

Fano's inequality for random variables

Sebastien Gerchinovitz, Pierre Ménard, Gilles Stoltz

▶ To cite this version:

Sebastien Gerchinovitz, Pierre Ménard, Gilles Stoltz. Fano's inequality for random variables. 2017. hal-01470862v1

HAL Id: hal-01470862 https://hal.science/hal-01470862v1

Preprint submitted on 17 Feb 2017 (v1), last revised 4 Jun 2019 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fano's inequality for random variables

Sébastien Gerchinovitz Pierre Ménard

IMT — Université Paul Sabatier Toulouse, France

sebastien.gerchinovitz@math.univ-toulouse.fr pierre.menard@math.univ-toulouse.fr

&

Gilles Stoltz

GREGHEC — HEC Paris, CNRS Jouy-en-Josas, France

stoltz@hec.fr

February 20, 2017

Abstract

We extend Fano's inequality, which controls the average probability of (disjoint) events in terms of the average of some Kullback-Leibler divergences, to work with arbitrary [0, 1]-valued random variables. Our simple two-step methodology is general enough to cover the case of an arbitrary (possibly continuously infinite) family of distributions as well as [0, 1]-valued random variables not necessarily summing up to 1. Several novel applications are provided, in which the consideration of random variables is particularly handy. The most important applications deal with the problem of Bayesian posterior concentration (minimax or distribution-dependent) rates and with a lower bound on the regret in non-stochastic sequential learning. We also improve in passing some earlier fundamental results: in particular, we provide a simple and enlightening proof of the refined Pinsker's inequality of Ordentlich and Weinberger [2005] and derive a sharper Bretagnolle and Huber [1978, 1979] inequality.

MSC 2000 subject classifications. Primary-62B10; secondary-62F15, 68T05.

Keywords: Multiple-hypotheses testing, Lower bounds, Information theory, Bayesian posterior concentration

1. Introduction

Fano's inequality is a popular information-theoretical result that provides a lower bound on worstcase error probabilities in multiple-hypotheses testing problems. It has important consequences in information theory [Cover and Thomas, 2006] and related fields. In mathematical statistics, it has become a key tool to derive lower bounds on minimax (worst-case) rates of convergence for various statistical problems such as nonparametric density estimation, regression, and classification (see, e.g., Tsybakov, 2009, Massart, 2007).

Multiple variants of Fano's inequality have been derived in the literature. They can handle a finite, countable, or even continuously infinite number of hypotheses. Depending on the community, it has been stated in various ways. In this article, we focus on statistical versions of Fano's inequality. For instance, its most classical version states that for all sequences of $N \ge 2$ probability distributions $\mathbb{P}_1, \ldots, \mathbb{P}_N$ on the same measurable space (Ω, \mathcal{F}) , and all events A_1, \ldots, A_N forming a partition of Ω ,

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{\frac{1}{N}\inf_{\mathbb{Q}}\sum_{i=1}^{N}\mathrm{KL}(\mathbb{P}_{i},\mathbb{Q}) + \ln(2)}{\ln(N)},$$

where the infimum in the right-hand side is over all probability distributions \mathbb{Q} over (Ω, \mathcal{F}) . The following alternative version is popular among statisticians and is due to Birgé [2005] and Massart [2007]: with the same notation and conditions,

$$\min_{1 \leqslant i \leqslant N} \mathbb{P}_i(A_i) \leqslant \max\left\{c, \ \frac{\overline{K}}{\ln(N)}\right\} \qquad \text{where} \qquad \overline{K} = \frac{1}{N-1} \sum_{i=2}^N \mathrm{KL}(\mathbb{P}_i, \mathbb{P}_1) \tag{1}$$

for some universal constant $c \in (0, 1)$. In both cases, the link to multiple-hypotheses testing is the following: when applied to events of the form $A_i = \{\widehat{\theta} = i\}$, the last inequality provides a lower bound on the worst-case error probability $\max_{1 \le i \le N} \mathbb{P}_i(\widehat{\theta} \ne i)$ for any estimator $\widehat{\theta}$.

Several extensions to more complex settings were derived in the past. For example, Han and Verdú [1994] addressed the case of countably infinitely many probability distributions, while Duchi and Wainwright [2013] and Chen et al. [2016] further generalized Fano's inequality to continuously infinitely many distributions. Gushchin [2003] extended Fano's inequality in another direction, by considering [0, 1]-valued random variables Z_i such that $Z_1 + \ldots + Z_N = 1$, instead of the special case $Z_i = \mathbb{1}_{A_i}$. All these extensions, as well as others recalled in Section 6, provide a variety of tools that adapt nicely to the variety of statistical problems.

Main contributions. In this article, we revisit Fano's inequality and make the following three sets of contributions. First, we extend Fano's inequality to both continuously many distributions \mathbb{P}_{θ} and arbitrary [0, 1]-valued random variables Z_{θ} that are not required to sum up (or integrate) to 1. We also point out that the alternative distribution \mathbb{Q} could vary with θ . Despite the high degree of generality, the proofs of these results are simple thanks to a reduction to Bernoulli distributions.

Second, we provide new statistical applications, illustrating in particular that it is handy to be able to consider random variables (not necessarily summing up to 1). The two main such applications deal with Bayesian posterior concentration lower bounds and a regret lower bound in non-stochastic sequential learning.

Finally, as a by-product of our simplified analysis, we highlight a direct connection between Fano's and Pinsker's inequalities. We prove a common bound that both implies Pinsker's inequality for N = 2 and a Fano-type inequality for all $N \ge 2$. These two inequalities were classically thought to be useful in distinct regimes (Pinsker's inequality for N = 2, Fano's inequality for $N \ge 3$). This is one reason why Birgé [2005] designed his alternative version (1) of the most classical version of Fano's inequality in order to make it nontrivial even for N = 2.

Content and outline of this article. The main body of this article contains new results and a new look at some older results (that we sometimes generalize), while the appendix contains omitted technical derivations and discussions, or even some known material (which we provide for the sake of self-completeness).

More precisely, Sections 2 and 3 explain our two-step methodology to obtain several versions of Fano's inequality, at various degrees of generality. These inequalities are discussed and compared to the literature later in the article, in Section 6. Before that, we present in Section 4 our two main applications: lower bounds for minimax Bayesian posterior concentration and for non-stochastic sequential learning. Section 5 presents two other applications which—perhaps surprisingly—follow from the special case N = 1 in Fano's inequality. One of these applications is about distribution-dependent lower bounds on Bayesian posterior concentration (elaborating on results by Hoffmann et al., 2015). Section 7 concludes the main body of the article and provides new and simpler proofs of some important bounds on the Kullback-Leibler divergence, the main contributions being a short and enlightening proof of the refined Pinsker's inequality by Ordentlich and Weinberger [2005], and a sharper Bretagnolle and Huber [1978, 1979] inequality.

The appendix of the present article contains the following material. In Section A, we present two toy applications of our continuous Fano's inequality for Gaussian mean-parameter estimation and in a setting of nonparametric regression. Section B provides some background on the problem of Bayesian posterior concentration. Section C carefully discusses the popular version of Fano's inequality proved by Birgé [2005] and Massart [2007]. Section D is a reminder of basic properties of f-divergences (such as the data-processing inequality), and Section E explains how our two-step methodology readily extends to f-divergences. Finally, Section F states and proves a version of Jensen's inequality tailored to the needs of the present article: that holds for general convex sets and for possibly infinite-valued convex functions.

Notation. Let \mathbb{P}, \mathbb{Q} be two probability distributions on the same measurable space (Ω, \mathcal{F}) . We write $\mathbb{P} \ll \mathbb{Q}$ to indicate that \mathbb{P} is absolutely continuous with respect to \mathbb{Q} . Moreover, the Kullback-Leibler divergence $\mathrm{KL}(\mathbb{P}, \mathbb{Q})$ is defined by

$$\mathrm{KL}(\mathbb{P},\mathbb{Q}) = \begin{cases} \int_{\Omega} \ln\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) \mathrm{d}\mathbb{P} & \text{if } \mathbb{P} \ll \mathbb{Q}; \\ +\infty & \text{otherwise.} \end{cases}$$

We write $\operatorname{Ber}(p)$ for the Bernoulli distribution with parameter p. We also use the usual measuretheoretic conventions in $\mathbb{R} \cup \{+\infty\}$; in particular $0 \times (+\infty) = 0$ and $1/0 = +\infty$, as well as 0/0 = 0. We also set $\ln(0) = -\infty$ and $0 \ln(0) = 0$.

2. How to derive a Fano-type inequality: an example

In this section we explain on an example the methodology to derive Fano-type inequalities. We will present the generalization of the approach and the resulting bounds in Section 3, but the proof below already contains the two key arguments: a reduction to Bernoulli distributions, and a lower bound on the kl function. We discuss how novel (or not novel) our results and approaches are in Section 6.

Proposition 1. Given an underlying measurable space, for all probability pairs \mathbb{P}_i , \mathbb{Q}_i and all events A_i (non necessarily disjoint), where $i \in \{1, \ldots, N\}$, with $0 < \frac{1}{N} \sum_{i=1}^N \mathbb{Q}_i(A_i) < 1$, we have

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{\frac{1}{N}\sum_{i=1}^{N}\mathrm{KL}(\mathbb{P}_{i},\mathbb{Q}_{i}) + \ln(2)}{-\ln\left(\frac{1}{N}\sum_{i=1}^{N}\mathbb{Q}_{i}(A_{i})\right)}$$

In particular, if $N \ge 2$ and the A_i form a partition,

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{\frac{1}{N}\inf_{\mathbb{Q}}\sum_{i=1}^{N}\mathrm{KL}(\mathbb{P}_{i},\mathbb{Q}) + \ln(2)}{\ln(N)}.$$

The proof uses the Kullback-Leibler divergence function kl between Bernoulli distributions: for all $(p,q) \in [0,1]^2$,

$$\mathrm{kl}(p,q) \stackrel{\mathrm{def}}{=} \mathrm{KL}\big(\mathrm{Ber}(p),\mathrm{Ber}(q)\big) = p \ln\left(\frac{p}{q}\right) + (1-p) \ln\left(\frac{1-p}{1-q}\right),$$

with the usual measure-theoretic conventions. The proof also relies on the two following informationtheoretic tools, which are proved in Appendix D. Lemma 2 indicates that transforming the data at hand can only reduce the ability to distinguish between two probability distributions. Corollary 3 states that the Kullback-Leilber divergence is jointly convex.

Lemma 2 (Contraction of entropy; also known as data-processing inequality). Let \mathbb{P} and \mathbb{Q} be two probability distributions over the same measurable space (Ω, \mathcal{F}) , and let X be any random variable on (Ω, \mathcal{F}) . Denote by \mathbb{P}^X and \mathbb{Q}^X the laws of X under \mathbb{P} and \mathbb{Q} respectively. Then,

$$\mathrm{KL}ig(\mathbb{P}^X,\mathbb{Q}^Xig)\leqslant\mathrm{KL}(\mathbb{P},\mathbb{Q})$$
 .

Corollary 3 (Joint convexity of KL). The Kullback-Leibler divergence KL is jointly convex, i.e., for all probability distributions $\mathbb{P}_1, \mathbb{P}_2$ and $\mathbb{Q}_1, \mathbb{Q}_2$ over the same measurable space (Ω, \mathcal{F}) , and all $\lambda \in (0, 1)$,

$$\mathrm{KL}(\lambda \mathbb{P}_1 + (1-\lambda)\mathbb{P}_2, \lambda \mathbb{Q}_1 + (1-\lambda)\mathbb{Q}_2) \leq \lambda \mathrm{KL}(\mathbb{P}_1, \mathbb{Q}_1) + (1-\lambda) \mathrm{KL}(\mathbb{P}_2, \mathbb{Q}_2)$$

Proof (of Proposition 1): Our first step is to reduce the problem to Bernoulli distributions. Using first the joint convexity of the Kullback-Leibler divergence (Corollary 3), and second the data-processing inequality with the indicator functions $X = \mathbb{1}_{A_i}$ (Lemma 2), we get

$$\operatorname{kl}\left(\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}), \frac{1}{N}\sum_{i=1}^{N}\mathbb{Q}_{i}(A_{i})\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{kl}\left(\mathbb{P}_{i}(A_{i}), \mathbb{Q}_{i}(A_{i})\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{KL}(\mathbb{P}_{i}, \mathbb{Q}_{i}).$$
(2)

Therefore, we have $kl(\overline{p}, \overline{q}) \leq \overline{K}$ with

$$\overline{p} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{P}_i(A_i) \qquad \overline{q} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{Q}_i(A_i) \qquad \overline{K} = \frac{1}{N} \sum_{i=1}^{N} \mathrm{KL}(\mathbb{P}_i, \mathbb{Q}_i) \,. \tag{3}$$

Our second and last step is to lower bound $\operatorname{kl}(\overline{p},\overline{q})$ to extract an upper bound on \overline{p} . Noting that $\overline{p}\ln(\overline{p}) + (1-\overline{p})\ln(1-\overline{p}) \ge -\ln(2)$, we have, by definition of $\operatorname{kl}(\overline{p},\overline{q})$,

$$\operatorname{kl}(\overline{p},\overline{q}) \ge \overline{p}\ln(1/\overline{q}) - \ln(2), \quad \text{thus} \quad \overline{p} \le \frac{\operatorname{kl}(\overline{p},\overline{q}) + \ln(2)}{\ln(1/\overline{q})}.$$

$$\tag{4}$$

Substituting the upper bound $kl(\bar{p}, \bar{q}) \leq \bar{K}$ in (4) concludes the proof.

3. Various Fano-type inequalities, with the same two ingredients

We extend the approach of Section 2 and derive a broad family of Fano-type inequalities, which will be of the form

$$\overline{p} \leqslant \psi(\overline{q}, \overline{K}) ,$$

where the average quantities \overline{p} , \overline{q} and \overline{K} are described in Section 3.1 and where the functions ψ are described in Section 3.2. The simplest example that we considered in Section 2 was given, for some probability distributions $\mathbb{P}_i, \mathbb{Q}_i$ and some events A_i , by

$$\overline{p} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{P}_i(A_i) \qquad \overline{q} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{Q}_i(A_i) \qquad \overline{K} = \frac{1}{N} \sum_{i=1}^{N} \mathrm{KL}(\mathbb{P}_i, \mathbb{Q}_i).$$

But we also address here the more general case where the finite averages are replaced with integrals over any measurable space Θ , and where the indicator functions $\mathbb{1}_{A_i}$ are replaced with arbitrary [0,1]-valued random variables Z_{θ} , where $\theta \in \Theta$.

Section 3.3 states some examples of such Fano-type inequalities, based on a choice of averages picked in Section 3.1 and a choice of functions ψ picked in Section 3.2.

We recall that the novelty (or lack of novelty) of our results will be discussed in detail in Section 6.

3.1. Reduction to Bernoulli distributions

As in Section 2, we can use the contraction of relative entropy to lower bound any Kullback-Leibler divergence by that of suitably chosen Bernoulli distributions. We present four such reductions, in increasing degree of generality. We only recall how to prove the first one, since they are all similar.

Finitely many distributions; uniform averages. We consider some underlying measurable space, N pairs of probability distributions \mathbb{P}_i , \mathbb{Q}_i on this space, and N events A_i , where $i \in \{1, \ldots, N\}$. The events A_i do not need to be disjoint. Recall from Section 2 that

$$\operatorname{kl}\left(\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}), \frac{1}{N}\sum_{i=1}^{N}\mathbb{Q}_{i}(A_{i})\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{kl}\left(\mathbb{P}_{i}(A_{i}), \mathbb{Q}_{i}(A_{i})\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{KL}(\mathbb{P}_{i}, \mathbb{Q}_{i}),$$
(5)

where the first inequality is by joint convexity of the Kullback-Leibler divergence (Corollary 3 above), and where the second inequality is by the data-processing inequality (Lemma 2 above), considering the indicator functions $X = \mathbb{1}_{A_i}$.

Countably many distributions; general averages. The argument above carries over to any convex combination $\alpha = (\alpha_1, \alpha_2, \ldots)$ of countably many pairs of probability distributions \mathbb{P}_i , \mathbb{Q}_i and events A_i , where $i \in \{1, 2, \ldots\}$. The convex combination α can be thought of as a prior distribution. Let $\delta_{(x,y)}$ denote the Dirac mass at $(x,y) \in \mathbb{R}^2$. Using the general form of Jensen's inequality stated in Lemma 28 (Appendix F) with $\varphi(p,q) = \mathrm{kl}(p,q)$ on the convex set $C = [0,1]^2$, together with the probability measure $\mu = \sum_i \alpha_i \delta_{(\mathbb{P}_i(A_i),\mathbb{Q}_i(A_i))}$, we get

$$\operatorname{kl}\left(\sum_{i\geq 1}\alpha_{i}\,\mathbb{P}_{i}(A_{i}),\,\sum_{i\geq 1}\alpha_{i}\,\mathbb{Q}_{i}(A_{i})\right)\leqslant\sum_{i\geq 1}\alpha_{i}\,\operatorname{kl}\left(\mathbb{P}_{i}(A_{i}),\mathbb{Q}_{i}(A_{i})\right)\leqslant\sum_{i\geq 1}\alpha_{i}\,\operatorname{KL}(\mathbb{P}_{i},\mathbb{Q}_{i})\,.\tag{6}$$

Distributions indexed by a possibly continuous set; general averages. We consider statistical models \mathbb{P}_{θ} , \mathbb{Q}_{θ} with a measurable parameter space (Θ, \mathcal{G}) , a prior probability distribution ν over Θ , and a collection A_{θ} of events (not necessarily disjoint) such that

 $\theta \in \Theta \longmapsto (\mathbb{P}_{\theta}(A_{\theta}), \mathbb{Q}_{\theta}(A_{\theta}))$ and $\theta \in \Theta \longmapsto \mathrm{KL}(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta})$

are \mathcal{G} -measurable. The reduction is this time (we use again the general form of Jensen's inequality in Appendix F, Lemma 28):

$$\mathrm{kl}\left(\int_{\Theta} \mathbb{P}_{\theta}(A_{\theta}) \,\mathrm{d}\nu(\theta), \int_{\Theta} \mathbb{Q}_{\theta}(A_{\theta}) \,\mathrm{d}\nu(\theta)\right) \leq \int_{\Theta} \mathrm{kl}\big(\mathbb{P}_{\theta}(A_{\theta}), \mathbb{Q}_{\theta}(A_{\theta})\big) \,\mathrm{d}\nu(\theta) \leq \int_{\Theta} \mathrm{KL}(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta}) \,\mathrm{d}\nu(\theta).$$
(7)

Random variables; general averages. In the reductions above, it was unnecessary that the sets A_i or A_{θ} form a partition or even be disjoint. It is therefore not surprising that the former reductions can be generalized by replacing the indicator functions $\mathbb{1}_{A_i}$ or $\mathbb{1}_{A_{\theta}}$ with arbitrary [0, 1]-valued random variables Z_i or Z_{θ} . The most elegant way of generalizing the reduction is the following consequence of Lemma 2 (extracted from Garivier et al. [2017] and proved again in Appendix D for the sake of self-completeness).

Corollary 4 (Contraction of entropy; with expectations of random variables). Let \mathbb{P} and \mathbb{Q} be two probability distributions over the same measurable space (Ω, \mathcal{F}) , and let X be any random variable on (Ω, \mathcal{F}) taking values in [0, 1]. Denote by $\mathbb{E}_{\mathbb{P}}[X]$ and $\mathbb{E}_{\mathbb{Q}}[X]$ the expectations of X under \mathbb{P} and \mathbb{Q} respectively. Then,

$$\mathrm{kl}(\mathbb{E}_{\mathbb{P}}[X], \mathbb{E}_{\mathbb{Q}}[X]) \leq \mathrm{KL}(\mathbb{P}, \mathbb{Q}).$$

We now state the reduction in the case of finitely many distributions and uniform averages, as well as in the case of distributions indexed by a possibly continuous set. In the first case, we consider a collection Z_1, \ldots, Z_N of random variables taking values in [0, 1] and denote by $\mathbb{E}_{\mathbb{P}_i}$ and $\mathbb{E}_{\mathbb{Q}_i}$ the expectations with respect to \mathbb{P}_i and \mathbb{Q}_i ; the reduction is

$$\operatorname{kl}\left(\frac{1}{N}\sum_{i=1}^{N}\mathbb{E}_{\mathbb{P}_{i}}\left[Z_{i}\right], \frac{1}{N}\sum_{i=1}^{N}\mathbb{E}_{\mathbb{Q}_{i}}\left[Z_{i}\right]\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{kl}\left(\mathbb{E}_{\mathbb{P}_{i}}\left[Z_{i}\right], \mathbb{E}_{\mathbb{Q}_{i}}\left[Z_{i}\right]\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{KL}(\mathbb{P}_{i}, \mathbb{Q}_{i}).$$
(8)

In the most general case, the [0,1]-valued random variables are denoted by Z_{θ} , where $\theta \in \Theta$, and expectations with respect to \mathbb{P}_{θ} and \mathbb{Q}_{θ} are denoted by $\mathbb{E}_{\mathbb{P}_{\theta}}$ and $\mathbb{E}_{\mathbb{Q}_{\theta}}$; we assume that

$$\theta \in \Theta \longmapsto \left(\mathbb{E}_{\mathbb{P}_{\theta}} [Z_{\theta}], \mathbb{E}_{\mathbb{Q}_{\theta}} [Z_{\theta}] \right) \quad \text{and} \quad \theta \in \Theta \longmapsto \mathrm{KL}(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta})$$

are \mathcal{G} -measurable. The reduction then is

$$\operatorname{kl}\left(\int_{\Theta} \mathbb{E}_{\mathbb{P}_{\theta}}\left[Z_{\theta}\right] \mathrm{d}\nu(\theta), \int_{\Theta} \mathbb{E}_{\mathbb{Q}_{\theta}}\left[Z_{\theta}\right] \mathrm{d}\nu(\theta)\right) \leqslant \int_{\Theta} \operatorname{kl}\left(\mathbb{E}_{\mathbb{P}_{\theta}}\left[Z_{\theta}\right], \mathbb{E}_{\mathbb{Q}_{\theta}}\left[Z_{\theta}\right]\right) \mathrm{d}\nu(\theta) \leqslant \int_{\Theta} \operatorname{KL}(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta}) \, \mathrm{d}\nu(\theta).$$

$$\tag{9}$$

This most general form of the reduction will be used in Section 4.1.

3.2. Any lower bound on kl leads to a Fano-type inequality

The section above indicates that after the reduction to the Bernoulli case, we get inequations of the form (\overline{p} is usually the unknown)

$$\mathrm{kl}(\overline{p},\overline{q}) \leqslant \overline{K}\,,$$

where \overline{K} is an average of Kullback-Leibler divergences, and \overline{p} and \overline{q} are averages of probabilities of events or expectations of [0, 1]-valued random variables.

We thus proceed by lower bounding the kl function. The first bound was already used in Section 2.

The most classical bound. For all $p \in [0, 1]$ and $q \in (0, 1)$,

$$\operatorname{kl}(p,q) \ge p \ln(1/q) - \ln(2), \quad \text{thus} \quad p \le \frac{\operatorname{kl}(p,q) + \ln(2)}{\ln(1/q)}.$$

$$(10)$$

This bound can be improved by replacing the term $\ln(2)$ with $\ln(2-q)$, which leads to a non-trivial bound even if q = 1/2 (as is the case in some applications).

A consequence of a convexity inequality. This bound was known and we recall its proof in Section 7.1. For all $p \in [0, 1]$ and $q \in (0, 1)$,

$$kl(p,q) \ge p \ln(1/q) - \ln(2-q),$$
 thus $p \le \frac{kl(p,q) + \ln(2-q)}{\ln(1/q)}.$ (11)

A (novel) consequence of this bound is that

$$p \leq 0.21 + 0.79 q + \frac{\mathrm{kl}(p,q)}{\ln(1/q)}$$
 (12)

11/

A final bound, of a similar flavor, is stated below. Note that, perhaps surprisingly, it makes a connection between Pinsker's and Fano's inequalities.

A consequence of a refined Pinsker's inequality. The first inequality was known, the second is a novel but straightforward consequence of it. We provide the proofs in Section 7.2. For all $p \in [0, 1]$ and $q \in (0, 1)$,

$$\operatorname{kl}(p,q) \ge \max\left\{\ln\left(\frac{1}{q}\right), 2\right\} (p-q)^2, \quad \text{thus} \quad p \le q + \sqrt{\frac{\operatorname{kl}(p,q)}{\max\left\{\ln(1/q), 2\right\}}}.$$
(13)

3.3. Examples of combinations

The combination of (8) and (13) ensures the following Fano-type inequality for finitely many random variables, whose sum does not need to be 1.

Lemma 5. Given an underlying measurable space, for all probability pairs \mathbb{P}_i , \mathbb{Q}_i and for all [0,1]valued random variables Z_i defined on this measurable space, where $i \in \{1, \ldots, N\}$, with

$$0 < \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}_{\mathbb{Q}_i} \left[Z_i \right] < 1$$

we have

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{E}_{\mathbb{P}_{i}}[Z_{i}] \leqslant \frac{1}{N}\sum_{i=1}^{N}\mathbb{E}_{\mathbb{Q}_{i}}[Z_{i}] + \sqrt{\frac{\frac{1}{N}\sum_{i=1}^{N}\mathrm{KL}(\mathbb{P}_{i},\mathbb{Q}_{i})}{-\ln\left(\frac{1}{N}\sum_{i=1}^{N}\mathbb{E}_{\mathbb{Q}_{i}}[Z_{i}]\right)}}$$

The combination of (7) and (10) yields a continuous version of Fano's inequality. (We discard again all measurability issues.)

