which concludes the proof.

References

- [Agr20] Rohit Agrawal. Finite-sample concentration of the multinomial in relative entropy. *IEEE Transactions on Information Theory*, 66(10):6297–6302, 2020.
- [Agr22] Rohit Agrawal. Finite-sample concentration of the empirical relative entropy around its mean. *Preprint arXiv:2203.00800*, 2022.
- [Bac24] Francis Bach. Learning Theory from First Principles. MIT Press, 2024.
- [Bar87] Andrew R. Barron. Are Bayes rules consistent in information? In Open Problems in Communication and Computation, pages 85–91. Springer, 1987.
- [BGPV21] Arnab Bhattacharyya, Sutanu Gayen, Eric Price, and N. V. Vinodchandran. Near-optimal learning of tree-structured distributions by Chow-Liu. In Proceedings of the 53rd annual ACM-SIGACT Symposium on Theory of Computing, pages 147–160, 2021.
- [BHBO17] Anna Ben-Hamou, Stéphane Boucheron, and Mesrob I. Ohannessian. Concentration inequalities in the infinite urn scheme for occupancy counts and the missing mass, with applications. *Bernoulli*, 23(1):249–287, 2017.
- [BK13] Daniel Berend and Aryeh Kontorovich. On the concentration of the missing mass. *Electronic Communications in Probability*, 18:1–7, 2013.
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. Concentration Inequalities: A Nonasymptotic Theory of Independence. Oxford University Press, 2013.
- [BP23] Alankrita Bhatt and Ankit Pensia. Sharp concentration inequalities for the centred relative entropy. Information and Inference: A Journal of the IMA, 12(1):524–550, 2023.
- [BS04] Dietrich Braess and Thomas Sauer. Bernstein polynomials and learning theory. Journal of Approximation Theory, 128(2):187–206, 2004.
- [Can20] Clément L. Canonne. A short note on learning discrete distributions. arXiv preprint arXiv:2002.11457, 2020.
- [Cat97] Olivier Catoni. The mixture approach to universal model selection. Technical report, École Normale Supérieure, 1997.
- [Cat04] Olivier Catoni. Statistical Learning Theory and Stochastic Optimization. Ecole d'Eté de Probabilités de Saint-Flour XXXI - 2001, volume 1851 of Lecture Notes in Mathematics. Springer-Verlag Berlin Heidelberg, 2004.
- [Cat12] Olivier Catoni. Challenging the empirical mean and empirical variance: a deviation study. Annales de l'Institut Henri Poincaré, Probabilités et Statistiques, 48(4):1148–1185, 2012.

[CG99]	Stanley F. Chen and Joshua Goodman. An empirical study of smoothing techniques for language modeling. Computer Speech & Language, $13(4):359-394$, 1999.
[CKT24]	Julien Chhor, Olga Klopp, and Alexandre Tsybakov. Generalized multi-view model: Adaptive density estimation under low-rank constraints. <i>Preprint arXiv:2404.17209</i> , 2024.
[Csi98]	Imre Csiszár. The method of types. <i>IEEE Transactions on Information Theory</i> , 44(6):2505–2523, 1998.
[CSS23]	Clément L. Canonne, Ziteng Sun, and Ananda T. Suresh. Concentration bounds for discrete distribution estimation in KL divergence. In <i>Proceedings of the 2023 IEEE International Symposium on Information Theory</i> , pages 2093–2098, 2023.
[CT06]	Thomas M. Cover and Joy A. Thomas. <i>Elements of Information Theory</i> . Wiley-Interscience, 2nd edition, 2006.
[DLLO16]	Luc Devroye, Matthieu Lerasle, Gábor Lugosi, and Roberto I. Oliveira. Sub-gaussian mean estimators. <i>The Annals of Statistics</i> , 44(6):2695–2725, 2016.
[Dur10]	Rick Durrett. <i>Probability: theory and examples.</i> Cambridge University Press, 5th edition, 2010.
[FOOP17]	Moein Falahatgar, Mesrob I. Ohannessian, Alon Orlitsky, and Venkatadheeraj Pichapati. The power of absolute discounting: all-dimensional distribution estimation. In Advances in Neural Information Processing Systems 30, 2017.
[Gas18]	Élisabeth Gassiat. Universal Coding and Order Identification by Model Selection Methods. Springer Monographs in Mathematics. Springer, 2018.
[Goo53]	Irving J. Good. The population frequencies of species and the estimation of population parameters. $Biometrika$, 40(3-4):237–264, 1953.
[GR20]	F. Richard Guo and Thomas S. Richardson. Chernoff-type concentration of empirical probabilities in relative entropy. <i>IEEE Transactions on Information Theory</i> , 67(1):549–558, 2020.
[HJW15]	Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Minimax estimation of discrete distributions under ℓ_1 loss. <i>IEEE Transactions on Information Theory</i> , 61(11):6343–6354, 2015.
[HJW23]	Yanjun Han, Soham Jana, and Yihong Wu. Optimal prediction of Markov chains with and without spectral gap. <i>IEEE Transactions on Information Theory</i> , 69(6):3920–3959, 2023.
[Hut13]	Marcus Hutter. Sparse adaptive Dirichlet-multinomial-like processes. In <i>Proceedings of the 26th Annual Conference on Learning Theory</i> , pages 432–459, 2013.
[IH81]	Ildar Abdulovič Ibragimov and Rafail Zalmanovich Has'minskii. <i>Statistical estimation: asymptotic theory.</i> Springer Science & Business Media, 1981.
[JDP83]	Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. <i>The Annals of Statistics</i> , 11(1):286–295, 1983.
[JM25]	Daniel Jurafsky and James H. Martin. Speech and Language Processing. Jan. 12 Draft, 3rd edition, 2025.
[KN95]	Reinhard Kneser and Hermann Ney. Improved backing-off for M-gram language modeling. In <i>Proceedings of the international conference on Acoustics, Speech, and Signal Processing</i> , pages 181–184. IEEE, 1995.
[KOPS15]	Sudeep Kamath, Alon Orlitsky, Dheeraj Pichapati, and Ananda Theertha Suresh. On learning distributions from their samples. In <i>Proceedings of the 28th Conference on Learning Theory</i> , pages 1066–1100, 2015.
[KT81]	Raphail Krichevsky and Victor Trofimov. The performance of universal encoding. <i>IEEE Transactions on Information Theory</i> , 27(2):199–207, 1981.