Lemma 6. We consider a mesurable space (Θ, \mathcal{E}) equipped with a probability distribution ν . Given an underlying measurable space (Ω, \mathcal{F}) , for all two collections $\mathbb{P}_{\theta}, \mathbb{Q}_{\theta}$, of probability distributions over this space and all collections of events A_{θ} of (Ω, \mathcal{F}) , where $\theta \in \Theta$, with

$$0 < \int_{\Theta} \mathbb{Q}_{\theta}(A_{\theta}) \,\mathrm{d}\nu(\theta) < 1\,,$$

we have

$$\int_{\Theta} \mathbb{P}_{\theta}(A_{\theta}) \, \mathrm{d}\nu(\theta) \leqslant \frac{\int_{\Theta} \mathrm{KL}(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta}) \, \mathrm{d}\nu(\theta) + \ln(2)}{-\ln\left(\int_{\Theta} \mathbb{Q}_{\theta}(A_{\theta}) \, \mathrm{d}\nu(\theta)\right)}$$

3.4. Extensions to *f*-divergences

Gushchin [2003] generalized Fano-type inequalities with the Kullback-Leibler divergence (as above) to arbitrary f-divergences, in the case where finitely many [0, 1]-valued random variables $Z_1 + \ldots + Z_N =$ 1 are considered. As we discuss in Appendix E, the main reason why this generalization was possible is that f-divergences also satisfy a data-processing inequality. We show that all the reductions to Bernoulli distributions discussed in Section 3.1 go through for f-divergences as well.

3.5. On the sharpness of the obtained bounds

The reductions of Section 3.1 are sharp in the sense that they can hold with equality (they cannot be improved at this level of generality). Now, we want to draw from the result of this first reduction, which is of the form $\operatorname{kl}(\overline{p}, \overline{q}) \leq \overline{K}$, an upper bound on \overline{p} . We introduce the generalized inverse of kl in its second argument: for all $q \in [0, 1]$ and all $y \ge 0$,

$$kl(\cdot, q)^{(-1)}(y) \stackrel{\text{def}}{=} \sup\{p \in [0, 1] : kl(p, q) \leq y\};$$

when $q \in (0, 1)$, it is thus equal to the largest root q of the equation kl(p, q) = y if $y \leq ln(1/q)$ or to 1 otherwise. We then get

$$\overline{p} \leqslant \operatorname{kl}(\cdot, \overline{q})^{(-1)}(\overline{K}).$$

This formulation should be reminiscent of Birgé [2005, Theorem 2], but has one major practical drawback: it is unreadable, and this is why we considered the lower bounds of Section 3.2.

Question is now how sharp these lower bounds on kl are. They are all (in spirit) of the form

$$\mathrm{kl}(p,q) \geqslant p \ln \left(\frac{1}{q}\right) - \dots \qquad \text{thus} \qquad p \leqslant \frac{\mathrm{kl}(p,q)}{\ln(1/q)} + \dots,$$

where the ... on the right refer to terms that vanish when $q \to 0$. In the applications, q is typically small and the main term $kl(p,q)/\ln(1/q)$ is of the order of a constant. Therefore, the lemma below explains that up to the ... terms, the bounds of Section 3.2 are essentially optimal.

Lemma 7. For all $q \in (0,1)$ and $p \in [0,1]$, whenever $p \ge q$, we have

$$\operatorname{kl}(p,q) \leqslant p \ln\left(\frac{1}{q}\right) \qquad thus \qquad p \geqslant \frac{\operatorname{kl}(p,q)}{\ln(1/q)}$$

Proof: We note that when $p \ge q$, we have $(1-p)/(1-q) \le 1$, so that

$$\mathrm{kl}(p,q) = p \ln\left(\frac{1}{q}\right) + \underbrace{p \ln(p)}_{\leqslant 0} + (1-p) \underbrace{\ln\left(\frac{1-p}{1-q}\right)}_{\leqslant 0} \leqslant p \ln\left(\frac{1}{q}\right),$$

hence the first inequality.

4. Main applications

We present two new applications of Fano's inequality, with [0, 1]-valued random variables Z_i or Z_{θ} . The topics covered are:

- Bayesian posterior concentration rates, for which we use the reduction (9);
- robust sequential learning (prediction of individual sequences) in the case of sparse losses, which relies on the reduction (8).

As can be seen below, the fact that we are now able to consider arbitrary [0, 1]-valued random variables Z_{θ} on a continuous parameter space Θ makes the proof of the Bayesian posterior concentration lower bound quite simple.

For pedagogical purposes, we also illustrate in Appendix A how to use the continuous Fano's inequality for Gaussian mean-parameter estimation or for some nonparametric regression. Two more applications will also be presented in Section 5; they have a different technical flavor, as they rely on only one pair of distributions, i.e., N = 1.

4.1. Lower bounds on Bayesian posterior concentration rates

In the next paragraphs we show how our continuous Fano's inequality can be used in a simple fashion to derive lower bounds for posterior concentration rates.

Setting and Bayesian terminology. We consider the following density estimation setting: we observe a sample of independent and identically distributed random variables $X_{1:n} = (X_1, \ldots, X_n)$ drawn from a probability distribution P_{θ} on $(\mathcal{X}, \mathcal{F})$, with a fixed but unknown $\theta \in \Theta$. We assume that the measurable parameter space (Θ, \mathcal{G}) is equipped with a prior distribution π and that all $P_{\theta'}$ have a density $p_{\theta'}$ with respect to some reference measure \mathfrak{m} on $(\mathcal{X}, \mathcal{F})$. We also assume that $(x, \theta') \mapsto p_{\theta'}(x)$ is $\mathcal{F} \otimes \mathcal{G}$ -measurable. We can thus consider the transition kernel $(x_{1:n}, A) \mapsto \mathbb{P}_{\pi}(A | x_{1:n})$ defined for all $x_{1:n} \in \mathcal{X}^n$ and all sets $A \in \mathcal{G}$ by

$$\mathbb{P}_{\pi}(A \mid x_{1:n}) = \frac{\int_{A} \prod_{i=1}^{n} p_{\theta'}(x_i) \,\mathrm{d}\pi(\theta')}{\int_{\Theta} \prod_{i=1}^{n} p_{\theta'}(x_i) \,\mathrm{d}\pi(\theta')}$$
(14)

if the denominator lies in $(0, +\infty)$; if it is null or infinite, we set, e.g., $\mathbb{P}_{\pi}(A | x_{1:n}) = \pi(A)$. The resulting random measure $\mathbb{P}_{\pi}(\cdot | X_{1:n})$ is known as the *posterior* distribution.

Let $\ell : \Theta \times \Theta \to \mathbb{R}_+$ be a measurable loss function that we assume to be a pseudo-metric¹. A posterior concentration rate with respect to ℓ is a sequence $(\varepsilon_n)_{n \ge 1}$ of positive real numbers such that, for all $\theta \in \Theta$,

$$\mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi} \big(\theta' : \ell(\theta', \theta) \leqslant \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \longrightarrow 1 \qquad \text{as } n \to +\infty \,,$$

where \mathbb{E}_{θ} denotes the expectation with respect to $X_{1:n}$ where each X_j has the P_{θ} law. The above convergence guarantee means that, as the size n of the sample increases, the posterior mass concentrates in expectation on an ε_n -neighborhood of the true parameter θ . Several variants of this definition exist (e.g., convergence in probability or almost surely; or ε_n that may depend on θ). Though most of these definitions can be handled with the techniques provided below, we only consider this one for the sake of conciseness.

¹The only difference with a metric is that we allow $\ell(\theta, \theta') = 0$ for $\theta \neq \theta'$.

Minimax posterior concentration rate. As our sequence $(\varepsilon_n)_{n\geq 1}$ does not depend on the specific $\theta \in \Theta$ at hand, we may study uniform posterior concentration rates: sequences $(\varepsilon_n)_{n\geq 1}$ such that

$$\inf_{\theta \in \Theta} \mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi} \big(\theta' : \ell(\theta', \theta) \leqslant \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \longrightarrow 1 \qquad \text{as } n \to +\infty \,.$$
(15)

The minimax posterior concentration rate is given by a sequence $(\varepsilon_n)_{n \ge 1}$ such that (15) holds for some prior π while there exists a constant $\gamma \in (0, 1)$ such that for all priors π' on Θ ,

$$\limsup_{n \to +\infty} \inf_{\theta \in \Theta} \mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi'} \big(\theta' : \ell(\theta', \theta) \leqslant \gamma \varepsilon_n \, \big| \, X_{1:n} \big) \Big] < 1.$$

We focus on proving the latter statement and provide a general technique to do so. Though we only illustrate it in the finite-dimensional Gaussian setting, adapting it to, e.g., the nonparametric regression problem of Appendix A.2 would add no technical difficulty.

Proposition 8 (A posterior concentration lower bound in the finite-dimensional Gaussian model). Let $d \ge 1$ be the ambient dimension, $n \ge 1$ the sample size, and $\sigma > 0$ the standard deviation. Assume we observe an *n*-sample $X_{1:n} = (X_1, \ldots, X_n)$ distributed according to $\mathcal{N}(\theta, \sigma^2 I_d)$ for some unknown $\theta \in \mathbb{R}^d$. Let π' be any prior distribution on \mathbb{R}^d . Then the posterior distribution $\mathbb{P}_{\pi'}(\cdot | X_{1:n})$ defined in (14) satisfies, for the Euclidean loss $\ell(\theta', \theta) = \|\theta' - \theta\|_2$ and for $\varepsilon_n = (\sigma/8)\sqrt{d/n}$,

$$\inf_{\theta \in \mathbb{R}^d} \mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi'} \big(\theta' : \| \theta' - \theta \|_2 \leqslant \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \leqslant c_d \,,$$

where $(c_d)_{d \ge 1}$ is a decreasing sequence such that $c_1 \le 0.55$, $c_2 \le 0.37$, and $c_d \to 0.21$ as $d \to +\infty$.

This proposition indicates that the best possible posterior concentration rate is at best $\sigma \sqrt{d/n}$ up to a multiplicative constant; actually, this order of magnitude is the best achievable posterior concentration rate, see, e.g., Le Cam and Yang [2000, Chapter 8].

There are at least two ways to prove the lower bound of Proposition 8. A first one is to use a well-known conversion of "good" Bayesian posteriors into "good" point estimators, which indicates that lower bounds for point estimation can be turned into lower bounds for posterior concentration. For the sake of completeness, we recall this conversion in Appendix B and provide a nonasymptotic variant of Theorem 2.5 by Ghosal et al. [2000].

The second method—followed in the proof below—is however more direct. We use our most general continuous Fano's inequality with the random variables $Z_{\theta} = \mathbb{P}_{\pi'}(\theta' : \|\theta' - \theta\|_2 \leq \varepsilon_n |X_{1:n}) \in [0, 1].$

Proof: We may assume, with no loss of generality, that the probability space on which $X_{1:n}$ is defined is $(\mathbb{R}^d)^n$ endowed with its Borel σ -field and the probability measure $\mathbb{P}_{\theta} = \mathcal{N}(\theta, \sigma^2)^{\otimes n}$. Let ν denote the uniform distribution on the Euclidean ball $B(0, \rho \varepsilon_n) = \{u \in \mathbb{R}^d : ||u||_2 \leq \rho \varepsilon_n\}$ for some $\rho > 1$ to be determined by the analysis. Then, by the continuous Fano inequality in the form given by the combination of (9) and (13), with $\mathbb{Q}_{\theta} = \mathbb{P}_0 = \mathcal{N}(0, \sigma^2)^{\otimes n}$, where 0 denotes the null vector of \mathbb{R}^d , and with the [0, 1]-valued random variables $Z_{\theta} = \mathbb{P}_{\pi'}(\theta' : ||\theta' - \theta||_2 \leq \varepsilon_n |X_{1:n})$, we have

$$\inf_{\theta \in \mathbb{R}^{d}} \mathbb{E}_{\theta} \left[Z_{\theta} \right] \leq \int_{B(0,\rho\varepsilon_{n})} \mathbb{E}_{\theta} \left[Z_{\theta} \right] d\nu(\theta) \leq \int_{B(0,\rho\varepsilon_{n})} \mathbb{E}_{0} \left[Z_{\theta} \right] d\nu(\theta) + \sqrt{\frac{\int_{B(0,\rho\varepsilon_{n})} \mathrm{KL}(\mathbb{P}_{\theta},\mathbb{P}_{0}) d\nu(\theta)}{-\ln \int_{B(0,\rho\varepsilon_{n})} \mathbb{E}_{0} \left[Z_{\theta} \right] d\nu(\theta)}} \\ \leq \left(\frac{1}{\rho} \right)^{d} + \sqrt{\frac{n\rho^{2}\varepsilon_{n}^{2}/(2\sigma^{2})}{d\ln \rho}}, \qquad (16)$$

where the last inequality follows from (17) and (18) below. First note that, by independence, $\operatorname{KL}(\mathbb{P}_{\theta}, \mathbb{P}_0) = n \operatorname{KL}(\mathcal{N}(\theta, \sigma^2), \mathcal{N}(0, \sigma^2)) = n \|\theta\|_2^2 / (2\sigma^2)$, so that

$$\int_{B(0,\rho\varepsilon_n)} \mathrm{KL}\big(\mathbb{P}_{\theta},\mathbb{P}_0\big) \,\mathrm{d}\nu(\theta) = \frac{n}{2\sigma^2} \int_{B(0,\rho\varepsilon_n)} \|\theta\|_2^2 \,\mathrm{d}\nu(\theta) \leqslant \frac{n\rho^2\varepsilon_n^2}{2\sigma^2} \,. \tag{17}$$

Second, using the Fubini-Tonelli theorem (twice) and the definition of

$$Z_{\theta} = \mathbb{P}_{\pi'} \left(\theta' : \| \theta' - \theta \|_2 \leqslant \varepsilon_n \, \big| \, X_{1:n} \right) = \mathbb{E}_{\theta' \sim \mathbb{P}_{\pi'}(\cdot | X_{1:n})} \left[\mathbb{1}_{\{\| \theta' - \theta \|_2 \leqslant \varepsilon_n\}} \right],$$

we can see that

$$q \stackrel{\text{def}}{=} \int_{B(0,\rho\varepsilon_n)} \mathbb{E}_0 \left[Z_\theta \right] d\nu(\theta) = \mathbb{E}_0 \left[\int_{B(0,\rho\varepsilon_n)} \mathbb{E}_{\theta' \sim \mathbb{P}_{\pi'}(\cdot \mid X_{1:n})} \left[\mathbb{1}_{\{ \|\theta' - \theta\|_2 \leqslant \varepsilon_n\}} \right] d\nu(\theta) \right] \\ = \mathbb{E}_0 \left[\mathbb{E}_{\theta' \sim \mathbb{P}_{\pi'}(\cdot \mid X_{1:n})} \left[\int_{B(0,\rho\varepsilon_n)} \mathbb{1}_{\{ \|\theta' - \theta\|_2 \leqslant \varepsilon_n\}} d\nu(\theta) \right] \right] \\ = \mathbb{E}_0 \left[\mathbb{E}_{\theta' \sim \mathbb{P}_{\pi'}(\cdot \mid X_{1:n})} \left[\nu \left(B(\theta',\varepsilon_n) \cap B(0,\rho\varepsilon_n) \right) \right] \right] \leqslant \left(\frac{1}{\rho}\right)^d, \quad (18)$$

where to get the last inequality we used the fact that $\nu(B(\theta', \varepsilon_n) \cap B(0, \rho \varepsilon_n))$ is the ratio of the volume of the (possibly truncated) Euclidean ball $B(\theta', \varepsilon_n)$ of radius ε_n and center θ' with the volume of the support of ν , namely, the larger Euclidean ball $B(0, \rho \varepsilon_n)$, in dimension d.

The proof is then concluded by recalling that $\rho > 1$ was a parameter of the analysis and by picking, e.g., $\varepsilon_n = (\sigma/8)\sqrt{d/n}$: by (16), we have

$$\inf_{\theta \in \mathbb{R}^d} \mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi} \big(\theta' : \| \theta' - \theta \|_2 \leqslant \varepsilon_n \, \big| \, X_{1:n} \big) \Big] = \inf_{\theta \in \mathbb{R}^d} \mathbb{E}_{\theta} \big[Z_{\theta} \big] \leqslant \inf_{\rho \ge 1} \left\{ \left(\frac{1}{\rho} \right)^d + \frac{\rho}{8\sqrt{2\ln\rho}} \right\} \stackrel{\text{def}}{=} c_d \, A_{1:n} \left(\frac{1}{\rho} \right)^d + \frac{\rho}{8\sqrt{2\ln\rho}} \Big\}$$

We can see that $c_1 \leq 0.55$ and $c_2 \leq 0.37$ via the respective choices $\rho = 5$ and $\rho = 3$, while the fact that the limit is smaller than (and actually equal to) $\sqrt{e}/8 \leq 0.21$ follows from the choice $\rho = \sqrt{e}$.

Note that, when using (13) above, we implicitly assumed that the quantity q in (18) lies in (0, 1). The fact that q < 1 follows directly from the upper bound $(1/\rho)^d$ and from $\rho > 1$. Besides, the condition q > 0 is met as soon as $\mathbb{P}_0(\mathbb{P}_{\pi'}(B(0,\varepsilon_n) | X_{1:n}) > 0) > 0$; indeed, for $\theta' \in B(0,\varepsilon_n)$, we have $\nu(B(\theta',\varepsilon_n) \cap B(0,\rho\varepsilon_n)) > 0$ and thus q appears in the last equality of (18) as being lower bounded by the expectation of a positive function over a set with positive probability. If on the contrary $\mathbb{P}_0(\mathbb{P}_{\pi'}(B(0,\varepsilon_n) | X_{1:n}) > 0) = 0$, then $\mathbb{P}_0(Z_0 > 0) = 0$, so that $\inf_{\theta} \mathbb{E}_{\theta}[Z_{\theta}] = \mathbb{E}_0[Z_0] = 0$, which immediately implies the bound of Proposition 8.

Remark 1. Though the lower bound of Proposition 8 is only stated for the posterior distributions $\mathbb{P}_{\pi'}(\cdot | X_{1:n})$, it is actually valid for any transition kernel $Q(\cdot | X_{1:n})$. This is because the proof above relies on general information-theoretic arguments and does not use the particular form of $\mathbb{P}_{\pi'}(\cdot | X_{1:n})$. This is in the same spirit as for minimax lower bounds for point estimation.

In Section 5.2 we derive another type of posterior concentration lower bound that is no longer uniform. More precisely, we prove a distribution-dependent lower bound that specifies how the posterior mass fails to concentrate on ε_n -neighborhoods of θ for every $\theta \in \Theta$.

4.2. Lower bounds in robust sequential learning with sparse losses

We consider a framework of robust sequential learning called prediction of individual sequences. Its origins and core results are described in the monography by Cesa-Bianchi and Lugosi [2006]. In its

simplest version, a decision-maker and an environment play repeatedly as follows: at each round $t \ge 1$, and simultaneously, the environment chooses a vector of losses $\ell_t = (\ell_{1,t}, \ldots, \ell_{N,t}) \in [0,1]^N$ while the decision-maker picks an index $I_t \in \{1, \ldots, N\}$, possibly at random. Both players then observe ℓ_t and I_t . The decision-maker wants to minimize her cumulative regret, the difference between her cumulative loss and the cumulative loss associated with the best constant choice of an index: for $T \ge 1$,

$$R_T = \sum_{t=1}^T \ell_{I_t,t} - \min_{k=1,\dots,N} \sum_{t=1}^T \ell_{k,t} \,.$$

In this setting the optimal regret in the worst-case is of the order of $\sqrt{T \ln(N)}$. Cesa-Bianchi et al. [1997] exhibited an asymptotic lower bound of $\sqrt{T \ln(N)/2}$, based on the central limit theorem and on the fact that the expectation of the maximum of N independent standard Gaussian random variables is of the order of $\sqrt{\ln(N)}$. To do so, they considered stochastic environments drawing independently the loss vectors ℓ_t according to a well-chosen distribution.

Cesa-Bianchi et al. [2005] extended this result to a variant called label-efficient prediction, in which loss vectors are observed upon choosing and with a budget constraint: no more than m observations within T rounds. They prove an optimal and non-asymptotic lower bound on the regret of the order of $T\sqrt{\ln(N)/m}$, based on several applications of Fano's inequality to deterministic strategies of the decision-maker, and then, an application of Fubini's theorem to handle general, randomized, strategies. Our re-shuffled proof technique below shows that a single application of Fano's inequality to general strategies would be sufficient there (details omitted).

Recently, Kwon and Perchet [2016] considered a setting of sparse loss vectors, in which at each round at most s of the N components of the loss vectors ℓ_t are different from zero. They prove an optimal and asymptotic lower bound on the regret of the order of $\sqrt{Ts \ln(N)/N}$, which generalizes the result for the basic framework, in which s = N. Their proof is an extension of the proof of Cesa-Bianchi et al. [1997] and is based on the central limit theorem together with additional technicalities, e.g., the use of Slepian's lemma to deal with some dependencies arising from the sparsity assumption.

The aim of this section is to provide a short and elementary proof of this optimal $\sqrt{Ts \ln(N)/N}$ bound. As a side result, our bound will even be non-asymptotic. The expectation in the statement below is with respect to the internal randomization used by the decision-maker's strategy.

Theorem 9. For all strategies of the decision-maker, for all $N \ge 2$ and all $T > N \ln(N)/(16s)$, there exists a fixed-in-advance sequence of loss vectors ℓ_1, \ldots, ℓ_T in $[0, 1]^N$ that are each s-sparse such that

$$\mathbb{E}[R_T] = \sum_{t=1}^T \mathbb{E}[\ell_{I_t,t}] - \min_{k=1,\dots,N} \sum_{t=1}^T \ell_{k,t} \ge \frac{1}{32} \sqrt{T\frac{s}{N} \ln N}.$$

Proof: We fix $\varepsilon \in (0, s/(2N))$ and consider, as Kwon and Perchet [2016] did, independent and identically distributed loss vectors $\ell_t \in [0, 1]^N$, drawn according to one distribution among P_i , where $1 \leq i \leq N$. Each distribution P_i over $[0, 1]^N$ is defined as the law of a random vector L drawn in two steps as follows. We pick s components uniformly at random among $\{1, \ldots, N\}$. Then, the components k not picked are associated with zero losses, $L_k = 0$. The losses L_k for picked components $k \neq i$ are drawn according to a Bernoulli distribution with parameter 1/2. If component i is picked, its loss L_i is drawn according to a Bernoulli distribution with parameter $1/2 - \varepsilon N/s$. The loss vector $L \in [0, 1]^N$ thus generated is indeed s-sparse. We denote by P_i^T the T-th product distribution $P_i \otimes \cdots \otimes P_i$. We will actually identify the underlying probability and the law P_i^T . Finally, we denote the expectation under P_i^T by \mathbb{E}_i .

Now, under P_i^T , the components $\ell_{k,t}$ of the loss vectors are all distributed according to Bernoulli distributions, with parameters s/(2N) if $k \neq i$ and $s/(2N) - \varepsilon$ if k = i. The expected regret, where the expectation \mathbb{E} is with respect to the strategy's internal randomization and the expectation \mathbb{E}_i is

with respect to the random choice of the loss vectors, is thus larger than

$$\mathbb{E}_{i}\left[\mathbb{E}[R_{T}]\right] \geqslant \sum_{t=1}^{T} \mathbb{E}_{i}\left[\mathbb{E}[\ell_{I_{t},t}]\right] - \min_{k=1,\dots,N} \sum_{t=1}^{T} \mathbb{E}_{i}[\ell_{k,t}] = \sum_{t=1}^{T} \frac{s}{2N} \left(1 - \varepsilon \mathbb{E}_{i}\left[\mathbb{E}[\mathbb{1}_{\{I_{t}=i\}}]\right]\right) - T\left(\frac{s}{2N} - \varepsilon\right)$$
$$= T\varepsilon \left(1 - \mathbb{E}_{i}\left[\mathbb{E}[F_{i}(T)]\right]\right),$$

where

$$F_i(T) = \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{\{I_t=i\}}.$$

All in all, we copied almost word for word the (standard) beginning of the proof by Kwon and Perchet [2016], whose first lower bound is exactly

$$\sup_{\ell_1,\dots,\ell_t} \mathbb{E}[R_T] \ge \frac{1}{N} \sum_{i=1}^N \mathbb{E}_i \Big[\mathbb{E}[R_T] \Big] \ge T \varepsilon \bigg(1 - \frac{1}{N} \sum_{i=1}^N \mathbb{E}_i \Big[\mathbb{E}[F_i(T)] \Big] \bigg).$$
(19)

The main differences arise now: we replace a long asymptotic argument (based on the central limit theorem and the study of the limit via Slepian's lemma) by a single application of Fano's inequality.

We introduce the distribution Q over $[0,1]^N$ corresponding to the same randomization scheme as for the P_i , except that no picked component is favored and that all their corresponding losses are drawn according to the Bernoulli distribution with parameter 1/2. We also denote by \mathbb{P} the probability distribution that underlies the internal randomization of the strategy. An application of Lemma 5 with $\mathbb{P}_i = \mathbb{P} \otimes P_i^T$ and $\mathbb{Q}_i = \mathbb{P} \otimes Q^T$, using that $F_1(T) + \ldots + F_N(T) = 1$ and thus $(1/N) \sum_{i=1}^N \mathbb{E}_Q[\mathbb{E}[F_i(T)]] = 1/N$, yields

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{E}_{i}\Big[\mathbb{E}\big[F_{i}(T)\big]\Big] \leqslant \frac{1}{N} + \sqrt{\frac{1}{N\ln(N)}\sum_{i=1}^{N}\mathrm{KL}\big(\mathbb{P}\otimes P_{i}^{T}, \mathbb{P}\otimes Q^{T}\big)}.$$
(20)

By independence, we get, for all i,

$$\mathrm{KL}(\mathbb{P}\otimes P_i^T, \mathbb{P}\otimes Q^T) = \mathrm{KL}(P_i^T, Q^T) = T \, \mathrm{KL}(P_i, Q) \,.$$

$$(21)$$

We now show that

$$\operatorname{KL}(P_i, Q) \leqslant \frac{s}{N} \operatorname{kl}\left(\frac{1}{2} - \varepsilon \frac{N}{s}, \frac{1}{2}\right).$$
 (22)

Indeed, both P_i and Q can be seen as uniform convex combinations of probability distributions of the following form, indexed by the subsets of $\{1, \ldots, N\}$ with s elements and up to permutations of the Bernoulli distributions in the products below (which does not change the value of the Kullback-Leibler divergences between them):

$$\begin{pmatrix} N-1\\ s-1 \end{pmatrix} \text{ distributions of the form (when } i \text{ is picked}) \\ \operatorname{Ber}\left(\frac{1}{2} - \varepsilon \frac{N}{s}\right) \otimes \bigotimes_{k=2}^{s} \operatorname{Ber}\left(\frac{1}{2}\right) \otimes \bigotimes_{k=s+1}^{N} \delta_{0} \quad \text{and} \quad \bigotimes_{k=1}^{s} \operatorname{Ber}\left(\frac{1}{2}\right) \otimes \bigotimes_{k=s+1}^{N} \delta_{0} \,,$$

where δ_0 denotes the Dirac mass at 0, and

$$\begin{pmatrix} N-1\\s \end{pmatrix} \text{ distributions of the form (when } i \text{ is not picked}) \\ \bigotimes_{k=1}^{s} \operatorname{Ber}\left(\frac{1}{2}\right) \otimes \bigotimes_{k=s+1}^{N} \delta_{0} \quad \text{ and } \quad \bigotimes_{k=1}^{s} \operatorname{Ber}\left(\frac{1}{2}\right) \otimes \bigotimes_{k=s+1}^{N} \delta_{0} \,.$$

Only the first set of distributions contributes to the Kullback-Leibler divergence. By convexity of the Kullback-Leibler divergence (Corollary 3), we thus get the inequality

$$\operatorname{KL}(P_i, Q) \leqslant \frac{\binom{N-1}{s-1}}{\binom{N}{s}} \operatorname{KL}\left(\operatorname{Ber}\left(\frac{1}{2} - \varepsilon \frac{N}{s}\right) \otimes \bigotimes_{k=2}^{s} \operatorname{Ber}\left(\frac{1}{2}\right) \otimes \bigotimes_{k=s+1}^{N} \delta_0, \ \bigotimes_{k=1}^{s} \operatorname{Ber}\left(\frac{1}{2}\right) \otimes \bigotimes_{k=s+1}^{N} \delta_0\right)$$
$$= \frac{s}{N} \operatorname{kl}\left(\frac{1}{2} - \varepsilon \frac{N}{s}, \frac{1}{2}\right),$$

where the last equality is again by independence. Finally, the lemma stated right after this proof shows that

$$\operatorname{kl}\left(\frac{1}{2} - \varepsilon \frac{N}{s}, \frac{1}{2}\right) \leqslant \frac{4N^2 \varepsilon^2}{s^2} \,. \tag{23}$$

Combining (19)–(23), we proved

$$\sup_{\ell_1,\dots,\ell_t} \mathbb{E}[R_T] \ge T\varepsilon \left(1 - \frac{1}{N} - \sqrt{\frac{4NT\varepsilon^2}{s\ln(N)}}\right) \ge T\varepsilon \left(\frac{1}{2} - c\varepsilon\right),$$

where we used $1/N \leq 1/2$ and denoted $c = 2\sqrt{NT}/\sqrt{s\ln(N)}$. A standard optimization suggest the choice $\varepsilon = 1/(4c)$, which is valid, i.e., is indeed $\langle s/(2N) \rangle$ as required, as soon as $T > N \ln(N)/(16s)$. We get a lower bound $T\varepsilon/4$, which is the claimed bound.

Lemma 10. For all $p \in (0,1)$, for all $\varepsilon \in (0,p)$,

$$\operatorname{kl}(p-\varepsilon,p) \leqslant \frac{\varepsilon^2}{p(1-p)}$$

Proof: This result is a special case of the fact that the KL divergence is upper bounded by the χ^2 -divergence. We recall, in our particular case, how this is seen:

$$kl(p-\varepsilon,p) = (p-\varepsilon)\ln\left(1-\frac{\varepsilon}{p}\right) + (1-p+\varepsilon)\ln\left(1+\frac{\varepsilon}{1-p}\right) \leq (p-\varepsilon)\frac{-\varepsilon}{p} + (1-p+\varepsilon)\frac{\varepsilon}{1-p} = \frac{\varepsilon^2}{p} + \frac{\varepsilon^2}{1-p},$$

where we used $\ln(1+u) \leq u$ for all $u > -1$ to get the stated inequality.

where we used $\ln(1+u) \leq u$ for all u > -1 to get the stated inequality.

5. Other applications, with N = 1 pair of distributions

Interestingly, Proposition 1 can be useful even for N = 1 pair of distributions. Rewriting it slightly differently, we indeed have, for all distributions \mathbb{P}, \mathbb{Q} and all events A,

$$\mathbb{P}(A)\ln\left(\frac{1}{\mathbb{Q}(A)}\right) \leqslant \mathrm{KL}(\mathbb{P},\mathbb{Q}) + \ln(2).$$

Solving for $\mathbb{Q}(A)$ —and not for $\mathbb{P}(A)$ as was previously the case—we get

$$\mathbb{Q}(A) \ge \exp\left(-\frac{\mathrm{KL}(\mathbb{P},\mathbb{Q}) + \ln(2)}{\mathbb{P}(A)}\right),\tag{24}$$

where the above inequality is true even if $\mathbb{P}(A) = 0$ or $\mathrm{KL}(\mathbb{P},\mathbb{Q}) = +\infty$. More generally, for all distributions \mathbb{P}, \mathbb{Q} and all [0, 1]-valued random variables Z, we have, by Corollary 4,

$$\mathbb{E}_{\mathbb{Q}}[Z] \ge \exp\left(-\frac{\mathrm{KL}(\mathbb{P},\mathbb{Q}) + \ln(2)}{\mathbb{E}_{\mathbb{P}}[Z]}\right),\tag{25}$$

where again the above inequality is true even if $\mathbb{E}_{\mathbb{P}}[Z] = 0$ or $\mathrm{KL}(\mathbb{P}, \mathbb{Q}) = +\infty$.

The bound (24) is similar in spirit to (a consequence of) the Bretagnolle-Huber inequality, recalled and actually improved in Section 7.3; see details therein, and in particular its consequence (43). Both bounds can indeed be useful when $KL(\mathbb{P}, \mathbb{Q})$ is larger than a constant and $\mathbb{P}(A)$ is close to 1.

Next we show two applications of (24) and (25): a simple proof of a large deviation lower bound for Bernoulli distributions, and a distribution-dependent posterior concentration lower bound.

5.1. A simple proof of Cramér's theorem for Bernoulli distributions

The next proposition is a well-known large deviation result on the sample mean of independent and identically distributed Bernoulli random variables. It is a particular case of Cramér's theorem that dates back to Cramér [1938], Chernoff [1952]; see also Cerf and Petit [2011] for further references and a proof in a very general context. Thanks to Fano's inequality (24), the proof of the lower bound that we provide below avoids any explicit change of measure (see the remark after the proof). We are grateful to Aurélien Garivier for suggesting this proof technique to us.

Proposition 11 (Cramér's theorem for Bernoulli distributions). Let $\theta \in (0, 1)$. Assume that X_1, \ldots, X_n are independent and identically distributed random variables drawn from $\text{Ber}(\theta)$. Denoting by \mathbb{P}_{θ} the underlying probability measure, we have, for all $x \in (\theta, 1)$,

$$\lim_{n \to +\infty} \frac{1}{n} \ln \mathbb{P}_{\theta}\left(\frac{1}{n} \sum_{i=1}^{n} X_i > x\right) = -\operatorname{kl}(x, \theta) \,.$$

Proof: We set $\overline{X}_n \stackrel{\text{def}}{=} n^{-1} \sum_{i=1}^n X_i$. For the convenience of the reader we first briefly recall how to prove the upper bound, and then proceed with a new proof for the lower bound.

Upper bound: By the Cramér-Chernoff method and the duality formula for the Kullback-Leibler divergence between Bernoulli distributions (see, e.g., Boucheron et al. 2013, pages 21–24), we have, for all $n \ge 1$,

$$\mathbb{P}_{\theta}\left(\overline{X}_{n} > x\right) \leq \exp\left(-n\sup_{\lambda>0}\left\{\lambda x - \ln \mathbb{E}_{\theta}\left[e^{\lambda X_{1}}\right]\right\}\right) = \exp\left(-n\operatorname{kl}(x,\theta)\right),\tag{26}$$

that is,

$$\forall n \ge 1, \qquad \frac{1}{n} \ln \mathbb{P}_{\theta} \left(\overline{X}_n > x \right) \leqslant -\operatorname{kl}(x, \theta).$$

Lower bound: Choose $\varepsilon > 0$ small enough such that $x + \varepsilon < 1$. As in the proof of Proposition 18, we may assume with no loss of generality that the underlying distribution is $\mathbb{P}_{\theta} = \text{Ber}(\theta)^{\otimes n}$. By Fano's inequality in the form (24) with the distributions $\mathbb{P} = \mathbb{P}_{x+\varepsilon}$ and $\mathbb{Q} = \mathbb{P}_{\theta}$, and the event $A = \{\overline{X}_n > x\}$, we have

$$\mathbb{P}_{\theta}(\overline{X}_n > x) \ge \exp\left(-\frac{\mathrm{KL}(\mathbb{P}_{x+\varepsilon}, \mathbb{P}_{\theta}) + \ln(2)}{\mathbb{P}_{x+\varepsilon}(\overline{X}_n > x)}\right)$$

Noting that $\operatorname{KL}(\mathbb{P}_{x+\varepsilon},\mathbb{P}_{\theta}) = n \operatorname{kl}(x+\varepsilon,\theta)$ we get

$$\mathbb{P}_{\theta}\left(\overline{X}_{n} > x\right) \ge \exp\left(-\frac{n\,\mathrm{kl}(x+\varepsilon,\theta) + \ln 2}{\mathbb{P}_{x+\varepsilon}(\overline{X}_{n} > x)}\right) \ge \exp\left(-\frac{n\,\mathrm{kl}(x+\varepsilon,\theta) + \ln 2}{1-e^{-n\,\mathrm{kl}(x,x+\varepsilon)}}\right),\tag{27}$$

where the last bound follows from $\mathbb{P}_{x+\varepsilon}(\overline{X}_n > x) = 1 - \mathbb{P}_{x+\varepsilon}(\overline{X}_n \leq x) \ge 1 - e^{-n \operatorname{kl}(x,x+\varepsilon)}$ by a derivation similar to (26) above. Taking the logarithms of both sides and letting $n \to +\infty$ finally yields

$$\liminf_{n \to +\infty} \frac{1}{n} \ln \mathbb{P}_{\theta} \left(\overline{X}_n > x \right) \ge - \operatorname{kl}(x + \varepsilon, \theta) \,.$$

We conclude the proof by letting $\varepsilon \to 0$, and by combining the upper and lower bounds.

Comparison with an historical proof. A classical proof for the lower bound relies on the same change of measure as the one used above, i.e., that transports the measure $\text{Ber}(\theta)^{\otimes n}$ to $\text{Ber}(x+\varepsilon)^{\otimes n}$. The bound (26), or any other large deviation inequality, is also typically used therein. However, the change of measure is usually carried out explicitly by writing

$$\mathbb{P}_{\theta}\left(\overline{X}_{n} > x\right) = \mathbb{E}_{\theta}\left[\mathbb{1}_{\left\{\overline{X}_{n} > x\right\}}\right] = \mathbb{E}_{x+\varepsilon}\left[\mathbb{1}_{\left\{\overline{X}_{n} > x\right\}} \frac{\mathrm{d}\mathbb{P}_{\theta}}{\mathrm{d}\mathbb{P}_{x+\varepsilon}}(X_{1}, \dots, X_{n})\right] = \mathbb{E}_{x+\varepsilon}\left[\mathbb{1}_{\left\{\overline{X}_{n} > x\right\}}e^{-n\,\widehat{\mathrm{KL}}_{n}}\right],$$

where the empirical Kullback-Leibler divergence $\widehat{\mathrm{KL}}_n$ is defined by

$$\widehat{\mathrm{KL}}_n \stackrel{\text{def}}{=} \frac{1}{n} \ln\left(\frac{\mathrm{d}\mathbb{P}_{x+\varepsilon}}{\mathrm{d}\mathbb{P}_{\theta}}(X_1,\ldots,X_n)\right) = \frac{1}{n} \sum_{i=1}^n \left(\mathbbm{1}_{\{X_i=1\}} \ln\left(\frac{x+\varepsilon}{\theta}\right) + \mathbbm{1}_{\{X_i=0\}} \ln\left(\frac{1-(x+\varepsilon)}{1-\theta}\right)\right).$$

The empirical Kullback-Leibler divergence $\widehat{\mathrm{KL}}_n$ is then compared to its limit $\mathrm{kl}(x + \varepsilon, \theta)$ via the law of large numbers. On the contrary, our short proof above bypasses any call to the law of large numbers and does not perform the change of measure explicitly, in the same spirit as for the bandit lower bounds derived by Kaufmann et al. [2016] and Garivier et al. [2017]. Note that the different and more general proof of Cerf and Petit [2011] also bypassed any call to the law of large numbers thanks to other convex duality arguments.

5.2. Distribution-dependent posterior concentration lower bounds

In this section we consider the same Bayesian setting as the one described at the beginning of Section 4.1. In addition, we define the global modulus of continuity between KL and ℓ around $\theta \in \Theta$ and at scale $\varepsilon_n > 0$ by

$$\psi(\varepsilon_n, \theta, \ell) \stackrel{\text{def}}{=} \inf \left\{ \text{KL}(P_{\theta'}, P_{\theta}) : \ell(\theta', \theta) \ge 2\varepsilon_n, \ \theta' \in \Theta \right\};$$

the infimum is set to $+\infty$ if the set is empty.

Next we provide a distribution-dependent lower bound for posterior concentration rates, that is, a lower bound that holds true for every $\theta \in \Theta$, as opposed to the minimax lower bound of Section 4.1. Note however that we are here in a slightly different regime than in Section 4.1, where we addressed cases for which the uniform posterior concentration condition (29) below was proved to be impossible at scale ε_n (and actually took place at a slightly larger scale ε'_n).

Theorem 12 (Distribution-dependent posterior concentration lower bound). Assume that the posterior distribution $\mathbb{P}_{\pi}(\cdot | X_{1:n})$ satisfies the uniform concentration condition

$$\inf_{\theta \in \Theta} \mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi} \big(\theta' : \ell(\theta', \theta) < \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \longrightarrow 1 \qquad as \ n \to +\infty.$$

Then, for all $\theta \in \Theta$ and c > 1, for all n large enough,

$$\mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi} \big(\theta' : \ell(\theta', \theta) > \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \geqslant 2^{-c} \, \exp \Big(-c \, n \, \psi(\varepsilon_n, \theta, \ell) \Big) \,. \tag{28}$$

The conclusion can be stated equivalently as: for all $\theta \in \Theta$,

$$\liminf_{n \to +\infty} \frac{\ln\left(\mathbb{E}_{\theta}\left[\mathbb{P}_{\pi}\left(\theta' : \ell(\theta', \theta) > \varepsilon_{n} | X_{1:n}\right)\right]\right)}{\ln(2) + n \,\psi(\varepsilon_{n}, \theta, \ell)} \ge -1.$$

The above theorem is greatly inspired from Theorem 2.1 by Hoffmann et al. [2015]. Our Fano's inequality (25) however makes the proof more direct: the change-of-measure carried out by Hoffmann et al. [2015] is now implicit, and no proof by contradiction is required. We also bypass one technical assumption (see the discussion after the proof).

Proof: We fix $\theta \in \Theta$ and c > 1. By the uniform concentration condition, there exists $n_0 \ge 1$ such that, for all $n \ge n_0$,

$$\inf_{\theta^{\star} \in \Theta} \mathbb{E}_{\theta^{\star}} \Big[\mathbb{P}_{\pi} \big(\theta' : \ell(\theta', \theta^{\star}) < \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \geqslant \frac{1}{c} \,.$$
⁽²⁹⁾

We now fix $n \ge n_0$ and consider any $\theta^* \in \Theta$ such that $\ell(\theta^*, \theta) \ge 2\varepsilon_n$. Using Fano's inequality in the form of (25) with the distributions $\mathbb{P} = P_{\theta^*}^{\otimes n}$ and $\mathbb{Q} = P_{\theta}^{\otimes n}$, together with the [0, 1]-valued random variable $Z_{\theta} = \mathbb{P}_{\pi}(\theta' : \ell(\theta', \theta) > \varepsilon_n | X_{1:n})$, we get

$$\mathbb{E}_{\theta}\left[Z_{\theta}\right] \ge \exp\left(-\frac{\mathrm{KL}\left(P_{\theta^{\star}}^{\otimes n}, P_{\theta}^{\otimes n}\right) + \ln 2}{\mathbb{E}_{\theta^{\star}}\left[Z_{\theta}\right]}\right) = \exp\left(-\frac{n\,\mathrm{KL}\left(P_{\theta^{\star}}, P_{\theta}\right) + \ln 2}{\mathbb{E}_{\theta^{\star}}\left[Z_{\theta}\right]}\right).\tag{30}$$

By the triangle inequality and the assumption $\ell(\theta^*, \theta) \ge 2\varepsilon_n$ we can see that $\{\theta' : \ell(\theta', \theta) > \varepsilon_n\} \supseteq \{\theta' : \ell(\theta', \theta^*) < \varepsilon_n\}$, so that

$$\mathbb{E}_{\theta^{\star}} \left[Z_{\theta} \right] \geqslant \mathbb{E}_{\theta^{\star}} \left[\mathbb{P}_{\pi} \left(\theta' : \ell(\theta', \theta^{\star}) < \varepsilon_n \big| X_{1:n} \right) \right] \geqslant \frac{1}{c}$$

by the uniform lower bound (29). Substituting the above inequality into (30) then yields

$$\mathbb{E}_{\theta}[Z_{\theta}] \ge \exp\left(-c\left(n\operatorname{KL}(P_{\theta^{\star}}, P_{\theta}) + \ln 2\right)\right).$$

To conclude the proof, it suffices to take the supremum of the right-hand side over all $\theta^* \in \Theta$ such that $\ell(\theta^*, \theta) \ge 2\varepsilon_n$, and to identify the definition of $\psi(\varepsilon_n, \theta, \ell)$.

Note that, at first sight, our result may seem a little weaker than Hoffmann et al. [2015, Theorem 2.1], because we only define $\psi(\varepsilon_n, \theta, \ell)$ in terms of KL instead of a general pre-metric d: in other words, we only consider the case $d(\theta, \theta') = \sqrt{\text{KL}(P_{\theta'}, P_{\theta})}$. However, it is still possible to derive a bound in terms of an arbitrary pre-metric d by comparing d and KL after applying Theorem 12.

In the case of the pre-metric $d(\theta, \theta') = \sqrt{\text{KL}(P_{\theta'}, P_{\theta})}$, we bypass an additional technical assumption used for the the similar lower bound of Hoffmann et al. [2015, Theorem 2.1]; namely, that there exists a constant C > 0 such that

$$\sup_{\theta,\theta'} P_{\theta'}^{\otimes n} \Big(\mathcal{L}_n(\theta') - \mathcal{L}_n(\theta) \ge Cn \operatorname{KL}(P_{\theta'}, P_{\theta}) \Big) \longrightarrow 0 \quad \text{as } n \to +\infty,$$

where the supremum is over all $\theta, \theta' \in \Theta$ satisfying $\psi(\varepsilon_n, \theta, \ell) \leq \operatorname{KL}(P_{\theta'}, P_{\theta}) \leq 2\psi(\varepsilon_n, \theta, \ell)$, and where $\mathcal{L}_n(\theta) = \sum_{i=1}^n \ln(dP_{\theta}/d\mathfrak{m})(X_i)$ denotes the log-likelihood function with respect to a common dominating measure \mathfrak{m} . Besides, we get an improved constant in the exponential in (28), with respect to Hoffmann et al. [2015, Theorem 2.1]: by a factor of 3C/c, which, since $C \geq 1$ in most cases, is $3C/c \approx 3C \geq 3$ when $c \approx 1$. (A closer look at their proof can yield a constant arbitrarily close to 2C, which is still larger than our c by a factor of $2C/c \approx 2C \geq 2$.)

6. References and comparison to the literature

We discuss in this section how novel (or not novel) our results and approaches are.

Main innovations. We could find no reference indicating that the alternative distributions \mathbb{Q}_i and \mathbb{Q}_{θ} could vary and do not need to be set to a fixed alternative \mathbb{Q}_0 , nor that arbitrary [0,1]valued random variables Z_i or Z_{θ} could be considered. In particular, to the best of our knowledge, reduction (9) is a new result. We provide two novel applications with [0,1]-valued random variables in Section 4.

Also, as we discuss below in detail when referring to the work of Birgé [2005], results like Lemma 5 provide an interpolation between the most classical versions of Fano's inequality with a $\ln(2)$ factor and Pinsker's inequality. Typically, depending on $N \ge 3$ or N = 2, one or the other lemma had to be used, while Lemma 5 can be used in all cases.

What on the contrary was already known. The inequalities (10) are folklore knowledge. The first inequality in (11) can be found in Guntuboyina [2011]; the second inequality is a new (immediate) consequence. The first inequality in (13) is a consequence, which we derived on our own, of a refined Pinsker's inequality stated by Ordentlich and Weinberger [2005], while the second inequality is ours again.

Reduction (9) is new, as we indicated, but all other reductions were known, though sometimes proved in a more involved way. Reduction (2) and (6) were already known and used by Han and Verdú [1994, Theorems 2, 7 and 8]. Reduction (7) is stated in spirit by Chen et al. [2016] with a constant alternative $\mathbb{Q}_{\theta} \equiv \mathbb{Q}$; see also a detailed discussion and comparison below between their approach and the general approach we took in Section 3. We should also mention that Duchi and Wainwright [2013] provided preliminary (though more involved) results towards the continuous reduction (7). Reduction (8) is stated in a special case in Gushchin [2003], where $Z_1 + \ldots + Z_N = 1$. We also note that while we only discussed Kullback-Leibler divergences so far, all reductions (2) and (6)–(9) extend to f-divergences, as noted already by Gushchin [2003], see also Chen et al. [2016]. We state this extension to f-divergences in Section E of the appendix.

That the sets A_i considered in the reductions (2) and (6) form a partition of the underlying measurable space or that the random variables Z_i sum up to 1 in (8) were typical requirements in the literature until recently. Chen et al. [2016] noted in spirit that the requirement of forming a partition was unnecessary, which we too had been aware of as early as Stoltz [2007], where we also already mentioned the fact that in particular the alternative distribution \mathbb{Q} had not to be fixed and could depend on i or θ .

Finally, the conjunction of a reduction (2) or (6)–(9) and a lower bound on the kl function was already present in Han and Verdú [1994]. Other, more information-theoretic statements and proof techniques of Fano's inequalities for finitely many hypotheses as in Proposition 1 can be found, e.g., in Cover and Thomas [2006, Theorem 2.11.1], Yu [1997, Lemma 3] or Ibragimov and Has'minskii [1981, Chapter VII, Lemma 1.1] (they resort to classical formulas on the Shannon entropy, the conditional entropy, and the mutual information).

6.1. On the "generalized Fanos's inequality" of Chen et al. [2016]

The Bayesian setting considered is the following; it generalizes the setting of Han and Verdú [1994], whose results we discuss in a remark after the proof of Proposition 13.

A parameter space (Θ, \mathcal{G}) is equipped with a prior probability measure ν . A family of probability distributions $(\mathbb{P}_{\theta})_{\theta \in \Theta}$ over a measurable space (Ω, \mathcal{F}) , some outcome space $(\mathcal{X}, \mathcal{E})$, e.g., $\mathcal{X} = \mathbb{R}^n$, and a random variable $X : (\Omega, \mathcal{F}) \to (\mathcal{X}, \mathcal{E})$ are considered. We denote by \mathbb{E}_{θ} the expectation under \mathbb{P}_{θ} . Of course we may have $(\Omega, \mathcal{F}) = (\mathcal{X}, \mathcal{E})$ and X be the identity, in which case \mathbb{P}_{θ} will be the law of X under \mathbb{P}_{θ} .

The goal is either to estimate θ or to take good actions: we consider a measurable target space $(\mathcal{A}, \mathcal{H})$, that may or may not be equal to Θ . The quality of a prediction or of an action is measured by a measurable loss function $L: \Theta \times \mathcal{A} \to [0, 1]$. The random variable X is our observation, based

on which we construct a $\sigma(X)$ -measurable random variable \hat{a} with values in \mathcal{A} . Putting as side all measurability issues (here and in the rest of this subsection), the risk of \hat{a} in this model equals

$$R(\hat{a}) = \int_{\Theta} \mathbb{E}_{\theta} \Big[L(\theta, \hat{a}) \Big] \, \mathrm{d}\nu(\theta)$$

and the Bayes risk in this model is the smallest such possible risk,

$$R_{\text{Bayes}} = \inf_{\widehat{a}} R(\widehat{a}) ,$$

where the infimum is over all $\sigma(X)$ -measurable random variables with values in \mathcal{A} .