- [Lap25] Pierre-Simon de Laplace. Essai philosophique sur les probabilités. Bachelier, fifth edition, 1825.
- [Lat97] Rafał Latała. Estimation of moments of sums of independent real random variables. *The* Annals of Probability, 25(3):1502–1513, 1997.
- [LCY00] Lucien Le Cam and Grace Lo Yang. Asymptotics in statistics: some basic concepts. Springer Series in Statistics. Springer-Verlag New York, 2000.
- [MG22] Jaouad Mourtada and Stéphane Gaïffas. An improper estimator with optimal excess risk in misspecified density estimation and logistic regression. *Journal of Machine Learning Research*, 23(31):1–49, 2022.
- [MJT⁺20] Jay Mardia, Jiantao Jiao, Ervin Tánczos, Robert D. Nowak, and Tsachy Weissman. Concentration inequalities for the empirical distribution of discrete distributions: beyond the method of types. *Information and Inference: A Journal of the IMA*, 9(4):813–850, 2020.
- [MO03] David McAllester and Luis Ortiz. Concentration inequalities for the missing mass and for histogram rule error. *Journal of Machine Learning Research*, 4(Oct):895–911, 2003.
- [MS00] David McAllester and Robert E. Schapire. On the convergence rate of Good-Turing estimators. In *Proceedings of the 13th annual conference on Computational Learning Theory*, pages 1–6, 2000.
- [MU17] Michael Mitzenmacher and Eli Upfal. Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis. Cambridge University Press, 2017.
- [NEK94] Hermann Ney, Ute Essen, and Reinhard Kneser. On structuring probabilistic dependences in stochastic language modelling. *Computer Speech & Language*, 8(1):1–38, 1994.
- [OD12] Mesrob I. Ohannessian and Munther A. Dahleh. Rare probability estimation under regularly varying heavy tails. In *Proceedings of the 25th Annual Conference on Learning Theory*, pages 21.1–21.24, 2012.
- [OS15] Alon Orlitsky and Ananda Theertha Suresh. Competitive distribution estimation: Why is Good-Turing good. In Advances in Neural Information Processing Systems 28, 2015.
- [PW23] Yury Polyanskiy and Yihong Wu. Information Theory: From Coding to Learning. Cambridge University Press, 2023.
- [Rob51] Herbert Robbins. Asymptotically subminimax solutions of compound statistical decision problems. In Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, pages 131–149, 1951.
- [Teh06] Yee Whye Teh. A hierarchical Bayesian language model based on Pitman-Yor processes. In Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, pages 985–992, 2006.
- [Tsy09] Alexandre B. Tsybakov. Introduction to nonparametric estimation. Springer, 2009.
- [Vap00] Vladimir Vapnik. The Nature of Statistical Learning Theory. Springer, 2000.
- [vdG99] Sara van de Geer. Empirical Processes in M-estimation. Cambridge University Press, Cambridge, 1999.
- [vdHZCB23] Dirk van der Hoeven, Nikita Zhivotovskiy, and Nicolò Cesa-Bianchi. High-probability risk bounds via sequential predictors. *Preprint arXiv:2308.07588*, 2023.
- [vdV98] Aad van der Vaart. Asymptotic statistics. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1998.
- [Wai19] Martin J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- [YB99] Yuhong Yang and Andrew R. Barron. Information-theoretic determination of minimax rates of convergence. *The Annals of Statistics*, 27(5):1564–1599, 1999.