Chen et al. [2016] call their main result (Corollary 5) a "generalized Fano's inequality;" we state it below not only for $\{0, 1\}$ -valued loss functions L as in the original article but for any [0, 1]-valued loss functions, as we are able to prove it for any such loss function. The reason behind this extension is that we not only have the reduction (7) with events, but we also have the reduction (9) with [0, 1]-valued random variables. We also feel that our proof technique is more direct and more natural.

We only deal with with Kullback-Leibler divergences, but the result and proof below readily extend to f-divergences.

Proposition 13. In the setting described above, the Bayes risk is always larger than

$$R_{\text{Bayes}} \ge 1 + \frac{\left(\inf_{\mathbb{Q}} \int_{\Theta} \text{KL}(\mathbb{P}_{\theta}, \mathbb{Q}) \, \mathrm{d}\nu(\theta)\right) + \ln\left(1 + \inf_{a \in \mathcal{A}} \int_{\Theta} L(\theta, a) \, \mathrm{d}\nu(\theta)\right)}{\ln\left(1 - \inf_{a \in \mathcal{A}} \int_{\Theta} L(\theta, a) \, \mathrm{d}\nu(\theta)\right)},$$

where the infimum in the numerator is over all probability measures \mathbb{Q} over (Ω, \mathcal{F}) .

Proof: We fix \hat{a} and an alternative \mathbb{Q} . The combination of (9) and (11), with $Z_{\theta} = 1 - L(\theta, \hat{a})$, yields

$$1 - \int_{\Theta} \mathbb{E}_{\theta} \left[L(\theta, \widehat{a}) \right] d\nu(\theta) \leqslant \frac{\int_{\Theta} \mathrm{KL}(\mathbb{P}_{\theta}, \mathbb{Q}) \, d\nu(\theta) + \ln\left(2 - q_{\widehat{a}}\right)}{\ln\left(1/q_{\widehat{a}}\right)}, \tag{31}$$

where $\mathbb{E}_{\mathbb{Q}}$ denotes the expectation with respect to \mathbb{Q} and

$$q_{\widehat{a}} = 1 - \int_{\Theta} \mathbb{E}_{\mathbb{Q}} \left[L(\theta, \widehat{a}) \right] d\nu(\theta) \,.$$

As $q \mapsto 1/\ln(1/q)$ and $q \mapsto \ln(2-q)/\ln(1/q)$ are both increasing, taking the supremum over the $\sigma(X)$ -measurable random variables \hat{a} in both sides of (31) gives

$$1 - R_{\text{Bayes}} \leqslant \frac{\int_{\Theta} \text{KL}(\mathbb{P}_{\theta}, \mathbb{Q}) \, \mathrm{d}\nu(\theta) + \ln(2 - q^{\star})}{\ln(1/q^{\star})}$$
(32)

where

$$q^{\star} = \sup_{\widehat{a}} q_{\widehat{a}} = 1 - \inf_{\widehat{a}} \int_{\Theta} \mathbb{E}_{\mathbb{Q}} \Big[L(\theta, \widehat{a}) \Big] \, \mathrm{d}\nu(\theta) = 1 - \inf_{a \in \mathcal{A}} \int_{\Theta} L(\theta, a) \, \mathrm{d}\nu(\theta) \,, \tag{33}$$

as is proved below. Taking the infimum of the right-hand side of (32) over \mathbb{Q} and rearranging concludes the proof.

It only remains to prove the last inequality of (33) and actually, as constant elements $a \in \mathcal{A}$ are special cases of random variables \hat{a} , we only need to prove that

$$\inf_{\widehat{a}} \int_{\Theta} \mathbb{E}_{\mathbb{Q}} \Big[L(\theta, \widehat{a}) \Big] \, \mathrm{d}\nu(\theta) \ge \inf_{a \in \mathcal{A}} \int_{\Theta} L(\theta, a) \, \mathrm{d}\nu(\theta) \,.$$
(34)

Now, each \hat{a} that is $\sigma(X)$ -measurable can be rewritten $\hat{a} = \bar{a}(X)$ for some measurable function $\overline{a}: \mathcal{X} \to \mathcal{A}$; then, by the Fubini-Tonelli theorem:

$$\int_{\Theta} \mathbb{E}_{\mathbb{Q}} \Big[L(\theta, \widehat{a}) \Big] \, \mathrm{d}\nu(\theta) = \int_{\mathcal{X}} \left(\int_{\Theta} L(\theta, \overline{a}(x)) \, \mathrm{d}\nu(\theta) \right) \, \mathrm{d}\mathbb{Q}(x) \ge \int_{\mathcal{X}} \left(\inf_{a \in \mathcal{A}} \int_{\Theta} L(\theta, a) \, \mathrm{d}\nu(\theta) \right) \, \mathrm{d}\mathbb{Q}(x) \,,$$

hich proves (34).

which proves (34).

Remark 2. As mentioned by Chen et al. [2016], one of the major results of Han and Verdú [1994], namely, their Theorem 8, is a special case of Proposition 13, with $\Theta = \mathcal{A}$ and the loss function $L(\theta, \theta') = \mathbb{1}_{\{\theta \neq \theta'\}}$. The (opposite of the) denominator in the lower bound on the Bayes risk then takes the simple form

$$-\ln\left(1-\inf_{\theta'\in\Theta}\int_{\Theta}L(\theta,\theta')\,\mathrm{d}\nu(\theta)\right)=\ln\left(\sup_{\theta\in\Theta}\nu(\{\theta\})\right)\stackrel{\mathrm{def}}{=}H_{\infty}(\nu)\,,$$

which is called the infinite-order Rényi entropy of the probability distribution ν . Han and Verdú [1994] only dealt with the case of discrete sets Θ but the extension to continuous Θ is immediate, as we showed in Section 3.

6.2. Comparison to Birgé [2005]

This version of Fano's inequality is extremely popular among statisticians. We state here a slightly simplified version of the main result by Birgé [2005] (his Corollary 1), inspired by a previous (looser) simplification by Massart [2007]: in Appendix C these two alternative statements are stated, proved, and compared to Theorem 14. In contrast, the proof of Theorem 14 is provided at the end of the present subsection; it of course follows the methodology described in Section 3.

The bounds by Birgé [2005] only deal with events A_1, \ldots, A_N forming a partition of the underlying measurable space. As should be clear from their proof this assumption is crucial.

Theorem 14 (Birgé's lemma). Given an underlying measurable space (Ω, \mathcal{F}) , for all $N \ge 2$, for all probability distributions $\mathbb{P}_1, \ldots, \mathbb{P}_N$, for all events A_1, \ldots, A_N forming a partition of Ω ,

$$\min_{1 \le i \le N} \mathbb{P}_i(A_i) \le \max\left\{c_N, \ \frac{\overline{K}}{\ln(N)}\right\} \qquad where \qquad \overline{K} = \frac{1}{N-1} \sum_{i=2}^N \mathrm{KL}(\mathbb{P}_i, \mathbb{P}_1)$$

and where $(c_N)_{N\geq 2}$ is a decreasing sequence, where each term c_N is defined as the unique $c \in (0,1)$ such that

$$\frac{-(c\ln(c) + (1-c)\ln(1-c))}{c} + \ln(1-c) = \ln\left(\frac{N-1}{N}\right).$$
(35)

We have, for instance, $c_2 \approx 0.7587$ and $c_3 \approx 0.7127$, while $\lim c_N = 0.63987$.

The aim of this subsection is to compare this bound to the versions of Fano's inequality following from the kl lower bounds (11), (10), and (13), in this order. In the setting of the theorem above and by picking constant alternatives \mathbb{Q} , these lower bounds on kl respectively lead to

 $\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{1}{N} + \sqrt{\frac{\frac{1}{N}\inf_{\mathbb{Q}}\sum_{i=1}^{N}\mathrm{KL}(\mathbb{P}_{i},\mathbb{Q})}{\max\{\ln(N),\,2\}}} \,.$

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{\frac{1}{N}\inf_{\mathbb{Q}}\sum_{i=1}^{N}\mathrm{KL}(\mathbb{P}_{i},\mathbb{Q}) + \ln\left(2-\frac{1}{N}\right)}{\ln(N)} \leqslant \frac{\frac{1}{N}\inf_{\mathbb{Q}}\sum_{i=1}^{N}\mathrm{KL}(\mathbb{P}_{i},\mathbb{Q}) + \ln(2)}{\ln(N)}, \quad (36)$$

and

The main point of Birgé [2005] was that the most classical version of Fano's inequality, that is, the right-most side of (36), was quite unpractical for small values of N, and even useless when N = 2. In the latter case N = 2, the statistical doxa had it that one should rather resort to Pinsker's inequality, which is exactly (37) when N = 2. One of his main motivation was therefore to get an inequality that would be useful for all $N \ge 2$, so that one does not have to decide which of the classical Pinsker's inequality or the classical Fano's inequality should be applied. A drawback, however, of his bound is the \overline{K} term, in which one cannot pick a convenient \mathbb{Q} as in the bounds (36)–(37). Also, the result is about the minimum of the $\mathbb{P}_i(A_i)$, not about their average.

Now, we note that both the middle term in (36) and the bound (37) yield useful bounds, even for N = 2. The middle term in (36) was derived—with a different formulation—by Chen et al. [2016], see Proposition 13 above. Our contribution is to note that our inequality (37) provides an interpolation between Pinsker's and Fano's inequalities. More precisely, (37) implies both Pinsker's inequality and, lower bounding the maximum by $\ln(N)$, a bound as useful as Theorem 14. Indeed, in practice, the additional additive 1/N term and the additional square root do not prevent from obtaining the desired lower bounds, as illustrated in Section 4.2.

We close this subsection with a proof of Theorem 14.

Proof (of Theorem 14): We denote by $h: p \in [0,1] \mapsto -(p \ln(p) + (1-p) \ln(1-p))$ the binary entropy function. The existence of c_N follows from the fact that $c \in (0,1) \mapsto h(c)/c + \ln(1-c)$ is continuous and decreasing, as the sum of two such functions; its respective limits are $+\infty$ and $-\infty$ at 0 and 1.

Reduction (2) with $\mathbb{Q}_i = \mathbb{P}_1$ for all $i \ge 2$ indicates that $\mathrm{kl}(\tilde{p}, \tilde{q}) \leqslant \overline{K}$ where

$$\widetilde{p} \stackrel{\text{def}}{=} \frac{1}{N-1} \sum_{i=2}^{N} \mathbb{P}_i(A_i), \qquad \widetilde{q} \stackrel{\text{def}}{=} \frac{1}{N-1} \sum_{i=2}^{N} \mathbb{P}_1(A_i) = \frac{1-\mathbb{P}_1(A_1)}{N-1}, \qquad \overline{K} = \frac{1}{N-1} \sum_{i=2}^{N} \text{KL}(\mathbb{P}_i, \mathbb{P}_1);$$

note that we used the assumption of a partition to get the alternative definition of the \tilde{q} quantity. We use the following lower bound on kl, which follows from calculations similar to the ones performed in (4), using that $c_N \ge 1/2$ and that the binary entropy $h: p \mapsto -(p \ln(p) + (1-p) \ln(1-p))$ is decreasing on [1/2, 1]: for $p \ge c_N$,

$$\operatorname{kl}(p,q) \ge p \ln\left(\frac{1}{q}\right) - h(c_N) \ge p \ln\left(\frac{1}{q}\right) - p \frac{h(c_N)}{c_N},$$

where $\ln(1/q) - h(c_N)/c_N > 0$ for $q < \exp(-h(c_N)/c_N)$. Hence,

$$\forall p \in [0,1], \quad \forall q \in \left(0, \exp\left(-h(c_N)/c_N\right)\right), \qquad p \leqslant \max\left\{c_N, \frac{\mathrm{kl}(p,q)}{\ln(1/q) - h(c_N)/c_N}\right\}.$$
(38)

Now, we set $a = \min_{1 \le i \le N} \mathbb{P}_i(A_i)$ and may assume $a \ge c_N$ (otherwise, the stated bound is obtained).

(37)

We have, by the very definition of a as a minimum and by the definition (35) of c_N ,

$$a \leqslant \widetilde{p}$$
 and $\widetilde{q} \leqslant \frac{1-a}{N-1} \leqslant \frac{1-c_N}{N-1} = \frac{1}{N} \exp\left(-\frac{h(c_N)}{c_N}\right),$ (39)

while $\tilde{q} > 0$ unless $\mathbb{P}_1(A_1) = a = 0$, in which case there is nothing to prove. We may therefore combine $\mathrm{kl}(\tilde{p},\tilde{q}) \leq \overline{K}$ with (38) to get

$$a \leq \widetilde{p} \leq \max\left\{c_N, \frac{\operatorname{kl}(\widetilde{p}, \widetilde{q})}{\ln(1/\widetilde{q}) - h(c_N)/c_N}\right\} \leq \max\left\{c_N, \frac{\overline{K}}{\ln(N)}\right\},\$$

where, for the last inequality, we used the upper bound on \tilde{q} in (39).

7. Proofs of the stated lower bounds on kl (and of an improved Bretagnolle-Huber inequality)

We prove in this section the convexity inequalities (11) and (12) as well as the refined Pinsker's inequality and its consequence (13). Using the same techniques and methodology as for establishing these bounds, we also improve in passing the Bretagnolle-Huber inequality.

7.1. Proofs of the convexity inequalities (11) and (12)

Proof: Inequality (12) follows from (11) via a function study of $q \in (0, 1) \mapsto \ln(2-q)/\ln(1/q)$, which is dominated by $0.21 + 0.79 \,\mathrm{q}$.

Now, the shortest proof of (11) notes that the duality formula for the Kullback-Leibler divergence between Bernoulli distributions—already used in (26)—ensures that, for all $p \in [0, 1]$ and $q \in (0, 1]$,

$$\operatorname{kl}(p,q) = \sup_{\lambda \in \mathbb{R}} \left\{ \lambda p - \ln\left(q\left(e^{\lambda} - 1\right) + 1\right) \right\} \ge p \ln\left(\frac{1}{q}\right) - \ln(2 - q)$$

for the choice $\lambda = \ln(1/q)$.

An alternative, longer but more elementary proof uses a direct convexity argument, as in Guntuboyina [2011, Example II.4], which already included the inequality of interest in the special case when q = 1/N; see also Chen et al. [2016]. We deal separately with p = 0 and p = 1, and thus restrict our attention to $p \in (0, 1)$ in the sequel. For $q \in (0, 1)$, as $p \mapsto kl(p, q)$ is convex and differentiable on (0, 1), we have

$$\forall (p, p_0) \in (0, 1)^2, \qquad kl(p, q) - kl(p_0, q) \ge \underbrace{\ln\left(\frac{p_0(1-q)}{(1-p_0)q}\right)}_{\frac{\partial}{\partial p} kl(p_0, q)} (p-p_0). \tag{40}$$

The choice $p_0 = 1/(2-q)$ is such that

$$\frac{p_0}{1-p_0} = \frac{1}{1-q}$$
, thus $\ln\left(\frac{p_0(1-q)}{(1-p_0)q}\right) = \ln\left(\frac{1}{q}\right)$,

and

$$kl(p_0,q) = \frac{1}{2-q} \ln\left(\frac{1/(2-q)}{q}\right) + \frac{1-q}{2-q} \ln\left(\frac{(1-q)/(2-q)}{1-q}\right) = \frac{1}{2-q} \ln\left(\frac{1}{q}\right) + \ln\left(\frac{1}{2-q}\right).$$

Inequality (40) becomes

$$\forall p \in (0,1), \qquad \operatorname{kl}(p,q) - \frac{1}{2-q} \ln\left(\frac{1}{q}\right) + \ln(2-q) \ge \left(p - \frac{1}{2-q}\right) \ln\left(\frac{1}{q}\right),$$

which proves as well the bound (11).

7.2. Proofs of the refined Pinsker's inequality and of its consequence (13)

The next theorem is a stronger version of Pinsker's inequality for Bernoulli distributions, that was proved² by Ordentlich and Weinberger [2005]. Indeed, note that the function φ defined below satisfies min $\varphi = 2$, so that the next theorem always yields an improvement over the most classical version of Pinsker's inequality: $kl(p,q) \ge 2(p-q)^2$.

We provide below an alternative elementary proof for Bernoulli distributions of this refined Pinsker's inequality. The extension to the case of general distributions, via the contraction-of-entropy property, is stated at the end of this section.

Theorem 15 (A refined Pinsker's inequality by Ordentlich and Weinberger [2005]). For all $p, q \in [0, 1]$,

$$\operatorname{kl}(p,q) \ge \frac{\ln((1-q)/q)}{1-2q} \, (p-q)^2 \stackrel{\text{def}}{=} \varphi(q) \, (p-q)^2 \,,$$

where the multiplicative factor $\varphi(q) = (1-2q)^{-1} \ln((1-q)/q)$ is defined for all $q \in [0,1]$ by extending it by continuity as $\varphi(1/2) = 2$ and $\varphi(0) = \varphi(1) = +\infty$.

The proof shows that $\varphi(q)$ is the optimal multiplicative factor in front of $(p-q)^2$ when the bounds needs to hold for all $p \in [0, 1]$; the proof also provides a natural explanation for the value of φ .

Proof: The stated inequality is satisfied for $q \in \{0, 1\}$ as $kl(p,q) = +\infty$ in these cases unless p = q. The special case q = 1/2 is addressed at the end of the proof. We thus fix $q \in (0,1) \setminus \{1/2\}$ and set $f(p) = kl(p,q)/(p-q)^2$ for $p \neq q$, with a continuity extension at p = q. We exactly show that f attains its minimum at p = 1 - q, from which the result (and its optimality) follow by noting that

$$f(1-q) = \frac{\mathrm{kl}(1-q,q)}{(1-2q)^2} = \frac{\mathrm{ln}((1-q)/q)}{1-2q} = \varphi(q) \,.$$

Given the form of f, it is natural to perform a second-order Taylor expansion of kl(p,q) around q. We have

$$\frac{\partial}{\partial p}\operatorname{kl}(p,q) = \ln\left(\frac{p(1-q)}{(1-p)q}\right) \quad \text{and} \quad \frac{\partial^2}{\partial^2 p}\operatorname{kl}(p,q) = \frac{1}{p(1-p)} \stackrel{\text{def}}{=} \psi(p), \quad (41)$$

so that Taylor's formula with integral remainder reveals that for $p \neq q$,

$$f(p) = \frac{\mathrm{kl}(p,q)}{(p-q)^2} = \frac{1}{(p-q)^2} \int_q^p \frac{\psi(t)}{1!} (p-t)^1 \,\mathrm{d}t = \int_0^1 \psi(q+u(p-q))(1-u) \,\mathrm{d}u \,.$$

This rewriting of f shows that f is strictly convex (as ψ is so). Its global minimum is achieved at the unique point where its derivative vanishes. But by differentiating under the integral sign, we have, at p = 1 - q,

$$f'(1-q) = \int_0^1 \psi'(q+u(1-2q)) u(1-u) \, \mathrm{d}u = 0;$$

the equality to 0 follows from the fact that the function $u \mapsto \psi'(q+u(1-2q))u(1-u)$ is antisymmetric around u = 1/2 (essentially because ψ' is antisymmetric itself around 1/2). As a consequence, the convex function f attains its global minimum at 1-q, which concludes the proof for the case where $q \in (0,1) \setminus \{1/2\}$.

It only remains to deal with q = 1/2: we use the continuity of kl (p, \cdot) and φ to extend the obtained inequality from $q \in [0, 1] \setminus \{1/2\}$ to q = 1/2.

We now prove the second inequality of (13). A picture is helpful, see Figure 1.

 $^{^{2}}$ We also refer the reader to Kearns and Saul [1998, Lemma 1] and Berend and Kontorovich [2013, Theorem 3.2] for dual inequalities upper bounding the moment-generating function of the Bernoulli distributions.



Figure 1: Plots of φ [left] and $x \in (0, 1) \mapsto \varphi(x) - \ln(1/x)$ [right].

Corollary 16. For all $q \in (0,1]$, we have $\varphi(q) \ge 2$ and $\varphi(q) \ge \ln(1/q)$. Thus, for all $p \in [0,1]$ and $q \in (0,1)$,

$$p \leqslant q + \sqrt{\frac{\mathrm{kl}(p,q)}{\max\{\ln(1/q),\,2\}}}$$

Slightly sharper bounds are possible, like $\varphi(q) \ge (1+q)(1+q^2)\ln(1/q)$ or $\varphi(q) \ge \ln(1/q) + 2.5q$, but we were unable to exploit these refinements in our applications.

General refined Pinsker's inequality. The following result, which improves on Pinsker's inequality, is due to Ordentlich and Weinberger [2005]. Our approach through Bernoulli distributions enables to derive it in an elementary (and enlightening) way: by combining Theorem 15 and the data-processing inequality (Lemma 2).

Theorem 17. Let \mathbb{P} and \mathbb{Q} be two probability distributions over the same measurable space (Ω, \mathcal{F}) . Then

$$\sup_{A \in \mathcal{F}} \left| \mathbb{P}(A) - \mathbb{Q}(A) \right| \leqslant \sqrt{\frac{\mathrm{KL}(\mathbb{P}, \mathbb{Q})}{\inf_{A \in \mathcal{F}} \varphi(\mathbb{Q}(A))}},$$

where $\varphi \ge 2$ is defined in the statement of Theorem 15.

7.3. An improved Bretagnolle-Huber inequality

The Bretagnolle-Huber inequality was introduced by Bretagnolle and Huber [1978, 1979]. The multiplicative factor $e^{-1/e} \ge 0.69$ in our statement (42) below is a slight improvement over the original 1/2 factor. For all $p, q \in [0, 1]$,

$$1 - |p - q| \ge e^{-1/e} e^{-kl(p,q)}, \quad \text{thus} \quad q \ge p - 1 + e^{-1/e} e^{-kl(p,q)}.$$
(42)

It is worth to note that Bretagnolle and Huber [1978] also proved the inequality

$$|p-q| \leqslant \sqrt{1 - \exp(-\operatorname{kl}(p,q))}$$

which improves as well upon the Bretagnolle-Huber inequality with the 1/2 factor, but which is neither better nor worse than (42).

Now, via the data-processing inequality (Lemma 2), we get from (42)

$$1 - \sup_{A \in \mathcal{F}} \left| \mathbb{P}(A) - \mathbb{Q}(A) \right| \ge e^{-1/e} e^{-KL(\mathbb{P},\mathbb{Q})}.$$

The left-hand side can be rewritten as $\inf_{A \in \mathcal{F}} \{\mathbb{P}(A) + \mathbb{Q}(A^c)\}$, where A^c denotes the complement of A. Therefore, the above inequality is a lower bound on the test affinity between \mathbb{P} and \mathbb{Q} . For the sake of comparison to (24), we can restate the general version of the Bretagnolle-Huber inequality as: for all $A \in \mathcal{F}$,

$$\mathbb{Q}(A) \ge \mathbb{P}(A) - 1 + e^{-1/e} e^{-\operatorname{KL}(\mathbb{P},\mathbb{Q})}.$$
(43)

We now provide a proof of (42); note that our improvement was made possible because we reduced the proof to very elementary arguments in the case of Bernoulli distributions.

Proof: The case where $p \in \{0, 1\}$ or $q \in \{0, 1\}$ can be handled separately; we consider $(p, q) \in (0, 1)^2$ in the sequel. The derivative of the function $x \in (0, 1) \mapsto x \ln(x/(1-q))$ equals $1 + \ln(x) - \ln(1-q)$, so that the function achieves its minimum at x = (1-q)/e, with value $-(1-q)/e \ge -1/e$. Therefore,

$$-\operatorname{kl}(p,q) = -p\ln\left(\frac{p}{q}\right) - (1-p)\ln\left(\frac{1-p}{1-q}\right) \leqslant -p\ln\left(\frac{p}{q}\right) + \frac{1}{e} = p\left(\ln\left(\frac{q}{p}\right) + \frac{1}{e}\right) + (1-p)\frac{1}{e}$$

Therefore, using the convexity of the exponential,

$$e^{-kl(p,q)} \leq p \exp\left(\ln\left(\frac{q}{p}\right) + \frac{1}{e}\right) + (1-p)e^{1/e} = (q + (1-p))e^{1/e},$$

which shows that

$$1 - (p - q) \ge e^{-1/e} e^{-kl(p,q)}$$
.

By replacing q by 1 - q and p by 1 - p, we also get

$$1 - (q - p) = 1 - ((1 - p) - (1 - q)) \ge e^{-1/e} e^{-kl(1 - p, 1 - q)} = e^{-1/e} e^{-kl(p,q)}.$$

This concludes the proof, as 1-|p-q| is equal to the smallest value between 1-(p-q) and 1-(q-p).

References

- S.M. Ali and S.D. Silvey. A general class of coefficients of divergence of one distribution from another. Journal of the Royal Statistical Society. Series B. Methodological, 28:131–142, 1966.
- D. Berend and A. Kontorovich. On the concentration of the missing mass. *Electronic Communications in Probability*, 18(3):1–7, 2013.
- L. Birgé. A new lower bound for multiple hypothesis testing. *IEEE Transactions on Information Theory*, 51(4):1611–1615, 2005.
- S. Boucheron, G. Lugosi, and P. Massart. Concentration Inequalities: A Nonasymptotic Theory of Independence. Oxford University Press, 2013.
- J. Bretagnolle and C. Huber. Estimation des densités : risque minimax. Séminaire de Probabilités de Strasbourg, 12:342–363, 1978.
- J. Bretagnolle and C. Huber. Estimation des densités : risque minimax. Zeitschrift f
 ür Wahrscheinlichkeitstheorie und verwandte Gebiete, 47(2):119–137, 1979.

- R. Cerf and P. Petit. A short proof of Cramér's theorem in R. *The American Mathematical Monthly*, 118(10):925–931, 2011.
- N. Cesa-Bianchi and G. Lugosi. Prediction, Learning, and Games. Cambridge University Press, 2006.
- N. Cesa-Bianchi, Y. Freund, D. Haussler, D.P. Helmbold, R. Schapire, and M. Warmuth. How to use expert advice. *Journal of the ACM*, 44(3):427–485, 1997.
- N. Cesa-Bianchi, G. Lugosi, and G. Stoltz. Minimizing regret with label-efficient prediction. *IEEE Transactions on Information Theory*, 51:2152–2162, 2005.
- X. Chen, A. Guntuboyina, and Y. Zhang. On Bayes risk lower bounds. Journal of Machine Learning Research, 17(219):1–58, 2016.
- H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952.
- T.M. Cover and J.A. Thomas. *Elements of information theory*. John Wiley & Sons, second edition, 2006.
- H. Cramér. Sur un nouveau théorème limite de la théorie des probabilités. Actualites Scientifiques et Industrielles, 736:5–23, 1938.
- I. Csiszár. Eine informationstheoretische Ungleichung und ihre Anwendung auf den Beweis der Ergodizität von Markoffschen Ketten. A Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei, 8:85–108, 1963.
- J. Duchi. Lecture Notes for Statistics 311/Electrical Engineering 377, chapter 4, pages 41-48. 2014. URL https://web.stanford.edu/class/stats311/Lectures/lec-05.pdf.
- J. Duchi and M.J. Wainwright. Distance-based and continuum Fano inequalities with applications to statistical estimation. 2013. arXiv:1311.2669.
- R. Durrett. Probability: Theory and Examples. Cambridge University Press, 4th edition, 2010.
- T.S. Ferguson. *Mathematical statistics: A decision theoretic approach*. Probability and Mathematical Statistics, Vol. 1. Academic Press, New York-London, 1967.
- A. Garivier, P. Ménard, and G. Stoltz. Explore first, exploit next: the true shape of regret in bandit problems. *Mathematics of Operations Research*, 2017. To appear.
- S. Ghosal, J.K. Ghosh, and A.W. van der Vaart. Convergence rates of posterior distributions. Annals of Statistics, 28(2):500–531, 2000.
- R.M. Gray. Entropy and Information Theory. Springer, second edition, 2011.
- A. Guntuboyina. Lower bounds for the minimax risk using-divergences, and applications. *IEEE Transactions on Information Theory*, 57(4):2386–2399, 2011.
- A.A. Gushchin. On fanos lemma and similar inequalities for the minimax risk. Probability Theory and Mathematical Statistics, 67:26–37, 2003.
- L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. A Distribution-Free Theory of Nonparametric Regression. Springer Series in Statistics. Springer-Verlag, New York, 2002.
- T.S. Han and S. Verdú. Generalizing the Fano inequality. *IEEE Transactions on Information Theory*, 40(4):1247–1251, 1994.

- M. Hoffmann, J. Rousseau, and J. Schmidt-Hieber. On adaptive posterior concentration rates. Annals of Statistics, 43(5):2259–2295, 2015.
- I.A. Ibragimov and R.Z. Has'minskii. *Statistical Estimation: Asymptotic Theory*, volume 16. Springer-Verlag New York, 1981.
- I.A. Ibragimov and R.Z. Has'minskii. Bounds for the risks of non-parametric regression estimates. Theory of Probability and its Applications, 27(1):84–99, 1982.
- I.A. Ibragimov and R.Z. Has'minskii. Asymptotic bounds on the quality of the nonparametric regression estimation in L_p . Journal of Mathematical Sciences, 24(5):540–550, 1984.
- E. Kaufmann, O. Cappé, and A. Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.
- M. Kearns and L. Saul. Large deviation methods for approximate probabilistic inference. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence (UAI'98)*, pages 311–319, 1998.
- J. Kwon and V. Perchet. Gains and losses are fundamentally different in regret minimization: The sparse case. *Journal of Machine Learning Research*, 17(229):1–32, 2016.
- L. Le Cam. Asymptotic methods in statistical decision theory. Springer Series in Statistics. Springer-Verlag, New York, 1986.
- L. Le Cam and G.L. Yang. *Asymptotics in statistics: some basic concepts*. Springer Series in Statistics. Springer-Verlag, New York, second edition, 2000.
- P. Massart. Concentration Inequalities and Model Selection, volume 1896 of Lecture Notes in Mathematics. Springer, 2007.
- E. Ordentlich and M.J. Weinberger. A distribution dependent refinement of Pinsker's inequality. IEEE Transactions on Information Theory, 51(5):1836–1840, 2005.
- L. Pardo. Statistical Inference Based on Divergence Measures. Chapman & Hall/CRC, 2006.
- R.T. Rockafellar. Convex Analysis. Princeton University Press, second edition, 1972.
- G. Stoltz. An introduction to the prediction of individual sequences: (1) oracle inequalities; (2) prediction with partial monitoring, 2007. Statistics seminar of Université Paris VI and Paris VII, Chevaleret, November 12 and 26, 2007; written version of the pair of seminar talks available upon request.
- A.B. Tsybakov. Introduction to Nonparametric Estimation. Springer, 2009.
- A. Xu and M. Raginsky. Information-theoretic lower bounds on Bayes risk in decentralized estimation. 2016. arXiv:1607.00550.
- Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, 27(5):1564–1599, 1999.
- B. Yu. Assouad, Fano, and Le Cam. In D. Pollard, E. Torgersen, and G.L. Yang, editors, *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*, pages 423–435. New York, NY, 1997.

Supplementary material for the article "Fano's inequality for random variables" by Gerchinovitz, Ménard, Stoltz

This supplementary material is provided for the sake of self-completeness and is not necessarily intended for journal publication (unless the reviewers decide that parts of it are of important enough interest).

However, it will be made available on arXiv and the main body of the article will eventually rather point to this arXiv supplementary material.

A. Two toy applications of the continuous Fano's inequality

We present here two toy applications of our continuous Fano's inequality, that the unfamiliar reader may study in complement to the new applications addressed in Sections 4 and 5. The two topics covered below are:

- parametric estimation of the mean in a multivariate Gaussian model, where we use the reduction (7);
- nonparametric regression with fixed design, which also relies on the reduction (7).

A.1. Parametric minimax lower bound in the mutivariate Gaussian model

The next result is a well-known minimax lower bound on the mean-estimation problem in the standard multivariate Gaussian model. Many proof techniques were used to derive this toy lower bound. The proof we provide below illustrates how to use Fano's inequality without relying on any discretization argument, thanks to a continuous version of Fano's inequality. Proofs of the same spirit were proposed by Duchi and Wainwright [2013] and Chen et al. [2016], though some discretizations were still used at some point in both references.

See the end of this subsection for extended comments and references, in particular with respect to other well-known proof techniques not resorting to discretization arguments like the use of Assouad's lemma.

Proposition 18 (Parametric lower bound with a continuous Fano's inequality). Let $d \ge 1$ be the ambient dimension and \mathbb{R}^d be the parameter space. Assume that we observe an *n*-sample X_1, \ldots, X_n distributed according to $\mathcal{N}(\theta, \sigma^2 I_d)$ for some unknown $\theta \in \mathbb{R}^d$, where $\sigma > 0$ and where I_d is the $d \times d$ identity matrix. Then, denoting by \mathbb{E}_{θ} the expectation when the unknown parameter is θ , we have the lower bound

$$\inf_{\widehat{\theta}} \sup_{\theta \in \mathbb{R}^d} \mathbb{E}_{\theta} \Big[\left\| \widehat{\theta} - \theta \right\|_2^2 \Big] \ge c_d \frac{\sigma^2 d}{n} \,,$$

where the infimum is over all \mathbb{R}^d -valued estimators $\hat{\theta} = \hat{\theta}(X_1, \ldots, X_n)$, and where $(c_d)_{d \ge 1}$ is an increasing sequence such that $c_1 \ge 0.01$, $c_2 \ge 0.025$, and $\lim_{d \to +\infty} c_d \ge 0.05$.

The proof of this result uses similar but simpler arguments than the one of Proposition 8.

Proof: We may assume with no loss of generality that the underlying probability space is $\Omega = (\mathbb{R}^d)^n$, that each X_i is the *i*-th projection map $(x_1, \ldots, x_n) \in \Omega \mapsto x_i \in \mathbb{R}^d$, and that the collection of probability distributions over Ω is formed by the $\mathbb{P}_{\theta} = \mathcal{N}(\theta, \sigma^2 I_d)^{\otimes n}$. We still denote by \mathbb{E}_{θ} the expectation under \mathbb{P}_{θ} . Fix any estimator $\hat{\theta} = \hat{\theta}(X_1, \ldots, X_n)$ and let $\varepsilon > 0$ be determined by the analysis. By Markov's inequality for the first inequality and considering any probability distribution ν over \mathbb{R}^d for the second inequality,

$$\sup_{\theta \in \mathbb{R}^{d}} \mathbb{E}_{\theta} \Big[\left\| \widehat{\theta} - \theta \right\|_{2}^{2} \Big] \ge \sup_{\theta \in \mathbb{R}^{d}} \varepsilon^{2} \mathbb{P}_{\theta} \Big(\left\| \widehat{\theta} - \theta \right\|_{2}^{2} > \varepsilon^{2} \Big) = \varepsilon^{2} \left(1 - \inf_{\theta \in \mathbb{R}^{d}} \mathbb{P}_{\theta} \Big(\left\| \widehat{\theta} - \theta \right\|_{2} \leqslant \varepsilon \Big) \right) \\ \ge \varepsilon^{2} \left(1 - \int_{\mathbb{R}^{d}} \mathbb{P}_{\theta} \Big(\left\| \widehat{\theta} - \theta \right\|_{2} \leqslant \varepsilon \Big) \, \mathrm{d}\nu(\theta) \right).$$
(44)

We take ν as the uniform distribution on the Euclidean ball $B(0, \rho \varepsilon) = \{u \in \mathbb{R}^d : ||u||_2 \leq \rho \varepsilon\}$ for some $\rho > 1$ to be determined by the analysis (as is also the case for ε). Fano's inequality in the form given by the combination of (7) and (13), with the fixed alternative distribution \mathbb{P}_0 (where 0 denotes the

null vector of \mathbb{R}^d) and the sets $A_{\theta} = \{ \|\widehat{\theta} - \theta\|_2 \leq \varepsilon \}$, indicates that

$$\int_{B(0,\rho\varepsilon)} \mathbb{P}_{\theta} \left(\left\| \widehat{\theta} - \theta \right\|_{2} \leqslant \varepsilon \right) d\nu(\theta) \\
\leqslant \int_{B(0,\rho\varepsilon)} \mathbb{P}_{0} \left(\left\| \widehat{\theta} - \theta \right\|_{2} \leqslant \varepsilon \right) d\nu(\theta) + \sqrt{\frac{\int_{B(0,\rho\varepsilon)} KL(\mathbb{P}_{\theta}, \mathbb{P}_{0}) d\nu(\theta)}{-\ln\left(\int_{B(0,\rho\varepsilon)} \mathbb{P}_{0} \left(\left\| \widehat{\theta} - \theta \right\|_{2} \leqslant \varepsilon \right) d\nu(\theta) \right)}} \\
\leqslant \left(\frac{1}{\rho} \right)^{d} + \sqrt{\frac{n\rho^{2}\varepsilon^{2}/(2\sigma^{2})}{d\ln\rho}},$$
(45)

where the second inequality follows from the inequalities (46) and (47) below. First note that, by independence, $\operatorname{KL}(\mathbb{P}_{\theta}, \mathbb{P}_0) = n \operatorname{KL}(\mathcal{N}(\theta, \sigma^2), \mathcal{N}(0, \sigma^2)) = n \|\theta\|_2^2 / (2\sigma^2)$, so that

$$\int_{B(0,\rho\varepsilon)} \operatorname{KL}(\mathbb{P}_{\theta}, \mathbb{P}_{0}) \,\mathrm{d}\nu(\theta) = \frac{n}{2\sigma^{2}} \int_{B(0,\rho\varepsilon)} \|\theta\|_{2}^{2} \,\mathrm{d}\nu(\theta) \leqslant \frac{n\rho^{2}\varepsilon^{2}}{2\sigma^{2}} \,. \tag{46}$$

Second, by the Fubini-Tonelli theorem,

$$q \stackrel{\text{def}}{=} \int_{B(0,\rho\varepsilon)} \mathbb{P}_0\Big(\|\widehat{\theta} - \theta\|_2 \leqslant \varepsilon \Big) \, \mathrm{d}\nu(\theta) = \mathbb{E}_0 \left[\int_{B(0,\rho\varepsilon)} \mathbb{1}_{\left\{ \|\widehat{\theta} - \theta\|_2 \leqslant \varepsilon \right\}} \, \mathrm{d}\nu(\theta) \right] \\ = \mathbb{E}_0 \Big[\nu\Big(B\big(\widehat{\theta},\varepsilon\big) \cap B\big(0,\rho\varepsilon\big) \Big) \Big] \leqslant \left(\frac{1}{\rho}\right)^d, \tag{47}$$

where to get the last inequality we used the fact that, almost surely, $\nu(B(\hat{\theta}, \varepsilon) \cap B(0, \rho\varepsilon))$ is the ratio of the volume of a (possibly truncated) Euclidean ball of radius ε with the volume of the support of ν , namely, the larger Euclidean ball $B(0, \rho\varepsilon)$, in dimension d.

We conclude the proof by combining (44) and (45) and recalling that $\rho > 1$ and $\varepsilon > 0$ were two parameters to get

$$\sup_{\theta \in \mathbb{R}^d} \mathbb{E}_{\theta} \Big[\left\| \widehat{\theta} - \theta \right\|_2^2 \Big] \ge \sup_{\rho > 1} \sup_{\varepsilon > 0} \varepsilon^2 \left(1 - \left(\frac{1}{\rho} \right)^d - \sqrt{\frac{n\rho^2 \varepsilon^2}{2\sigma^2 d \ln \rho}} \right) = \frac{d\sigma^2}{n} \underbrace{\sup_{\rho > 1} \frac{8\left(1 - (1/\rho)^d \right)^3 \ln(\rho)}{27\rho^2}}_{=c_d}$$

for the optimal choice of $\varepsilon = (2/3)(1 - (1/\rho)^d)\sqrt{2d\sigma^2 \ln(\rho)/(n\rho^2)}$. We see $c_1 \ge 0.01$ and $c_2 \ge 0.025$ via the respective choices $\rho = 4$ and $\rho = 2.5$, while the fact that the limit is larger than 0.05 follows, e.g., from the choice $\rho = 2$.

Note that, when using (13) above, we implicitly assumed that the quantity q defined in (47) lies in (0,1). The fact that q < 1 follows directly from the upper bound $(1/\rho)^d$ given $\rho > 1$. As for the condition q > 0, note that given the support of ν , we could rewrite (44) as

$$\begin{split} \sup_{\theta \in \mathbb{R}^d} \mathbb{E}_{\theta} \Big[\left\| \widehat{\theta} - \theta \right\|_2^2 \Big] \geqslant \sup_{\theta \in B(0,\rho\varepsilon)} \mathbb{E}_{\theta} \Big[\left\| \widehat{\theta} - \theta \right\|_2^2 \Big] \geqslant \sup_{\theta \in B(0,\rho\varepsilon)} \mathbb{E}_{\theta} \Big[\left\| \Pi(\widehat{\theta}) - \theta \right\|_2^2 \Big] \\ \geqslant \varepsilon^2 \left(1 - \int_{\mathbb{R}^d} \mathbb{P}_{\theta} \Big(\left\| \Pi(\widehat{\theta}) - \theta \right\|_2 \leqslant \varepsilon \Big) \, \mathrm{d}\nu(\theta) \right), \end{split}$$

where $\Pi(\widehat{\theta})$ is the projection of $\widehat{\theta}$ onto the closed convex set $B(0, \rho \varepsilon)$. Thus, we can assume without loss of generality that $\widehat{\theta} \in B(0, \rho \varepsilon)$ almost surely. In this case, the normalized volume $\nu(B(\widehat{\theta}, \varepsilon) \cap B(0, \rho \varepsilon))$ is almost surely positive, so that $q = \mathbb{E}_0[\nu(B(\widehat{\theta}, \varepsilon) \cap B(0, \rho \varepsilon))] > 0$. Comparison with other, historical proofs. Various types of proofs were proposed in the literature to derive a lower bound of order $d\sigma^2/n$ as above.

The proof technique that consists in lower bounding the minimax risk by the Bayes risk works surprisingly well in this simple estimation problem. It is indeed folklore knowledge that taking a Gaussian prior with covariance matrix $s^2 I_d$ and letting $s \to +\infty$ yields, after simple calculations, a lower bound of $d\sigma^2/n$ (see, e.g., Massart, 2007, page 106). Interestingly the multiplicative constant of 1 is even optimal because it matches the upper bound of $d\sigma^2/n$ satisfied by the empirical mean. However, this proof technique does not carry over easily to more complex settings such as, e.g., the same Gaussian model but with a bounded parameter space Θ , as is the case in the nonparametric regression problem of Section A.2 below. This is the reason why, even for this toy estimation problem, it is useful to provide alternative proofs that may be suboptimal in terms of the multiplicative constant, but that can be easily adapted to more intricate settings.

Another simple proof technique consists in using Assouad's lemma, which is very useful when the loss function can be decomposed as a sum over the d coordinates, as is the case here. Assouad's lemma reduces the estimation problem to d parallel two-hypotheses testing problems. (See, e.g., Yu, 1997, Example 2 for an application of Assouad's lemma.)

All alternative proofs that we know of are based on Fano's inequality and often involve a discretization argument. Historical proofs reduce the estimation problem to a multiple-hypotheses testing problem (with exponentially many hypotheses) by showing that every estimator must fail to identify at least one θ in a subset of the *d*-dimensional rescaled hypercube $\{0, c\sigma/\sqrt{n}\}^d$ (this subset is obtained via a combinatorial tool known as Varshamov-Gilbert's lemma). See, e.g., Yu [1997, Example 2] or Massart [2007, Proposition 4.8] for such applications of Fano's inequality. More recently Duchi and Wainwright [2013] and Chen et al. [2016] provided a continuous version of Fano's inequality (of which Lemma 6 above is a generalization to some extent) to avoid the discretization step mentioned earlier. Instead they directly addressed a multiple testing problem with continuously many hypotheses. This provides a nice interpretation of the factor d in the lower bound $d\sigma^2/n$ as the log ratio of the volumes of two Euclidean balls in \mathbb{R}^d , as in (47) above. Note however that both articles use a discretization argument at some point: Duchi and Wainwright [2013] prove their continuous Fano inequality (Proposition 2 therein) via an unnecessarily involved grid-based approximation argument. Chen et al. [2016] later proved a continuous Fano's inequality (cf. Corollary 3.5 and Theorem 4.1 therein) without any discretization argument, but the way they use it in Example 5.3 for the Gaussian model relies on an unnecessary calculation of covering numbers. On the contrary, the proof we provided above uses no discretization whatsoever.

We finally mention the lower bound that Xu and Raginsky [2016] derived for the Bayes risk with a uniform prior on a Euclidean ball (as in the proof above). The proof of their Corollary 3, which uses a generalized Fano's inequality, also bypasses any discretization step. It is however only asymptotic in n, and it requires longer calculations than above since log ratios of densities have to be manipulated explicitly.

A.2. A minimax lower bound for nonparametric regression

In this subsection we revisit a well-known lower bound within the nonparametric regression model with fixed design, which unfolds as follows. We observe an *n*-sample $(x_1, Y_1), \ldots, (x_n, Y_n)$, where $x_i = (i-1)/n \in [0,1]$ and

$$Y_i = f(x_i) + \varepsilon_i$$
, where $1 \leq i \leq n$,

for some unknown function $f \in \mathcal{F}_L \stackrel{\text{def}}{=} \{g : [0,1] \to \mathbb{R}, g \text{ is } L\text{-Lipschitz}\}$ and ε_i that are independent and identically distributed according to $\mathcal{N}(0, \sigma^2)$. The goal is to estimate f; the parameters are σ^2 , L and f. For the sake of notation, we only focus on f and denote by \mathbb{P}_f and \mathbb{E}_f the probability and expectation underlying the random vector (Y_1, \ldots, Y_n) . Actually, as in the previous section and with no loss of generality, we may identify the law of (Y_1, \ldots, Y_n) and the underlying probability. We assess the accuracy of any estimator $\widehat{f} \in \mathbb{L}^2([0,1])$ via its expected quadratic risk,

$$\mathbb{E}_f \Big[\|\widehat{f} - f\|_2^2 \Big],$$

where the squared Euclidean norm of any $g \in \mathbb{L}^2 \big([0, 1] \big)$ is defined as $\|g\|_2^2 = \int_0^1 g(x)^2 \, \mathrm{d}x.$

The next lower bound is well known; see Ibragimov and Has'minskii [1982, 1984], Tsybakov [2009, Theorem 2.8] or Duchi [2014, Theorem 4.4] for proofs based on either Fano's inequality (with a discretization argument) or Assouad's lemma. We illustrate below how to use the continuous Fano's inequality.

Proposition 19 (Nonparametric lower bound with no discretization). Fix $\sigma^2 > 0$ and L > 0, two quantities possibly known to the statistician. In the nonparametric regression model described above, we have, for all $n \ge \max\{L/(\sigma\sqrt{\ln(2)}), 64\ln(2)\sigma^2/L^2\}$,

$$\inf_{\widehat{f}} \sup_{f \in \mathcal{F}_L} \mathbb{E}_f \Big[\|\widehat{f} - f\|_2^2 \Big] \ge C \left(\frac{\sigma^2 L}{n} \right)^{2/3},$$

where the infimum is over all estimators $\hat{f} = \hat{f}(Y_1, \ldots, Y_n) \in \mathbb{L}^2([0,1])$, and where C is a universal positive constant; e.g., C = 0.001 works.

The lower bound of Proposition 19 is tight in the sense that there exist estimators \hat{f} such that

$$\sup_{f \in \mathcal{F}_L} \mathbb{E}_f \Big[\|\widehat{f} - f\|_2^2 \Big] \leqslant \widetilde{C} \left(\frac{\sigma^2 L}{n} \right)^{2/3}$$

for some universal constant \widetilde{C} ; see, e.g., Tsybakov [2009, Theorem 1.7] or Duchi [2014, Corollary 4.3].

Proof: We start as in the proof of Duchi [2014, Theorem 4.4]. Let φ be the function defined on \mathbb{R} by $\varphi(x) = (1/2 - |x - 1/2|)_+$, where $y_+ = \max\{0, y\}$. Note that φ is a 1–Lipschitz function with support (0, 1); in addition,

$$\|\varphi\|_2^2 = 2 \int_0^{1/2} x^2 \, \mathrm{d}x = \frac{1}{12}.$$

Thus, for any integer $d \ge 2$ (to be determined by the analysis), the functions $f_j : [0,1] \to \mathbb{R}$ defined for all $j \in \{1, \ldots, d\}$ by

$$f_j(x) = \frac{L}{d} \varphi\left(d\left(x - \frac{j-1}{d}\right)\right)$$

belong to \mathcal{F}_L and form an orthogonal system in $\mathbb{L}^2([0,1])$, since the f_j have pairwise disjoint supports. We define $f_{\theta} = \sum_{j=1}^d \theta_j f_j$ for all $\theta \in \mathbb{R}^d$. Note that, again because of the disjoint supports of the f_j , the mapping $\theta \mapsto f_{\theta}$ is an injection from Θ into \mathcal{F}_L , where

$$\Theta \stackrel{\mathrm{def}}{=} \left\{ \theta \in \mathbb{R}^d : \|\theta\|_{\infty} \leqslant 1 \right\}.$$

Moreover, for all $\theta, \theta' \in \Theta$, we have the norm relationship

$$\|f_{\theta} - f_{\theta'}\|_{2}^{2} = \sum_{j=1}^{d} (\theta_{j} - \theta_{j}')^{2} \|f_{j}\|_{2}^{2} = \frac{L^{2}}{d^{3}} \|\varphi\|_{2}^{2} \|\theta - \theta'\|_{2}^{2} = \frac{L^{2}}{12 d^{3}} \|\theta - \theta'\|_{2}^{2}.$$
 (48)

Next we reduce the nonparametric problem to a parametric one and then proceed as in the proof of Proposition 18, avoiding any discretization argument. To that end, we write abusively $\mathbb{P}_{\theta} = \mathbb{P}_{f_{\theta}}$ and $\mathbb{E}_{\theta} = \mathbb{E}_{f_{\theta}}$.

We set $\widehat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmin}} \|\widehat{f} - f_{\theta}\|_{2}^{2}$ and get

$$\sup_{f \in \mathcal{F}_L} \mathbb{E}_f \Big[\|\widehat{f} - f\|_2^2 \Big] \ge \sup_{\theta \in \Theta} \mathbb{E}_\theta \Big[\|\widehat{f} - f_\theta\|_2^2 \Big] \ge \sup_{\theta \in \Theta} \mathbb{E}_\theta \Big[\|f_{\widehat{\theta}} - f_\theta\|_2^2 \Big],$$
(49)

where we first used that the set $\mathcal{F}_{\Theta} = \{f_{\theta} : \theta \in \Theta\}$ is a subset of \mathcal{F}_L , and second, that $f_{\widehat{\theta}}$ is the projection of \widehat{f} onto the closed convex subset \mathcal{F}_{Θ} of $\mathbb{L}^2([0,1])$. Now, for all $\rho > 0$ (to be determined by the analysis), inequality (48) and Markov's inequality yield

$$\sup_{\theta \in \Theta} \mathbb{E}_{\theta} \left[\left\| f_{\widehat{\theta}} - f_{\theta} \right\|_{2}^{2} \right]$$

$$= \frac{L^{2}}{12 d^{3}} \sup_{\theta \in \Theta} \mathbb{E}_{\theta} \left[\left\| \widehat{\theta} - \theta \right\|_{2}^{2} \right] \geqslant \frac{L^{2}}{12 d^{3}} \sup_{\theta \in \Theta} \mathbb{E}_{\theta} \left[\rho d \mathbb{1}_{\left\{ \| \widehat{\theta} - \theta \|_{2}^{2} \geqslant \rho d \right\}} \right]$$

$$\geqslant \frac{\rho L^{2}}{12 d^{2}} \left(1 - \inf_{\theta \in \Theta} \mathbb{P}_{\theta} \left(\left\| \widehat{\theta} - \theta \right\|_{2}^{2} \leqslant \rho d \right) \right) \geqslant \frac{\rho L^{2}}{12 d^{2}} \left(1 - \int_{\Theta} \mathbb{P}_{\theta} \left(\left\| \widehat{\theta} - \theta \right\|_{2}^{2} \leqslant \rho d \right) d\nu(\theta) \right), \quad (50)$$

where the last inequality holds true for any prior ν on Θ .

Now we choose the uniform (Lebesgue) prior $\nu \stackrel{\text{def}}{=} \mathcal{U}(\Theta)$ and apply Fano's inequality in the form of Lemma 6, with the fixed alternative distribution \mathbb{P}_0 (where 0 denotes the null vector of \mathbb{R}^d):

$$\int_{\Theta} \mathbb{P}_{\theta} \left(\left\| \widehat{\theta} - \theta \right\|_{2}^{2} \leqslant \rho d \right) d\nu(\theta) \leqslant \frac{\ln(2) + \int_{\Theta} \mathrm{KL} \left(\mathbb{P}_{\theta}, \mathbb{P}_{0} \right) d\nu(\theta)}{-\ln\left(\int_{\Theta} \mathbb{P}_{0} \left(\left\| \widehat{\theta} - \theta \right\|_{2}^{2} \leqslant \rho d \right) d\nu(\theta) \right)} \leqslant \frac{1}{3} + \frac{(n+d)L^{2}}{8 d^{3} \sigma^{2} \ln(2)}, \quad (51)$$

where the second inequality follows from the inequalities (52) and (54) below, with the choice of $\rho = 1/(2\pi e)$. Note that these calculations also show that the integral in the denominator in (51) lies in (0,1), as required for Lemma 6.

First note that $\mathbb{P}_{\theta} = \bigotimes_{i=1}^{n} \mathcal{N}(f_{\theta}(x_i), \sigma^2)$ and $\mathbb{P}_0 = \mathcal{N}(0, \sigma^2)^{\otimes n}$, so that, by independence,

$$\mathrm{KL}(\mathbb{P}_{\theta}, \mathbb{P}_{0}) = \sum_{i=1}^{n} \mathrm{KL}\Big(\mathcal{N}\big(f_{\theta}(x_{i}), \sigma^{2}\big), \, \mathcal{N}(0, \sigma^{2})\Big) = \frac{1}{2\sigma^{2}} \sum_{i=1}^{n} f_{\theta}(x_{i})^{2},$$

which, since the f_j have pairwise disjoint supports, is also equal to

$$\sum_{i=1}^{n} f_{\theta}(x_i)^2 = \sum_{j=1}^{d} \theta_j^2 \sum_{i=1}^{n} f_j(x_i)^2 \leq \|\theta\|_2^2 \left(\frac{n}{d} + 1\right) \frac{L^2}{4 d^2} \leq \left(\frac{n}{d} + 1\right) \frac{L^2}{4 d},$$

where for the first inequality we used that $||f_j||_{\infty} \leq L/(2d)$ and that at most n/d + 1 design points x_i are in the support of a function f_j , while the second inequality follows from $||\theta||_2^2 \leq d||\theta||_{\infty}^2 \leq d$ as $\theta \in \Theta$. Summarizing, we proved

$$\forall \theta \in \Theta, \qquad \operatorname{KL}(\mathbb{P}_{\theta}, \mathbb{P}_0) \leqslant \frac{(n+d)L^2}{8 d^2 \sigma^2}.$$
(52)

Second, as in inequality (47), we write

$$\int_{\Theta} \mathbb{P}_{0} \left(\left\| \widehat{\theta} - \theta \right\|_{2}^{2} \leq \rho d \right) d\nu(\theta) = \mathbb{E}_{0} \left[\nu \left(B \left(\widehat{\theta}, \sqrt{\rho d} \right) \cap \Theta \right) \right]$$
$$\leq \frac{\left(\sqrt{\rho d} \right)^{d}}{2^{d}} \operatorname{Vol}_{\mathbb{R}^{d}} \left(B(0, 1) \right) = \left(\frac{\rho d}{4} \right)^{d/2} \frac{\pi^{d/2}}{\Gamma(1 + d/2)}, \tag{53}$$

where we used the formula for the volume of the unit Euclidean ball B(0,1) of \mathbb{R}^d . As in addition,

$$\Gamma\left(1+\frac{d}{2}\right) = \int_0^{+\infty} t^{d/2} \,\mathrm{e}^{-t} \,\mathrm{d}t \ge \int_{d/2}^{+\infty} (d/2)^{d/2} \,\mathrm{e}^{-t} \,\mathrm{d}t = \left(\frac{d/2}{\mathrm{e}}\right)^{d/2} \,,$$

we finally get with the choice $\rho=1/(2\pi e)$

$$\int_{\Theta} \mathbb{P}_0\left(\left\|\widehat{\theta} - \theta\right\|_2^2 \leqslant \rho d\right) d\nu(\theta) \leqslant \left(\frac{\rho d}{4}\right)^{d/2} \frac{\pi^{d/2}}{\Gamma(1 + d/2)} \leqslant \left(\frac{\rho \pi e}{2}\right)^{d/2} = 4^{-d/2} = 2^{-d} \leqslant \frac{1}{8}$$
(54)

for $d \ge 3$. For d = 2, using that $\Gamma(2) = 1$, we see that the final upper bound in (53) equals $1/(4e) \le 1/8 \le 2^{-d}$. We use the 2^{-d} upper bound for the second term in the left-hand side of (51) and the 1/8 upper bound for its first term.

Now, combining (49)–(51), we proved so far that, for any integer $d \ge 2$ to be determined,

$$\sup_{f \in \mathcal{F}_L} \mathbb{E}_f \Big[\|\widehat{f} - f\|_2^2 \Big] \ge \frac{L^2}{24 \,\pi \mathrm{e} \, d^2} \left(\frac{2}{3} - \frac{(n+d)L^2}{8 \, d^3 \sigma^2 \ln(2)} \right) \,. \tag{55}$$

In the rest of the proof we choose

$$d = \left\lceil \left(\frac{n L^2 / \sigma^2}{\ln(2)}\right)^{1/3} \right\rceil \,.$$

We consider sample sizes n such that

$$n \ge \max\left\{\frac{L}{\sigma\sqrt{\ln(2)}}, 64\ln(2)\frac{\sigma^2}{L^2}\right\};$$

the first condition ensures that $d \leq n$ while the second condition entails that $d \geq 4$. The bound $d \geq 2$ was required above, while the bound $d \leq n$ simplifies the lower bound (55) into

$$\frac{L^2}{24 \,\pi \mathrm{e} \, d^2} \left(\frac{2}{3} - \frac{(n+d)L^2}{8 \, d^3 \sigma^2 \ln(2)} \right) \geqslant \frac{L^2}{24 \,\pi \mathrm{e} \, d^2} \left(\frac{2}{3} - \frac{n \, L^2}{4 \, d^3 \sigma^2 \ln(2)} \right).$$

We substitute the value of d therein, using that

$$\left(\frac{n\,L^2/\sigma^2}{\ln(2)}\right)^{1/3} \leqslant d \leqslant 1 + \left(\frac{n\,L^2/\sigma^2}{\ln(2)}\right)^{1/3} \leqslant \frac{5}{4} \left(\frac{n\,L^2/\sigma^2}{\ln(2)}\right)^{1/3},$$

where the last inequality follows from $n \ge 64 \ln(2) \sigma^2 / L^2$. We get

$$\begin{aligned} \frac{L^2}{24 \,\pi \mathrm{e} \, d^2} \left(\frac{2}{3} - \frac{n \, L^2}{4 \, d^3 \sigma^2 \ln(2)}\right) &\geqslant \frac{L^2}{24 \,\pi \mathrm{e} \, d^2} \left(\frac{2}{3} - \frac{1}{4}\right) = \frac{L^2}{24 \,\pi \mathrm{e} \, d^2} \frac{5}{12} \\ &\geqslant \frac{L^2}{24 \,\pi \mathrm{e}} \left(\frac{1}{(5/4)^2} \left(\frac{n \, L^2/\sigma^2}{\ln(2)}\right)^{-2/3}\right) \frac{5}{12} \\ &= C \left(\frac{\sigma^2 L}{n}\right)^{2/3} \quad \text{where} \quad C = \frac{\left(\ln(2)\right)^{2/3}}{90 \,\pi \mathrm{e}}.\end{aligned}$$

A numerical computation shows that this value of C is larger than 0.001, as claimed. Collecting all bounds, the proof is concluded.

Supremum versus Euclidean norms. Though the general structure of the proof is similar to that of Proposition 18, we emphasize a technical difference: here, the support Θ of the prior ν is a ball in the supremum norm instead of the Euclidean norm. The reason is that, contrary to Proposition 18 where the choice of Θ was not constrained (so that we could choose Θ as a Euclidean ball of arbitrary radius), here, we should choose $\Theta \subset [-1,1]^d$ to ensure the inclusion $\mathcal{F}_{\Theta} \subset \mathcal{F}_L$. Taking a Euclidean ball of radius at most 1 (and ν a uniform prior on this ball) would have led to a choice ρ of the order of 1/d for calculations similar to (53) and (54) to upper bound the integral at hand by a numerical constant smaller than 1. This would result in a lower bound not of the right orders of magnitude in n, L and σ^2 . On the contrary, the ball Θ in the sup norm allowed us to choose ρ as a constant and hence get an optimal lower bound.

Extension to Hölder functions. We can easily generalize the proof above to the set of (β, L) -Hölder functions over [0,1] to get a lower bound of the order of $n^{-2\beta/(2\beta+1)}$, where $\beta > 0$. We first recall the definition of such functions. Let $(p, \alpha) \in \mathbb{N} \times (0, 1]$ be such that $\beta = p + \alpha$. A function $f:[0,1] \to \mathbb{R}$ is called (β, L) -Hölder if it is p times differentiable and if for all $(x, y) \in [0,1]^2$,

$$|f^{(p)}(x) - f^{(p)}(y)| \leq L|x - y|^{\alpha}$$

We can indeed use the same construction (but without any discretization) as in Tsybakov [2009, Section 2.6] by choosing a function $\varphi : \mathbb{R} \to \mathbb{R}$ that is infinitely differentiable, $(\beta, 1)$ -Hölder, and supported on (0, 1), and by considering an orthogonal system of the form

$$f_j(x) = \frac{L}{d^{\beta}} \varphi\left(d\left(x - \frac{j-1}{d}\right)\right), \qquad j \in \{1, \dots, d\},$$

for a well-chosen d. The proof then follows exactly the same lines as above, with the same functions of the form $f_{\theta} = \sum_{j} \theta_{j} f_{j}$, for a parameter θ in $\Theta = \{\theta \in \mathbb{R}^{d} : \|\theta\|_{\infty} \leq 1\}$.

Note that this proof technique also works in higher dimensions, i.e., for (β, L) -Hölder functions over the *m*-dimensional cube $[0, 1]^m$. A simple adaptation of the above arguments indeed yields a lower bound of the order of $n^{-2\beta/(2\beta+m)}$, in the same spirit as in the lower bound that Györfi et al. [2002, Theorem 3.2] derived in the regression model with random design.

B. From Bayesian posteriors to point estimators

We recall below a well-known result that indicates how to construct good point estimators from good Bayesian posteriors (Section B.1 below). One theoretical benefit is that this result can be used to convert known minimax lower bounds for point estimation into minimax lower bounds for posterior concentration rates (Section B.2 below). This technique is thus a—less direct—alternative to the method we presented in Section 4.1.

B.1. The conversion

The following statement is a nonasymptotic variant of Theorem 2.5 by Ghosal et al. [2000] (see also Chapter 12, Proposition 3 by Le Cam, 1986, as well as Section 5.1 by Hoffmann et al., 2015). We consider the same setting as in Section 4.1 and assume in particular that the underlying probability measure is given by $P_{\theta}^{\otimes n}$.

Proposition 20 (From Bayesian posteriors to point estimators).

Let $n \ge 1$, $\delta > 0$, and $\theta \in \Theta$. Let $\widehat{\theta}_n = \widehat{\theta}_n(X_1, \ldots, X_n)$ be any estimator satisfying, $P_{\theta}^{\otimes n}$ -almost surely,

$$\mathbb{P}_{\pi}\left(\theta': \ell(\theta', \widehat{\theta}_{n}) < \varepsilon_{n} \, \big| \, X_{1:n}\right) \geqslant \sup_{\widetilde{\theta} \in \Theta} \mathbb{P}_{\pi}\left(\theta': \ell\left(\theta', \widetilde{\theta}\right) < \varepsilon_{n} \, \big| \, X_{1:n}\right) - \delta \,.$$
(56)

Then,

$$P_{\theta}^{\otimes n} \left(\mathbb{P}_{\pi} \left(\theta' : \ell(\theta', \theta) \geqslant \varepsilon_n \, \big| \, X_{1:n} \right) \geqslant \frac{1 - \delta}{2} \right) \geqslant P_{\theta}^{\otimes n} \left(\ell(\widehat{\theta}_n, \theta) \geqslant 2\varepsilon_n \right) \,. \tag{57}$$

This result implies that if $\hat{\theta}_n$ is a center of a ball that almost maximizes the posterior mass—see assumption (56)—and if the posterior mass concentrates around θ at a rate $\varepsilon'_n < \varepsilon_n$ —so that the lefthand side of (57) vanishes by Markov's inequality—then $\hat{\theta}_n$ is $(2\varepsilon_n)$ -close to θ with high probability. Therefore, at least from a theoretical viewpoint, a good posterior distribution can be converted into a good point estimator, by defining $\hat{\theta}_n$ based on $\mathbb{P}_{\pi}(\cdot | X_{1:n})$ such that (56) holds, i.e., by taking an approximate argument of the supremum. A measurable such $\hat{\theta}_n$ exists as soon as Θ is a separable topological space and the function $\tilde{\theta} \mapsto \mathbb{P}_{\pi}(\theta' : \ell(\theta', \tilde{\theta}) < \varepsilon_n | x_{1:n})$ is lower-semicontinuous for $\mathfrak{m}^{\otimes n}$ almost every $x_{1:n} \in \mathcal{X}^n$ (see the end of the proof of Corollary 21 for more details).

Proof: Denote by $B_{\ell}(\theta, \varepsilon) \stackrel{\text{def}}{=} \{\theta' \in \Theta : \ell(\theta', \theta) < \varepsilon\}$ the open ℓ -ball of center θ and radius ε . By the triangle inequality we have the following inclusions of events:

$$\left\{ \ell(\widehat{\theta}_{n}, \theta) \geq 2\varepsilon_{n} \right\} \subseteq \left\{ B_{\ell}(\widehat{\theta}_{n}, \varepsilon_{n}) \cap B_{\ell}(\theta, \varepsilon_{n}) = \varnothing \right\}$$

$$\subseteq \left\{ \mathbb{P}_{\pi} \left(B_{\ell}(\widehat{\theta}_{n}, \varepsilon_{n}) \mid X_{1:n} \right) + \mathbb{P}_{\pi} \left(B_{\ell}(\theta, \varepsilon_{n}) \mid X_{1:n} \right) \leqslant 1 \right\}$$

$$\subseteq \left\{ \mathbb{P}_{\pi} \left(B_{\ell}(\theta, \varepsilon_{n}) \mid X_{1:n} \right) \leqslant \frac{1+\delta}{2} \right\}$$

$$(58)$$

$$= \left\{ 1 - \mathbb{P}_{\pi} \left(\theta' : \ell(\theta', \theta) < \varepsilon_n \, \big| \, X_{1:n} \right) \geqslant \frac{1 - \delta}{2} \right\}$$

$$= \left\{ \mathbb{P}_{\pi} \left(\theta' : \ell(\theta', \theta) \geqslant \varepsilon_n \, \big| \, X_{1:n} \right) \geqslant \frac{1 - \delta}{2} \right\},$$
(59)

where (58) follows from the lower bound $\mathbb{P}_{\pi}(B_{\ell}(\widehat{\theta}_n,\varepsilon_n) | X_{1:n}) \geq \mathbb{P}_{\pi}(B_{\ell}(\theta,\varepsilon_n) | X_{1:n}) - \delta$, which holds by assumption (56) on $\widehat{\theta}_n$. This concludes the proof.

B.2. Application to posterior concentration lower bounds

We explained above that a good posterior distribution can be converted into a good point estimator. As noted by Ghosal et al. [2000] this conversion can be used the other way around: if we have a lower bound on the minimax rate of estimation, then Proposition 20 provides a lower bound on the minimax posterior concentration rate, as formalized in the following corollary. Assumption (60) below corresponds to an in-probability minimax lower bound; it is for instance a consequence of (45) in the standard multivariate Gaussian model with Euclidean loss.

Corollary 21. Let $n \ge 1$. Consider the setting of Section 4.1, with underlying probability measure $P_{\theta}^{\otimes n}$ when the unknown parameter is θ . Assume that Θ is a separable topological space and that $\tilde{\theta} \mapsto \ell(\theta', \tilde{\theta})$ is continuous for all $\theta' \in \Theta$. Assume also that for some absolute constant c < 1, we have

for all estimators
$$\widehat{\theta}_n$$
, $\inf_{\theta \in \Theta} P_{\theta}^{\otimes n} \left(\ell(\widehat{\theta}_n, \theta) < 2\varepsilon_n \right) \leqslant c$. (60)

Then, for all priors π' on Θ ,

$$\inf_{\theta \in \Theta} \mathbb{E}_{\theta} \left[\mathbb{P}_{\pi'} \left(\theta' : \ell(\theta', \theta) < \varepsilon_n \, \big| \, X_{1:n} \right) \right] \leqslant \frac{1+c}{2} < 1 \,.$$
(61)

Proof: Let $\delta > 0$ be a parameter that we will later take arbitrarily small. Fix any estimator $\hat{\theta}_n$ satisfying (56) for the prior π' , i.e., that almost maximizes the posterior mass on an open ball

of radius ε_n . (See the end of the proof for details on why such a measurable $\hat{\theta}_n$ exists.) Then, Proposition 20 used for all $\theta \in \Theta$ entails that

$$\sup_{\theta\in\Theta} P_{\theta}^{\otimes n} \left(\mathbb{P}_{\pi'} \left(\theta' : \ell(\theta', \theta) \geqslant \varepsilon_n \, \big| \, X_{1:n} \right) \geqslant \frac{1-\delta}{2} \right) \geqslant \sup_{\theta\in\Theta} P_{\theta}^{\otimes n} \left(\ell(\widehat{\theta}_n, \theta) \geqslant 2\varepsilon_n \right) \geqslant 1-c \,,$$

where the last inequality follows from the assumption (60). Now we use Markov's inequality to upper bound the left-hand side above and obtain

$$\frac{2}{1-\delta} \sup_{\theta \in \Theta} \mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi'} \big(\theta' : \ell(\theta', \theta) \ge \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \ge \sup_{\theta \in \Theta} P_{\theta}^{\otimes n} \Big(\mathbb{P}_{\pi'} \big(\theta' : \ell(\theta', \theta) \ge \varepsilon_n \, \big| \, X_{1:n} \big) \ge \frac{1-\delta}{2} \Big) \ge 1-c \,.$$

Letting $\delta \to 0$ and dividing both sides by 2 yields

$$1 - \inf_{\theta \in \Theta} \mathbb{E}_{\theta} \Big[\mathbb{P}_{\pi'} \big(\theta' : \ell(\theta', \theta) < \varepsilon_n \, \big| \, X_{1:n} \big) \Big] \geqslant \frac{1 - c}{2}$$

Rearranging terms concludes the proof of (61). We now address the technical issue mentioned at the beginning of the proof.

Why a measurable $\hat{\theta}_n$ exists. Note that it is possible to choose $\hat{\theta}_n$ satisfying (56) with π' in a mesurable way as soon as Θ is a separable topological space and

$$\psi: \widetilde{\theta} \in \Theta \longmapsto \mathbb{P}_{\pi'}(\theta': \ell(\theta', \widetilde{\theta}) < \varepsilon_n \mid x_{1:n})$$

is lower-semicontinuous for $\mathfrak{m}^{\otimes n}$ -almost every $x_{1:n} \in \mathcal{X}^n$, and thus $\mathbb{P}_{\theta}^{\otimes n}$ -almost surely for all $\theta \in \Theta$. The reason is that, in that case, it is possible to equate the supremum of ψ over Θ to a supremum on a countable subset of Θ . Next, and thanks to the continuity assumption on ℓ , we prove that the desired lower-semicontinuity holds true for all $x_{1:n} \in \mathcal{X}^n$ (not just almost all of them).

To that end, we show the lower-semicontinuity at any fixed $\theta^* \in \Theta$. Consider any sequence $(\tilde{\theta}_i)_{i \ge 1}$ in Θ converging to θ^* . For all $x_{1:n} \in \mathcal{X}^n$, by Fatou's lemma applied to the well-defined probability distribution $\mathbb{P}_{\pi'}(\cdot | x_{1:n})$, we have,

$$\lim_{i \to +\infty} \mathbb{P}_{\pi'} \left(\theta' : \ell(\theta', \widetilde{\theta}_i) < \varepsilon_n \mid x_{1:n} \right) = \liminf_{i \to +\infty} \mathbb{E}_{\pi'} \left[\mathbb{1}_{\{\ell(\theta', \widetilde{\theta}_i) < \varepsilon_n\}} \mid x_{1:n} \right] \\
\geqslant \mathbb{E}_{\pi'} \left[\liminf_{\substack{i \to +\infty \\ = 1 \text{ if } \ell(\theta', \theta^*) < \varepsilon_n}} \mathbb{1}_{\{\ell(\theta', \theta^*) < \varepsilon_n\}} \mid x_{1:n} \right] \\
\geqslant \mathbb{P}_{\pi'} \left(\theta' : \ell(\theta', \theta^*) < \varepsilon_n \mid x_{1:n} \right),$$
(62)

where in (62) we identify that the lim inf equals 1 as soon as $\ell(\theta', \theta^*) < \varepsilon_n$ by continuity of $\tilde{\theta} \mapsto \ell(\theta', \tilde{\theta})$ at $\tilde{\theta} = \theta^*$.

C. Variations on Theorem 14

The original result by Birgé [2005, Corollary 1] reads, with the notation of Theorem 14:

$$\min_{1 \le i \le N} \mathbb{P}_i(A_i) \le \max\left\{ d_N, \ \frac{\overline{K}}{\ln(N)} \right\},\tag{63}$$

where $(d_N)_{N \ge 2}$ is a decreasing sequence, defined as follows, based on functions $r_N : [0,1) \to \mathbb{R}$:

$$r_N(a) = \operatorname{kl}\left(a, \frac{1-a}{N-1}\right) - a \ln(N)$$
 and $d_N = \max\left\{a \in [0,1]: r_N(a) \le 0\right\}$.

On the other hand, the simplification by Massart [2007, Section 2.3.4] leads to

$$\min_{1 \leqslant i \leqslant N} \mathbb{P}_i(A_i) \leqslant \max\left\{\frac{2\mathrm{e} - 1}{2\mathrm{e}}, \ \overline{\frac{K}{\ln(N)}}\right\}.$$
(64)

(The original constant was a larger 2e/(2e+1) in Massart, 2007, Section 2.3.4.)

Before proving these results, we compare them with Theorem 14. The values of the c_N of Theorem 1, of the d_N of (63) and of (2e - 1)/(2e) are given by (values rounded upwards)

$$\frac{2e-1}{2e} \approx 0.8161 \qquad \text{and} \qquad \begin{array}{ccccccc} N & 2 & 3 & 7 & +\infty \\ c_N & 0.7587 & 0.7127 & < 0.67 & 0.63987 \\ d_N & 0.7428 & 0.7009 & < 2/3 & 0.63987 \end{array}$$

The c_N and d_N are thus extremely close. While the c_N are slightly larger than the d_N (with, however, the same limit), they are easier to compute in practice. (See the closed-form expression for r_N below.) Also, the proof of Theorem 14 is simpler than the proof of Birgé [2005, Corollary 1]: they rely on the proof scheme but the former involves fewer calculations than the latter. Indeed, let us now prove again Birgé [2005, Corollary 1].

Proof of (63). We use the notation of the proof of Theorem 14 and its beginning. We can assume with no loss of generality that $a \ge 1/N$, so that, using the definition of a,

$$\widetilde{q} \leqslant \frac{1-a}{N-1} \leqslant a \leqslant \widetilde{p}; \tag{65}$$

therefore,

$$\operatorname{kl}(\widetilde{p},\widetilde{q}) \ge \operatorname{kl}(a,\widetilde{q}) \ge \operatorname{kl}\left(a,\frac{1-a}{N-1}\right),$$

since by convexity, $p \mapsto \text{kl}(p,q)$ is increasing on [q,1] and $q \mapsto \text{kl}(p,q)$ is decreasing on [0,p]. Combining this with $\overline{K} \ge \text{kl}(\widetilde{p},\widetilde{q})$, one has proved

$$\overline{K} \ge \operatorname{kl}\left(a, \frac{1-a}{N-1}\right) = a \ln(N) + r_N(a),$$

from which the conclusion follows after studying the variations of $r_N(a)$ in a and N. This last analytical part of the proof is tedious, as

$$r_N(a) = a \ln\left(\frac{a}{1-a}\right) + (1-a) \ln\left(\frac{1-a}{1-\frac{1-a}{N-1}}\right) + \left(a \ln(N-1) - a \ln(N)\right)$$
$$= \left(a \ln(a) + (1-2a) \ln(a)\right) + a \ln\left(\frac{N-1}{N}\right) + (1-a) \ln\left(\frac{N-1}{N-2+a}\right),$$

and we could overcome these heavy calculations in our proof of Theorem 14.

Proof of (64). For $p \ge \ln(2)$ and all $q \in [0, 1]$,

$$\operatorname{kl}(p,q) \ge p \ln\left(\frac{1}{q}\right) - \ln(2) \ge p \ln\left(\frac{1}{q}\right) - p = p \ln\left(\frac{1}{\operatorname{eq}}\right).$$
(66)

Equation (39) is adapted as

$$a \leqslant \widetilde{p}$$
 and $\widetilde{q} \leqslant \frac{1-a}{N-1} \leqslant \frac{2(1-a)}{N} \leqslant \frac{1}{eN}$

where we used respectively, for the last two inequalities, that $1/(N-1) \leq 2/N$ for $N \geq 2$ and that, with no loss of generality, $a \geq (2e-1)/(2e)$. In particular, $e\tilde{q} \leq 1/N$. Combining this with $\overline{K} \geq kl(\tilde{p}, \tilde{q})$ and (66), we have proved

$$\overline{K} \ge \widetilde{p} \ln\left(\frac{1}{\mathrm{e}\widetilde{q}}\right) \ge a \ln(N) \,,$$

which concludes the proof.

D. Proofs of basic facts about *f*-divergences (and thus, about Kullback-Leibler divergences)

The results recalled and re-proved in this section were stated in the main body of the article (Sections 2 and 3) for Kullback-Leibler divergences, which are a special case of f-divergences with $f(x) = x \ln x$. We restate them in greater generality and to that end, first recall the definition of f-divergences. Note that these f-divergences will be further studied in Section E below, where we show that the reductions and results of Section 3 extend in a straightforward manner to arbitrary f-divergences.

D.1. Definition of f-divergences and basic properties

The definition of the Kullback-Leibler divergence can be generalized as follows (see Csiszár, 1963, Ali and Silvey, 1966 and Gushchin, 2003 for further details). Let $f: (0, +\infty) \to \mathbb{R}$ be any convex function satisfying f(1) = 0. By convexity, we can define

$$f(0) \stackrel{\text{def}}{=} \lim_{t \downarrow 0} f(t) \in \mathbb{R} \cup \{+\infty\};$$

the extended function $f: [0, +\infty) \to \mathbb{R} \cup \{+\infty\}$ is still convex.

Maximal slope. Note that, for any x > 0, the limit

$$\lim_{t \to +\infty} \frac{f(t) - f(x)}{t - x} = \sup_{t > 0} \frac{f(t) - f(x)}{t - x} \in [0, +\infty]$$

exists since (by convexity) the slope (f(t) - f(x))/(t - x) is non-decreasing as t increases. Besides, this limit does not depend on x and equals

$$M_f \stackrel{\text{def}}{=} \lim_{t \to +\infty} \frac{f(t)}{t} \in [0, +\infty],$$

which thus represents the maximal slope of f. An inequality that we will repeatedly use and that follows from the two equations above with t = x + y is

$$\forall \, x > 0, \, \, y > 0 \,, \qquad \frac{f(x+y) - f(x)}{y} \leqslant M_f \,,$$

Put differently,

$$\forall x \ge 0, \ y \ge 0, \qquad f(x+y) \le f(x) + y M_f, \tag{67}$$

where the extension to y = 0 is immediate and the one to x = 0 follows by continuity of f on $(0, +\infty)$, which itself follows from its convexity.

Lebesgue decomposition of measures. We recall that \ll denotes the absolute continuity between measures and we let \perp denote the fact that two measures are singular. For distributions \mathbb{P} and \mathbb{Q} defined on the same measurable space (Ω, \mathcal{F}) , the Lebesgue decomposition of \mathbb{P} with respect to \mathbb{Q} is denoted by

$$\mathbb{P} = \mathbb{P}_{ac} + \mathbb{P}_{sing}, \quad \text{where} \quad \mathbb{P}_{ac} \ll \mathbb{Q} \quad \text{and} \quad \mathbb{P}_{sing} \bot \mathbb{Q}, \tag{68}$$

so that \mathbb{P}_{ac} and \mathbb{P}_{sing} are both sub-probabilities (positive measures with total mass smaller than or equal to 1) and, by definition,

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}} = \frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}}$$

Definition. The f-divergence $\operatorname{Div}_f(\mathbb{P}, \mathbb{Q})$ between \mathbb{P} and \mathbb{Q} is defined as

$$\operatorname{Div}_{f}(\mathbb{P},\mathbb{Q}) = \int_{\Omega} f\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) \mathrm{d}\mathbb{Q} + \mathbb{P}_{\operatorname{sing}}(\Omega) M_{f}.$$
(69)

The existence of the integral in the right-hand side follows from the general form of Jensen's inequality stated in Lemma 28 (Appendix F) with $\varphi = f$ and $C = [0, +\infty)$. This inequality, together with (67), also indicates that $\text{Div}_f(\mathbb{P}, \mathbb{Q}) \ge 0$. Indeed,

$$\int_{\Omega} f\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) \mathrm{d}\mathbb{Q} \ge f\left(\int_{\Omega} \frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}} \,\mathrm{d}\mathbb{Q}\right) = f\left(\mathbb{P}_{\mathrm{ac}}(\Omega)\right),$$

so that by (67),

 $\operatorname{Div}_{f}(\mathbb{P},\mathbb{Q}) \geq f(\mathbb{P}_{\operatorname{ac}}(\Omega)) + \mathbb{P}_{\operatorname{sing}}(\Omega) M_{f} \geq f(\mathbb{P}_{\operatorname{ac}}(\Omega) + \mathbb{P}_{\operatorname{sing}}(\Omega)) = f(1) = 0.$

Concrete and important examples of f-divergences, such as the Hellinger distance and the χ^2 divergence, are discussed in details below. The Kullback-Leibler divergence corresponds to the function $f: x \mapsto x \ln(x)$. We have $M_f = +\infty$ for the Kullback-Leibler and χ^2 -divergences, while $M_f = 1$ for the Hellinger distance.

D.2. The data-processing inequality and two major consequences

Lemma 22 (Data-processing inequality). Let \mathbb{P} and \mathbb{Q} be two probability distributions over the same measurable space (Ω, \mathcal{F}) , and let X be any random variable on (Ω, \mathcal{F}) . Denote by \mathbb{P}^X and \mathbb{Q}^X the laws of X under \mathbb{P} and \mathbb{Q} respectively. Then,

$$\operatorname{Div}_f(\mathbb{P}^X, \mathbb{Q}^X) \leq \operatorname{Div}_f(\mathbb{P}, \mathbb{Q}).$$

Corollary 23 (Data-processing inequality with expectations of random variables). Let \mathbb{P} and \mathbb{Q} be two probability distributions over the same measurable space (Ω, \mathcal{F}) , and let X be any random variable on (Ω, \mathcal{F}) taking values in [0, 1]. Denote by $\mathbb{E}_{\mathbb{P}}[X]$ and $\mathbb{E}_{\mathbb{Q}}[X]$ the expectations of X under \mathbb{P} and \mathbb{Q} respectively. Then,

$$\operatorname{div}_f(\mathbb{E}_{\mathbb{P}}[X], \mathbb{E}_{\mathbb{Q}}[X]) \leq \operatorname{Div}_f(\mathbb{P}, \mathbb{Q}),$$

where $\operatorname{div}_f(p,q)$ denotes the f-divergence between Bernoulli distributions with respective parameters p and q.

Corollary 24 (Joint convexity of Div_f). All f-divergences Div_f are jointly convex, i.e., for all probability distributions $\mathbb{P}_1, \mathbb{P}_2$ and $\mathbb{Q}_1, \mathbb{Q}_2$ over the same measurable space (Ω, \mathcal{F}) , and all $\lambda \in (0, 1)$,

$$\operatorname{Div}_f\left((1-\lambda)\mathbb{P}_1+\lambda\mathbb{P}_2,\,(1-\lambda)\mathbb{Q}_1+\lambda\mathbb{Q}_2\right)\leqslant (1-\lambda)\operatorname{Div}_f(\mathbb{P}_1,\mathbb{Q}_1)+\lambda\operatorname{Div}_f(\mathbb{P}_2,\mathbb{Q}_2)\,.$$

Lemma 22 and Corollary 24 are folklore knowledge; we provide here complete and elementary proofs mostly for the sake of self-completeness. These proofs are extracted from Ali and Silvey [1966, Section 4.2], see also Pardo [2006, Proposition 1.2]. They can be refined: Gray [2011, Lemmas 7.5 and 7.6] establishes (70) below and then derives some (stronger) data-processing equality (not inequality). These proof techniques do not seem to be well known; indeed, in the literature many proofs of the elementary properties above for the Kullback-Leibler divergence focus on the discrete case (Cover and Thomas, 2006) or use the duality formula for the Kullback-Leibler divergence (Massart, 2007 or Boucheron et al., 2013, in particular Exercise 4.10 therein).

Proof (of Lemma 22): We recall that $\mathbb{E}_{\mathbb{Q}}$ denotes the expectation with respect to a measure \mathbb{Q} . Let X be a random variable from (Ω, \mathcal{F}) to (Ω', \mathcal{F}') . We write the Lebesgue decomposition (68) of \mathbb{P} with respect to \mathbb{Q} .

We first show that $(\mathbb{P}_{ac})^X \ll \mathbb{Q}^X$ and that the Radon-Nikodym derivative of $(\mathbb{P}_{ac})^X$ with respect to \mathbb{Q}^X equals

$$\frac{\mathrm{d}(\mathbb{P}_{\mathrm{ac}})^X}{\mathrm{d}\mathbb{Q}^X} = \mathbb{E}_{\mathbb{Q}}\left[\frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}} \middle| X = \cdot\right] \stackrel{\mathrm{def}}{=} \gamma ; \tag{70}$$

i.e., γ is any measurable function such that \mathbb{Q} -almost surely, $\mathbb{E}_{\mathbb{Q}}[(d\mathbb{P}_{ac}/d\mathbb{Q})|X] = \gamma(X)$. Indeed, using that $\mathbb{P}_{ac} \ll \mathbb{Q}$, we have, for all $A \in \mathcal{F}'$,

$$(\mathbb{P}_{\mathrm{ac}})^{X}(A) = \mathbb{P}_{\mathrm{ac}}(X \in A) = \int_{\Omega} \mathbb{1}_{A}(X) \frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}} \,\mathrm{d}\mathbb{Q} = \int_{\Omega} \mathbb{1}_{A}(X) \,\mathbb{E}_{\mathbb{Q}}\left[\frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}} \,\Big| X\right] \,\mathrm{d}\mathbb{Q}$$
(71)
$$= \int_{\Omega} \mathbb{1}_{A}(X) \,\gamma(X) \,\mathrm{d}\mathbb{Q} = \int_{\Omega'} \mathbb{1}_{A} \,\gamma \,\mathrm{d}\mathbb{Q}^{X} \,,$$

where the last equality in (71) follows by the tower rule.

Second, by unicity of the Lebesgue decomposition, the decomposition of \mathbb{P}^X with respect to \mathbb{Q}^X is therefore given by

$$\begin{split} \mathbb{P}^{X} &= \left(\mathbb{P}^{X}\right)_{\mathrm{ac}} + \left(\mathbb{P}^{X}\right)_{\mathrm{sing}} & \text{where} & \left(\mathbb{P}^{X}\right)_{\mathrm{ac}} = \left(\mathbb{P}_{\mathrm{ac}}\right)^{X} + \left(\mathbb{P}_{\mathrm{sing}}\right)^{X}_{\mathrm{ac}} \\ & \text{and} & \left(\mathbb{P}^{X}\right)_{\mathrm{sing}} = \left(\mathbb{P}_{\mathrm{sing}}\right)^{X}_{\mathrm{sing}}. \end{split}$$

The inner ac and sing symbols refer to the pair \mathbb{P}, \mathbb{Q} while the outer ac and sing symbols refer to $\mathbb{P}^X, \mathbb{Q}^X$.

We use this decomposition for the first equality below and integrate (67) for the first inequality below:

$$\operatorname{Div}_{f}(\mathbb{P}^{X}, \mathbb{Q}^{X}) = \int_{\Omega'} f\left(\frac{\mathrm{d}(\mathbb{P}_{\mathrm{ac}})^{X}}{\mathrm{d}\mathbb{Q}^{X}} + \frac{\mathrm{d}(\mathbb{P}_{\mathrm{sing}})^{X}}{\mathrm{d}\mathbb{Q}^{X}}\right) \mathrm{d}\mathbb{Q}^{X} + \left(\mathbb{P}_{\mathrm{sing}}\right)^{X}_{\mathrm{sing}}(\Omega') M_{f}$$

$$\leq \int_{\Omega'} f\left(\frac{\mathrm{d}(\mathbb{P}_{\mathrm{ac}})^{X}}{\mathrm{d}\mathbb{Q}^{X}}\right) \mathrm{d}\mathbb{Q}^{X} + \left(\left(\mathbb{P}_{\mathrm{sing}}\right)^{X}_{\mathrm{ac}}(\Omega') + \left(\mathbb{P}_{\mathrm{sing}}\right)^{X}_{\mathrm{sing}}(\Omega')\right) M_{f}$$

$$= \int_{\Omega'} f(\gamma) \mathrm{d}\mathbb{Q}^{X} + (\mathbb{P}_{\mathrm{sing}})^{X}(\Omega') M_{f}$$

$$= \int_{\Omega} f\left(\gamma(X)\right) \mathrm{d}\mathbb{Q} + \mathbb{P}_{\mathrm{sing}}(\Omega) M_{f}$$

$$= \int_{\Omega} f\left(\mathbb{E}_{\mathbb{Q}}\left[\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}} \middle| X\right]\right) \mathrm{d}\mathbb{Q} + \mathbb{P}_{\mathrm{sing}}(\Omega) M_{f}$$

$$\leq \int_{\Omega} \mathbb{E}_{\mathbb{Q}}\left[f\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) \middle| X\right] \mathrm{d}\mathbb{Q} + \mathbb{P}_{\mathrm{sing}}(\Omega) M_{f}$$

$$(72)$$

$$= \int_{\Omega} f\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) \mathrm{d}\mathbb{Q} + \mathbb{P}_{\mathrm{sing}}(\Omega) M_f = \mathrm{Div}_f(\mathbb{P}, \mathbb{Q}) \,,$$

where the inequality in (72) is a consequence of the conditional Jensen's inequality in its general form stated in Appendix F, Lemma 29, with $\varphi = f$ and $C = [0, +\infty)$, and where the final equality follows from the tower rule.

We continue with the proof of Corollary 23, which is (almost) extracted from Garivier et al. [2017]: it was proved therein for Kullback-Leibler divergences.

Proof (of Corollary 23): We augment the underlying measurable space into $\Omega \times [0, 1]$, where [0, 1] is equipped with the Borel σ -algebra $\mathcal{B}([0, 1])$ and the Lebesgue measure \mathfrak{m} . We denote by $\mathbb{P} \otimes \mathfrak{m}$ and $\mathbb{Q} \otimes \mathfrak{m}$ the product distributions of \mathbb{P} and \mathfrak{m} , \mathbb{Q} and \mathfrak{m} . We write the Lebesgue decomposition $\mathbb{P} = \mathbb{P}_{ac} + \mathbb{P}_{sing}$ of \mathbb{P} with respect to \mathbb{Q} , and deduce from it the Lebesgue decomposition of $\mathbb{P} \otimes \mathfrak{m}$ with respect to $\mathbb{Q} \otimes \mathfrak{m}$: the absolutely continuous part is given by $\mathbb{P}_{ac} \otimes \mathfrak{m}$, with density

$$(\omega, x) \in \Omega \times [0, 1] \longmapsto \frac{\mathrm{d}(\mathbb{P}_{\mathrm{ac}} \otimes \mathfrak{m})}{\mathrm{d}(\mathbb{Q} \otimes \mathfrak{m})} (\omega, x) = \frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}} (\omega) \,,$$

while the singular part is given by $\mathbb{P}_{sing} \otimes \mathfrak{m}$, a subprobability with total mass $\mathbb{P}_{sing}(\Omega)$. In particular,

$$\operatorname{Div}_f(\mathbb{P}\otimes\mathfrak{m},\,\mathbb{Q}\otimes\mathfrak{m})=\operatorname{Div}_f(\mathbb{P},\mathbb{Q})$$

Now, for all events $E \in \mathcal{F} \otimes \mathcal{B}([0,1])$, the data-processing inequality (Lemma 22) ensures that

$$\operatorname{Div}_f(\mathbb{P}\otimes\mathfrak{m}, \mathbb{Q}\otimes\mathfrak{m}) \ge \operatorname{Div}_f((\mathbb{P}\otimes\mathfrak{m})^{\mathbb{1}_E}, (\mathbb{Q}\otimes\mathfrak{m})^{\mathbb{1}_E}) = \operatorname{div}_f((\mathbb{P}\otimes\mathfrak{m})(E), (\mathbb{Q}\otimes\mathfrak{m})(E)),$$

where the final equality is by mere definition of div_f as the f-divergence between Bernoulli distributions. The proof is concluded by noting that for the choice of $E = \{(\omega, x) \in \Omega \times [0, 1] : x \leq X(\omega)\}$, Tonelli's theorem ensures that

$$(\mathbb{P} \otimes \mathfrak{m})(E) = \int_{\Omega} \left(\int_{[0,1]} \mathbb{1}_{\left\{ x \leq X(\omega) \right\}} \, \mathrm{d}\mathfrak{m}(x) \right) \mathrm{d}\mathbb{P}(\omega) = \mathbb{E}_{\mathbb{P}}[X] \,,$$

and, similarly, $(\mathbb{Q} \otimes \mathfrak{m})(E) = \mathbb{E}_{\mathbb{Q}}[X]$.

The joint convexity of Div_f (Corollary 24) may be proved directly, in two steps. First, the log-sum inequality is generalized into the fact that the mapping $(p,q) \in [0, +\infty)^2 \mapsto q f(p/q)$ is jointly convex. Second, a common dominating measure like $\mu = \mathbb{P}_1 + \mathbb{P}_2 + \mathbb{Q}_1 + \mathbb{Q}_2$ is introduced, Radon-Nikodym derivatives p_j and q_j are introduced for the \mathbb{P}_j and \mathbb{Q}_j with respect to μ , and the generalized log-sum inequality is applied pointwise.

We suggest to see instead Corollary 24 as an elementary consequence of the data-processing inequality.

Proof (of Corollary 24): We augment the probability space Ω into $\Omega' = \{1, 2\} \times \Omega$ equipped with the σ -algebra \mathcal{F}' generated by the events $A \times B$, where $A \in \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$ and $B \in \mathcal{F}$. We define the random pair (J, X) on this space by the projections

$$X: (j,\omega) \in \{1,2\} \times \Omega \longmapsto \omega \qquad \text{and} \qquad J: (j,\omega) \in \{1,2\} \times \Omega \longmapsto j \,,$$

and denote by \mathbb{P} the joint distribution of the random pair (J, X) such that $J \sim 1 + \text{Ber}(\lambda)$ and $X|J \sim \mathbb{P}_J$. More formally, \mathbb{P} is the unique probability distribution on (Ω', \mathcal{F}') such that, for all $(j, B) \in \{1, 2\} \times \mathcal{F}$,

$$\mathbb{P}(\{j\} \times B) = ((1-\lambda)\mathbb{1}_{\{j=1\}} + \lambda\mathbb{1}_{\{j=2\}})\mathbb{P}_j(B).$$

Similarly we define the joint probability distribution \mathbb{Q} on (Ω', \mathcal{F}') using the conditional distributions \mathbb{Q}_1 and \mathbb{Q}_2 instead of \mathbb{P}_1 and \mathbb{P}_2 .

The corollary follows directly from the data-processing inequality $\operatorname{Div}_f(\mathbb{P}^X, \mathbb{Q}^X) \leq \operatorname{Div}_f(\mathbb{P}, \mathbb{Q})$, as the laws of X under \mathbb{P} and \mathbb{Q} are respectively given by

$$\mathbb{P}^X = (1 - \lambda)\mathbb{P}_1 + \lambda\mathbb{P}_2$$
 and $\mathbb{Q}^X = (1 - \lambda)\mathbb{Q}_1 + \lambda\mathbb{Q}_2$,

while elementary calculations show that $\operatorname{Div}_f(\mathbb{P}, \mathbb{Q}) = (1 - \lambda) \operatorname{Div}_f(\mathbb{P}_1, \mathbb{Q}_1) + \lambda \operatorname{Div}_f(\mathbb{P}_2, \mathbb{Q}_2).$

Indeed, for the latter point, we consider the Lebesgue decompositions of \mathbb{P}_j with respect to \mathbb{Q}_j , where $j \in \{1, 2\}$:

 $\mathbb{P}_j = \mathbb{P}_{j,\mathrm{ac}} + \mathbb{P}_{j,\mathrm{sing}}, \quad \text{where} \quad \mathbb{P}_{j,\mathrm{ac}} \ll \mathbb{Q}_j \text{ and } \mathbb{P}_{j,\mathrm{sing}} \perp \mathbb{Q}_j.$

The (unique) Lebesgue decomposition of $\mathbb{P} = \mathbb{P}_{ac} + \mathbb{P}_{sing}$ with respect to \mathbb{Q} is then given by

$$\frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}}(j,\omega) = \mathbbm{1}_{\{j=1\}} \ \frac{\mathrm{d}\mathbb{P}_{1,\mathrm{ac}}}{\mathrm{d}\mathbb{Q}_1}(\omega) + \mathbbm{1}_{\{j=2\}} \ \frac{\mathrm{d}\mathbb{P}_{2,\mathrm{ac}}}{\mathrm{d}\mathbb{Q}_2}(\omega)$$

and for all $(j, B) \in \{1, 2\} \times \mathcal{F}$,

$$\mathbb{P}_{\text{sing}}(\{j\} \times B) = ((1-\lambda)\mathbb{1}_{\{j=1\}} + \lambda \mathbb{1}_{\{j=2\}}) \mathbb{P}_{j,\text{sing}}(B).$$

This entails that

$$\begin{aligned} \operatorname{Div}_{f}(\mathbb{P},\mathbb{Q}) &= \int_{\{1,2\}\times\Omega} f\left(\frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}}(j,\omega)\right) \mathrm{d}\mathbb{Q}(j,\omega) + \mathbb{P}_{\mathrm{sing}}\big(\{1,2\}\times\Omega\big) M_{f} \\ &= (1-\lambda) \int_{\Omega} f\left(\frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}}(1,\omega)\right) \mathrm{d}\mathbb{Q}_{1}(\omega) + \lambda \int_{\Omega} f\left(\frac{\mathrm{d}\mathbb{P}_{\mathrm{ac}}}{\mathrm{d}\mathbb{Q}}(2,\omega)\right) \mathrm{d}\mathbb{Q}_{2}(\omega) \\ &+ \big((1-\lambda) \mathbb{P}_{1,\mathrm{sing}}(\Omega) + \lambda \mathbb{P}_{2,\mathrm{sing}}(\Omega)\big) M_{f} \end{aligned}$$
$$= (1-\lambda) \operatorname{Div}_{f}(\mathbb{P}_{1},\mathbb{Q}_{1}) + \lambda \operatorname{Div}_{f}(\mathbb{P}_{2},\mathbb{Q}_{2}).$$

E. Extensions of the reductions of Section 3 to *f*-divergences

We recall that f-divergences are based on convex functions $f:(0,+\infty)\to\mathbb{R}$ extended at 0 via

$$f(0) \stackrel{\text{def}}{=} \lim_{t \downarrow 0} f(t) \in \mathbb{R} \cup \{+\infty\}$$

and such that f(1) = 0.

E.1. Reduction to Bernoulli distributions

We denote by div_f the f-divergence between Bernoulli distributions: for all $(p,q) \in [0,1]^2$,

$$\operatorname{div}_f(p,q) \stackrel{\operatorname{def}}{=} \operatorname{Div}_f(\operatorname{Ber}(p),\operatorname{Ber}(q)).$$

Because f-divergences are also jointly convex and enjoy a data-processing inequality (see Lemma 22 and Corollaries 23 and 24 in Appendix D) the various reductions considered in Section 3.1 hold as

well. We only illustrate the reduction by considering the simplest one, stated in (5), and the most general one, stated in (9); with the notation used therein, we have

$$\operatorname{div}_{f}\left(\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}), \frac{1}{N}\sum_{i=1}^{N}\mathbb{Q}_{i}(A_{i})\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{div}_{f}\left(\mathbb{P}_{i}(A_{i}), \mathbb{Q}_{i}(A_{i})\right) \leqslant \frac{1}{N}\sum_{i=1}^{N}\operatorname{Div}_{f}(\mathbb{P}_{i}, \mathbb{Q}_{i}), \quad (73)$$

and (putting aside all measurability issues)

$$\operatorname{div}_{f}\left(\int_{\Theta} \mathbb{E}_{\mathbb{P}_{\theta}}\left[Z_{\theta}\right] \mathrm{d}\nu(\theta), \int_{\Theta} \mathbb{E}_{\mathbb{Q}_{\theta}}\left[Z_{\theta}\right] \mathrm{d}\nu(\theta)\right) \leqslant \int_{\Theta} \operatorname{div}_{f}\left(\mathbb{E}_{\mathbb{P}_{\theta}}\left[Z_{\theta}\right], \mathbb{E}_{\mathbb{Q}_{\theta}}\left[Z_{\theta}\right]\right) \mathrm{d}\nu(\theta)$$
$$\leqslant \int_{\Theta} \operatorname{Div}_{f}(\mathbb{P}_{\theta}, \mathbb{Q}_{\theta}) \mathrm{d}\nu(\theta).$$
(74)

It thus suffices to lower bound div_f to obtain bounds of interest, as we did for kl in Section 3.2. We propose below such lower bounds for the χ^2 divergence and the Hellinger distance.

E.2. Lower bound on div_f for the χ^2 divergence

This case corresponds to $f(x) = x^2 - 1$. The associated divergence equals, when $\mathbb{P} \ll \mathbb{Q}$,

$$\chi^2(\mathbb{P},\mathbb{Q}) = \int_{\Omega} \left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right)^2 \mathrm{d}\mathbb{Q} - 1.$$

A direct calculation indicates that for all $(p,q) \in [0,1]^2$,

$$\chi^2\big(\mathrm{Ber}(p),\mathrm{Ber}(q)\big) = \frac{(p-q)^2}{q(1-q)} \ge \frac{(p-q)^2}{q}$$

in this case. We get, for instance, the following result based on the reduction (73), which corresponds to Proposition 1 for Kullback-Leibler divergences.

Lemma 25. Given an underlying measurable space, for all probability pairs \mathbb{P}_i , \mathbb{Q}_i and all events A_i (non necessarily disjoint), where $i \in \{1, \ldots, N\}$, with $0 < \frac{1}{N} \sum_{i=1}^N \mathbb{Q}_i(A_i) < 1$, we have

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{1}{N}\sum_{i=1}^{N}\mathbb{Q}_{i}(A_{i}) + \sqrt{\frac{1}{N}\sum_{i=1}^{N}\mathbb{Q}_{i}(A_{i})}\sqrt{\frac{1}{N}\sum_{i=1}^{N}\chi^{2}(\mathbb{P}_{i},\mathbb{Q}_{i})}.$$
(75)

In particular, if $N \ge 2$ and the A_i form a partition,

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{1}{N} + \frac{1}{\sqrt{N}}\sqrt{\frac{1}{N}\inf_{\mathbb{Q}}\sum_{i=1}^{N}\chi^{2}(\mathbb{P}_{i},\mathbb{Q})}.$$

E.3. Lower bound on div_f for the Hellinger distance

This case corresponds to $f(x) = (\sqrt{x} - 1)^2$, for which $M_f = 1$. The associated divergence equals, when $\mathbb{P} \ll \mathbb{Q}$,

$$H^{2}(\mathbb{P},\mathbb{Q}) = \int_{\Omega} \left(\sqrt{\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}} - 1 \right)^{2} \mathrm{d}\mathbb{Q} = 2 \left(1 - \int_{\Omega} \sqrt{\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}} \, \mathrm{d}\mathbb{Q} \right)$$

and always lies in [0,2]. A direct calculation indicates that for all $p \in [0,1]$ and $q \in (0,1)$,

$$h^{2}(p,q) \stackrel{\text{def}}{=} H^{2}\left(\text{Ber}(p), \text{Ber}(q)\right) = 2\left(1 - \left(\sqrt{pq} + \sqrt{(1-p)(1-q)}\right)\right),$$

and further direct calculations in the cases q = 0 and q = 1 show that this formula remains valid in these cases.

Resorting to the Cauchy-Schwarz inequality. The Cauchy-Schwarz inequality indicates that

$$\sqrt{pq} + \sqrt{(1-q)(1-p)} \le \sqrt{\left(p + (1-q)\right)\left(q + (1-p)\right)} = \sqrt{1 - (p-q)^2}$$

or put differently, that $h^2(p,q) \ge 2\left(1 - \sqrt{1 - (p-q)^2}\right)$, thus

$$p \leqslant q + \sqrt{1 - \left(1 - h^2(p, q)/2\right)^2} = q + \sqrt{h^2(p, q)\left(1 - h^2(p, q)/4\right)},$$
(76)

which is one of Le Cam's inequalities. This bound is clean and clear enough for the reader to be able to state consequences of it, e.g., in the spirit of Proposition 1 for Kullback-Leibler divergences or Lemma 25 for χ^2 divergences. In particular, if the A_i form a partition,

$$\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}_{i}(A_{i}) \leqslant \frac{1}{N} + \inf_{\mathbb{Q}}\sqrt{\frac{1}{N}\sum_{i=1}^{N}H^{2}(\mathbb{P}_{i},\mathbb{Q})}\sqrt{1 - \frac{1}{4N}\sum_{i=1}^{N}H^{2}(\mathbb{P}_{i},\mathbb{Q})}.$$

Solving for p. This is the path followed by Guntuboyina [2011, Example II.6]; as we prove below, we get

$$p \leq q + (1 - 2q) h^2(p, q) \left(1 - h^2(p, q)/4\right) + 2\sqrt{q(1 - q)} \left(1 - h^2(p, q)/2\right) \sqrt{h^2(p, q) \left(1 - h^2(p, q)/4\right)} .$$
(77)

It can be seen that this bound is a general expression of the bound stated by Guntuboyina [2011, Example II.6]. This bound is slightly tighter than (76), by construction (as we solve exactly an equation and perform no bounding) but it is much less readable. It anyway leads to similar conclusions in practice.

Proof: Assuming that $\underline{h}^2 = h^2(p,q)$ is given and fixed, we consider the equation, for the unknown $x \in [0,1]$,

$$\underline{h}^{2} = 2\left(1 - \left(\sqrt{q}\sqrt{x} + \sqrt{1-q}\sqrt{1-x}\right)\right);$$

this equation is satisfied for x = p, by definition of $h^2(p,q)$. Rearranging it, we get the equivalent equation

$$(1-x)(1-q) = \left(1 - \frac{h^2}{2} - \sqrt{q}\sqrt{x}\right)^2 = \left(1 - \frac{h^2}{2}\right)^2 - 2\left(1 - \frac{h^2}{2}\right)\sqrt{q}\sqrt{x} + qx,$$

or equivalently again,

$$x - 2(1 - \underline{h}^2/2)\sqrt{q}\sqrt{x} + (1 - \underline{h}^2/2)^2 - 1 + q = 0.$$

Solving this second-order equation for \sqrt{x} , we see that all solutions \sqrt{x} , including \sqrt{p} , are smaller than the largest root; in particular,

$$\sqrt{p} \leqslant (1 - \underline{h}^2/2)\sqrt{q} + \underbrace{\sqrt{(1 - \underline{h}^2/2)^2 q - (1 - \underline{h}^2/2)^2 + 1 - q}}_{=\sqrt{(1 - q)\underline{h}^2(1 - \underline{h}^2/4)}} .$$

Put differently,

$$p \leq (1 - \underline{h}^2/2)^2 q + (1 - q) \underline{h}^2 (1 - \underline{h}^2/4) + 2\sqrt{q(1 - q)} (1 - \underline{h}^2/2) \sqrt{\underline{h}^2 (1 - \underline{h}^2/4)}$$

= $q + (1 - 2q) h^2(p, q) (1 - h^2(p, q)/4) + 2\sqrt{q(1 - q)} (1 - h^2(p, q)/2) \sqrt{h^2(p, q) (1 - h^2(p, q)/4)},$

which was the expression to obtain.

E.4. Finding a good constant alternative \mathbb{Q}

Consider, for example, the reduction based on a convex combination $\alpha = (\alpha_1, \ldots, \alpha_N)$, with all $\alpha_i > 0$:

$$\operatorname{div}_f\left(\sum_{i=1}^N \alpha_i \mathbb{P}_i(A_i), \sum_{i=1}^N \alpha_i \mathbb{Q}_i(A_i)\right) \leqslant \sum_{i=1}^N \alpha_i \operatorname{div}_f\left(\mathbb{P}_i(A_i), \mathbb{Q}_i(A_i)\right) \leqslant \sum_{i=1}^N \alpha_i \operatorname{Div}_f(\mathbb{P}_i, \mathbb{Q}_i),$$

which is more general than (73) but less general than (74).

We wonder, under the constraint that only one fixed alternative distribution $\mathbb{Q}_i = \mathbb{Q}$ is considered, which such alternative to pick. That is, we want to compute or at least upper bound

$$\inf_{\mathbb{Q}} \sum_{i=1}^{N} \alpha_i \operatorname{Div}_f(\mathbb{P}_i, \mathbb{Q});$$

distributions \mathbb{Q} that (approximatively) reach the infimum should be used, at least from a theoretical viewpoint. Sometimes calculations are easier in practice for some specific \mathbb{Q} , as we illustrated, for instance, in Section 4.2. Otherwise, the lemma below indicates a good candidate, given by the weighted average $\overline{\mathbb{P}}_{\alpha}$ of the distributions \mathbb{P}_i .

To appreciate its performance, we denote by

$$B_f(\alpha) = \max_{j=1,\dots,N} \operatorname{Div}_f(\delta_j, \alpha)$$

the maximal f-divergence between a Dirac mass δ_j at j and the convex combination α . This bound equals $\ln(1/\min\{\alpha_1,\ldots,\alpha_N\})$ for a Kullback-Leibler divergence and $1/\min\{\alpha_1,\ldots,\alpha_N\} - 1$ for the χ^2 -divergence.

Lemma 26. Let $\mathbb{P}_1, \ldots, \mathbb{P}_N$ be N probability distributions over the same measurable space (Ω, \mathcal{F}) and let $\alpha = (\alpha_1, \ldots, \alpha_N)$ be a convex combination made of positive weights. Then,

$$\inf_{\mathbb{Q}} \sum_{i=1}^{N} \alpha_{i} \operatorname{Div}_{f}(\mathbb{P}_{i}, \mathbb{Q}) \leqslant \sum_{i=1}^{N} \alpha_{i} \operatorname{Div}_{f}(\mathbb{P}_{i}, \overline{\mathbb{P}}_{\alpha}) \leqslant B_{f}(\alpha),$$

where the infimum is over all probability distributions \mathbb{Q} on (Ω, \mathcal{F}) and where $\overline{\mathbb{P}}_{\alpha} \stackrel{\text{def}}{=} \sum_{i=1}^{N} \alpha_{i} \mathbb{P}_{i}$.

The first inequality holds with equality in the case of the Kullback-Leibler divergence, as follows from the so-called compensation equality (see, e.g., Yang and Barron, 1999 or Guntuboyina, 2011, Example II.4): assuming with no loss of generality in this case (since $M_f = +\infty$) that $\mathbb{P}_j \ll \mathbb{Q}$ for all $j \in \{1, \ldots, N\}$, we have $\overline{\mathbb{P}}_{\alpha} \ll \mathbb{Q}$ and $d\mathbb{P}_j/d\mathbb{Q} = (d\mathbb{P}_j/d\overline{\mathbb{P}}_{\alpha})(d\overline{\mathbb{P}}_{\alpha}/d\mathbb{Q})$, which entails

$$\sum_{i=1}^{N} \alpha_i \operatorname{KL}(\mathbb{P}_i, \mathbb{Q}) = \sum_{i=1}^{N} \alpha_i \int \left(\ln \frac{\mathrm{d}\mathbb{P}_i}{\mathrm{d}\overline{\mathbb{P}}_{\alpha}} + \ln \frac{\mathrm{d}\overline{\mathbb{P}}_{\alpha}}{\mathrm{d}\mathbb{Q}} \right) \mathrm{d}\mathbb{P}_i = \left(\sum_{i=1}^{N} \alpha_i \operatorname{KL}(\mathbb{P}_i, \overline{\mathbb{P}}_{\alpha}) \right) + \operatorname{KL}(\overline{\mathbb{P}}_{\alpha}, \mathbb{Q}),$$

where we used that $\sum_{i=1}^{N} \alpha_i \, \mathrm{d}\mathbb{P}_i = \, \mathrm{d}\overline{\mathbb{P}}_{\alpha}$. So, indeed, the considered infimum is achieved at $\mathbb{Q} = \overline{\mathbb{P}}$.

Proof: The first inequality follows from the choice $\mathbb{Q} = \overline{\mathbb{P}}_{\alpha}$. For the second inequality, we proceed as in Corollary 24 and consider the following probability distributions over $\{1, \ldots, N\} \times \Omega$: for all $j \in \{1, \ldots, N\}$ and all $B \in \mathcal{F}$,

 $\widetilde{\mathbb{P}}(\{j\} \times B) = \alpha_j \mathbb{P}_j(B)$ and $\widetilde{\mathbb{Q}}(\{j\} \times B) = \alpha_j \overline{\mathbb{P}}_\alpha(B)$.

Note that because $\alpha_i > 0$ for all i, we have $\mathbb{P}_j \ll \overline{\mathbb{P}}_{\alpha}$ for all j. Thus, $\widetilde{\mathbb{P}} \ll \widetilde{\mathbb{Q}}$, with Radon-Nikodym derivative given by

$$(j,\omega) \in \{1,\ldots,N\} \times \Omega \longmapsto \frac{\mathrm{d}\widetilde{\mathbb{P}}}{\mathrm{d}\widetilde{\mathbb{Q}}}(j,\omega) = \frac{\mathrm{d}\mathbb{P}_j}{\mathrm{d}\overline{\mathbb{P}}_\alpha}(\omega) \stackrel{\mathrm{def}}{=} p_j(\omega).$$

By uniqueness and linearity of the Radon-Nikodym derivatives, we thus have, for $\overline{\mathbb{P}}_{\alpha}$ -almost all ω ,

$$\sum_{j=1}^{N} \alpha_j \, p_j(\omega) = \sum_{j=1}^{N} \alpha_j \, \frac{\mathrm{d}\mathbb{P}_j}{\mathrm{d}\overline{\mathbb{P}}_\alpha}(\omega) = \frac{\mathrm{d}\overline{\mathbb{P}}_\alpha}{\mathrm{d}\overline{\mathbb{P}}_\alpha}(\omega) = 1 \,, \qquad \text{where} \qquad \forall k \in \{1, \dots, N\}, \quad \alpha_k p_k(\omega) \ge 0 \,;$$

that is, $\alpha p(\omega) = (\alpha_j p_j(\omega))_{1 \leq j \leq N}$ is a probability distribution over $\{1, \ldots, N\}$. (It corresponds to the conditional distribution of j given ω in the probabilistic model $j \sim \alpha$ and $\omega | j \sim \mathbb{P}_j$.)

We now compute $\operatorname{Div}_f(\widetilde{\mathbb{P}}, \widetilde{\mathbb{Q}})$ in two different ways. All manipulations below are valid because all integrals defining f-divergences exist (see the comments after the statement of Definition D.1, as well as the first part of the proof of Lemma 28). Integrating over j first,

$$\operatorname{Div}_{f}(\widetilde{\mathbb{P}}, \widetilde{\mathbb{Q}}) = \int_{\{1, \dots, N\} \times \Omega} f\left(\frac{\mathrm{d}\widetilde{\mathbb{P}}}{\mathrm{d}\widetilde{\mathbb{Q}}}(j, \omega)\right) \mathrm{d}\widetilde{\mathbb{Q}}(j, \omega)$$
$$= \sum_{j=1}^{N} \alpha_{j} \int_{\Omega} f\left(\frac{\mathrm{d}\mathbb{P}_{j}}{\mathrm{d}\overline{\mathbb{P}}_{\alpha}}(\omega)\right) \mathrm{d}\overline{\mathbb{P}}_{\alpha}(\omega) = \sum_{j=1}^{N} \alpha_{j} \operatorname{Div}_{f}(\mathbb{P}_{j}, \overline{\mathbb{P}}_{\alpha})$$

On the other hand, integrating over ω first,

$$\operatorname{Div}_{f}(\widetilde{\mathbb{P}}, \widetilde{\mathbb{Q}}) = \int_{\Omega} \left(\sum_{j=1}^{N} f(p_{j}(\omega)) \alpha_{j} \right) d\overline{\mathbb{P}}_{\alpha}(\omega)$$
$$= \int_{\Omega} \left(\sum_{j=1}^{N} f\left(\frac{\alpha_{j} p_{j}(\omega)}{\alpha_{j}}\right) \alpha_{j} \right) d\overline{\mathbb{P}}_{\alpha}(\omega) = \int_{\Omega} \operatorname{Div}_{f}(\alpha p(\omega), \alpha) d\overline{\mathbb{P}}_{\alpha}(\omega) \leqslant B_{f}(\alpha),$$

where the last inequality follows by noting that, by joint convexity of Div_f (see Corollary 3),

$$\operatorname{Div}_f(\alpha p(\omega), \alpha) \leq \sum_{j=1}^n \alpha_j p_j(\omega) \operatorname{Div}_f(\delta_j, \alpha) \leq B_f(\alpha).$$

Comparing the two obtained expressions for $\operatorname{Div}_f(\widetilde{\mathbb{P}}, \widetilde{\mathbb{Q}})$ concludes the proof.

F. On Jensen's inequality

Classical statements of Jensen's inequality for convex functions φ on $C \subseteq \mathbb{R}^n$ either assume that the underlying probability measure is supported on a finite number of points or that the convex subset C is open. In the first case, the proof follows directly from the definition of convexity, while in the second case, it is a consequence of the existence of subgradients. In both cases, it is assumed that the function φ under consideration only takes finite values. In this article, Jensen's inequality is applied several times to non-open convex sets C, like $C = [0, 1]^2$ or $C = [0, +\infty)$ and/or convex functions φ that can possibly be equal to $+\infty$ at some points.

The restriction of C being open is easy to drop when the dimension equals n = 1, i.e., when C is an interval; it was dropped, e.g., by Ferguson [1967, pages 74–76] in higher dimensions, thanks to a proof by induction to address possible boundary effects with respect to the arbitrary convex set C. Let $\mathcal{B}(\mathbb{R}^n)$ denote the Borel σ -field of \mathbb{R}^n . **Lemma 27** (Jensen's inequality for general convex sets; Ferguson, 1967). Let $C \subseteq \mathbb{R}^n$ be any nonempty convex Borel subset of \mathbb{R}^n and $\varphi: C \to \mathbb{R}$ be any convex Borel function. Then, for all probability measures μ on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ such that $\mu(C) = 1$ and $\int ||x|| d\mu(x) < +\infty$, we have

$$\int x \, \mathrm{d}\mu(x) \in C \qquad and \qquad \varphi\left(\int x \, \mathrm{d}\mu(x)\right) \leqslant \int_C \varphi(x) \, \mathrm{d}\mu(x) \,, \tag{78}$$

where the integral of φ against μ is well-defined in $\mathbb{R} \cup \{+\infty\}$.

Our contribution is the following natural extension.

Lemma 28. The result of Lemma 27 also holds for any convex Borel function $\varphi : C \to \mathbb{R} \cup \{+\infty\}$.

We rephrase this extension in terms of random variables. Let $C \subseteq \mathbb{R}^n$ be any non-empty convex Borel subset of \mathbb{R}^n and $\varphi : C \to \mathbb{R} \cup \{+\infty\}$ be any convex Borel function. Let X be an integrable random variable from any probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$, such that $\mathbb{P}(X \in C) = 1$. Then

$$\mathbb{E}[X] \in C$$
 and $\varphi(\mathbb{E}[X]) \leq \mathbb{E}[\varphi(X)]$,

where $\mathbb{E}[\varphi(X)]$ is well-defined in $\mathbb{R} \cup \{+\infty\}$.

Proof: We first check that $\varphi_{-} = \max\{-\varphi, 0\}$ is μ -integrable on C, so that the integral of φ against μ is well-defined in $\mathbb{R} \cup \{+\infty\}$. To that end, we will prove that φ is lower bounded on C by an affine function: $\varphi(x) \ge a^{\mathrm{T}}x + b$ for all $x \in C$, where $(a, b) \in \mathbb{R}^2$, from which it follows that $\varphi_{-}(x) \le ||a|| ||x|| + ||b||$ for all $x \in C$ and thus

$$\int_{C} \varphi_{-}(x) \, \mathrm{d}\mu(x) \leqslant \int_{C} \left(\|a\| \|x\| + \|b\| \right) \mathrm{d}\mu(x) = \|a\| \int_{C} \|x\| \, \mathrm{d}\mu(x) + \|b\| < +\infty \,.$$

So, it only remains to prove the affine lower bound. If the domain $\{\varphi < +\infty\}$ is empty, any affine function is suitable. Otherwise, $\{\varphi < +\infty\}$ is a non-empty convex set, so that its relative interior R is also non-empty (see Rockafellar, 1972, Theorem 6.2); we fix $x_0 \in R$. But, by Rockafellar [1972, Theorem 23.4], the function φ admits a subgradient at x_0 , that is, there exists $a \in \mathbb{R}^n$ such that $\varphi(x) \ge \varphi(x_0) + a^T(x - x_0)$ for all $x \in C$. This concludes the first part of this proof.

In the second part, we show the inequality (78) via a reduction to the case of real-valued functions. Indeed, note that if $\mu(\varphi = +\infty) > 0$ then the desired inequality is immediate. We can thus assume that $\mu(\varphi < +\infty) = 1$. But, using Lemma 27 with the non-empty convex Borel subset $\widetilde{C} = \{\varphi < +\infty\}$ and the real-valued convex Borel function $\widetilde{\varphi} : \widetilde{C} \to \mathbb{R}$ defined by $\widetilde{\varphi}(x) = \varphi(x)$, we get, since $\mu(\widetilde{C}) = 1$:

$$\int x \, \mathrm{d}\mu(x) \in \widetilde{C} \qquad \text{and} \qquad \widetilde{\varphi}\left(\int x \, \mathrm{d}\mu(x)\right) \leqslant \int_{\widetilde{C}} \widetilde{\varphi}(x) \, \mathrm{d}\mu(x) \, .$$

Using the facts that $\tilde{\varphi}(x) = \varphi(x)$ for all $x \in \tilde{C}$ and that $\mu(C \setminus \tilde{C}) = 1 - 1 = 0$ entails (78).

We now complete our extension by tacking the conditional form of Jensen's inequality.

Lemma 29 (A general conditional Jensen's inequality). Let $C \subseteq \mathbb{R}^n$ be any non-empty convex Borel subset of \mathbb{R}^n and $\varphi : C \to \mathbb{R} \cup \{+\infty\}$ be any convex Borel function. Let X be an integrable random variable from any probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$, such that $\mathbb{P}(X \in C) = 1$. Then, for every sub- σ -field \mathcal{G} of \mathcal{F} , we have, \mathbb{P} -almost surely,

$$\mathbb{E}[X | \mathcal{G}] \in C \quad and \quad \varphi(\mathbb{E}[X | \mathcal{G}]) \leq \mathbb{E}[\varphi(X) | \mathcal{G}],$$

where $\mathbb{E}[\varphi(X) | \mathcal{G}]$ is \mathbb{P} -almost-surely well-defined in $\mathbb{R} \cup \{+\infty\}$.

Proof: The proof follows directly from the unconditional Jensen's inequality (Lemma 28 above) and from the existence of regular conditional distributions. More precisely, by Durrett [2010, Theorems 2.1.15 and 5.1.9] applied to the case where $(S, \mathcal{S}) = (\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$, there exists a regular conditional distribution of X given \mathcal{G} . That is, there exists a function $K : \Omega \times \mathcal{B}(\mathbb{R}^n) \to [0, 1]$ such that:

(P1) for every $B \in \mathcal{B}(\mathbb{R}^n)$, $\omega \in \Omega \mapsto K(\omega, B)$ is \mathcal{G} -measurable and $\mathbb{P}(X \in B \mid \mathcal{G}) = K(\cdot, B) \mathbb{P}$ -a.s.; (P2) for \mathbb{P} -almost all $\omega \in \Omega$, the mapping $B \mapsto K(\omega, B)$ is a probability measure over $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. Moreover, as a consequence of (P1),

(P1') for every Borel function $g: \mathbb{R}^n \to \mathbb{R}$ such that g(X) is \mathbb{P} -integrable or such that g is nonnegative,

$$\int g(x) K(\cdot, dx) = \mathbb{E}[g(X) \mid \mathcal{G}] \qquad \mathbb{P}-\text{a.s.}$$

Now, given our assumptions and thanks to (P1) and (P1'):

(P3) by $\mathbb{P}(X \in C) = 1$ we also have $K(\cdot, C) = \mathbb{P}(X \in C \mid \mathcal{G}) = 1$ \mathbb{P} -a.s.;

(P4) since X is \mathbb{P} -integrable, so is $\int ||x|| K(\cdot, dx) = \mathbb{E}[||X|| |\mathcal{G}]$, which is therefore \mathbb{P} -a.s. finite.

We apply Lemma 28 with the probability measures $\mu_{\omega} = K(\omega, \cdot)$, for those ω for which the properties stated in (P2), (P3) and (P4) actually hold; these ω are \mathbb{P} -almost all elements of Ω . We get, for these ω ,

$$\int x \, K(\omega, \, \mathrm{d}x) \in C \qquad \text{and} \qquad \varphi \left(\int x \, K(\omega, \, \mathrm{d}x) \right) \leqslant \int_C \varphi(x) \, K(\omega, \, \mathrm{d}x) \,,$$

where the integral in the right-hand side is well defined in $\mathbb{R} \cup \{+\infty\}$. Thanks to (P1'), and by decomposing $\varphi(X)$ into $\varphi_{-}(X)$, which is integrable (see the beginning of the proof of Lemma 28), and $\varphi_{+}(X)$, which is nonnegative, we thus have proved that \mathbb{P} -a.s.,

$$\mathbb{E}[X | \mathcal{G}] \in C$$
 and $\varphi(\mathbb{E}[X | \mathcal{G}]) \leq \mathbb{E}[\varphi(X) | \mathcal{G}],$

which concludes the proof.

Gerchinovitz, Ménard, Stoltz