

Minimum KL-divergence on complements of L_1 balls*

Daniel Berend, Peter Harremoës, and Aryeh Kontorovich

Abstract

Pinsker’s widely used inequality upper-bounds the total variation distance $\|P - Q\|_1$ in terms of the Kullback-Leibler divergence $D(P\|Q)$. Although in general a bound in the reverse direction is impossible, in many applications the quantity of interest is actually $D^*(v, Q)$ — defined, for an arbitrary fixed Q , as the infimum of $D(P\|Q)$ over all distributions P that are at least v -far away from Q in total variation. We show that $D^*(v, Q) \leq Cv^2 + O(v^3)$, where $C = C(Q) = 1/2$ for “balanced” distributions, thereby providing a kind of reverse Pinsker inequality. Some of the structural results obtained in the course of the proof may be of independent interest. An application to large deviations is given.

I. INTRODUCTION

A. Pinsker’s inequality

The inequality bearing Pinsker’s name states that for two distributions P and Q ,

$$D(P\|Q) \geq \frac{V^2(P, Q)}{2}, \quad (1)$$

where

$$D(P\|Q) = \int \ln \left(\frac{dP}{dQ} \right) dP$$

is the Kullback-Leibler divergence of P from Q and $V(P, Q) = \|P - Q\|_1$ is their total variation distance. Actually, the name is a bit of a misattribution, since the explicit form of (1) was obtained by Csiszár [8] and Kullback [20] in 1967 and is occasionally referred to by their names. Gradual improvements were obtained by [7], [14], [17], [18], [19], [23], [26], [27], [28] and others; see [25] for a detailed history and the “best possible Pinsker inequality”. Recent extensions to general f -divergences may be found in [15] and [25]. This inequality has become a ubiquitous tool in probability [2], [9], [21], information theory [1], and, more recently, machine learning [5]. It will be useful to define the function $\text{KL}_2 : (0, 1)^2 \rightarrow [0, \infty)$ by

$$\text{KL}_2(p, q) = p \ln \frac{p}{q} + (1 - p) \ln \frac{1 - p}{1 - q}$$

and the so-called Vajda’s tight lower bound L [28]:

$$L(v) = \inf_{P, Q: V(P, Q) = v} D(P\|Q).$$

In [14] an exact parametric equation of the curve $(v, L(v))_{0 < v < \infty}$ in \mathbb{R}^2 was given:

$$\begin{aligned} v(t) &= t \left[1 - \left(\coth t - \frac{1}{t} \right)^2 \right], \\ L(v(t)) &= \ln \frac{t}{\sinh t} + t \coth t - \left(\frac{t}{\sinh t} \right)^2. \end{aligned}$$

Some upper bounds on the KL-divergence in terms of other f -divergences are known [10], [11], [12], and in [13, Lemma 3.10] it is shown that, under some conditions, $D(P\|Q) \leq \|P - Q\|_1 \ln(1/\min Q)$. The latter estimate is vacuous for Q with infinite support. In general, it is impossible to upper-bound $D(P\|Q)$ in terms of $V(P, Q)$, since for every $v \in (0, 2]$ there is a pair of distributions P, Q with $V(P, Q) = v$ and $D(P\|Q) = \infty$ [14]. However, in many applications, the actual quantity of interest is not $D(P\|Q)$ for arbitrary P and Q , but rather

$$D^*(v, Q) = \inf_{P: V(P, Q) \geq v} D(P\|Q). \quad (2)$$

For example, Sanov’s Theorem [6], [9] (which we will say more about below) implies that the probability that the empirical distribution \hat{Q}_n , based on a sample of size n , deviates in ℓ_1 by more than v from the true distribution Q behaves asymptotically as $\exp(-nD^*(v, Q))$. Throughout this paper, we consider a (finite or σ -finite) measure space $(\Omega, \mathcal{F}, \mu)$, and all the distributions in question will be defined on this space and assumed absolutely continuous with respect to μ ; this set of distributions will be denoted by \mathcal{P} . We will consistently use upper-case letters for distributions $P \in \mathcal{P}$ and corresponding lower-case letters for their densities p with respect to μ . We will use standard asymptotic notation $O(\cdot)$ and $\Omega(\cdot)$.

*A previous version had the title “A Reverse Pinsker Inequality”.

B. Balanced and unbalanced distributions

In this paper, we show that for the broad class of “balanced” distributions, $D^*(v, Q) = v^2/2 + O(v^4)$, which matches the form of the bound in (1). For distributions not belonging to this class, we show that

$$D^*(v, Q) = \frac{v^2}{8\beta(1-\beta)} - O(v^3),$$

where β is a measure of the “imbalance” of Q defined below; this may also be interpreted as a reverse Pinsker inequality. The *range* of a distribution is

$$\mathcal{R}(Q) = \{Q(A) : A \in \mathcal{F}\}.$$

A distribution Q has *full range* if $\mathcal{R}(Q) = [0, 1]$. Non-atomic distributions on \mathbb{R} have full range. The *balance coefficient* of a distribution Q is

$$\beta = \inf\{x \in \mathcal{R}(Q) : x \geq 1/2\}.$$

A distribution is *balanced* if $\beta = 1/2$ and *unbalanced* otherwise. In particular, all distributions with full range are balanced. Note that the balance coefficient of a discrete distribution Q is bounded by¹

$$\beta \leq \frac{1}{2} + \frac{q_{\max}}{2}, \quad (3)$$

where $q_{\max} = \max_{\omega \in \Omega} q(\omega)$. Ordentlich and Weinberger [24] considered the following distribution-dependent refinement of Pinsker’s inequality. For a distribution Q with balance coefficient β , define $\varphi(Q)$ by

$$\varphi(Q) = \frac{1}{2\beta - 1} \ln \frac{\beta}{1 - \beta}$$

(for $\beta = 1/2$, $\varphi(Q) = 2$). It is shown in [24] that

$$D(P||Q) \geq \frac{\varphi(Q)}{4} V(P, Q)^2 \quad (4)$$

for all P, Q , and furthermore, that $\varphi(Q)/4$ is the best Q -dependent coefficient possible:

$$\inf_P \frac{D(P||Q)}{V(P, Q)^2} = \frac{\varphi(Q)}{4}. \quad (5)$$

Although the left-hand sides of (2) and (5) bear a superficial resemblance, the two quantities are quite different (in particular, the former is constrained by $V(P, Q) \geq v$). While distribution-independent versions of (4) exist (viz., (1)), our main result (Theorem 1) does not admit a distribution-independent form. Simply put, the result in [24] yields a lower bound on $D^*(v, Q)$, while we seek to upper-bound this quantity — and actually compute it exactly for unbalanced distributions.

II. MAIN RESULTS

We can now state our reverse Pinsker inequality:

Theorem 1. *Suppose $Q \in \mathcal{P}$ has balance coefficient β . Then:*

(a) *For $\beta \geq 1/2$ and $0 < v < 1$,*

$$L(v) \leq D^*(v, Q) \leq \text{KL}_2(\beta - v/2, \beta).$$

(b) *For $\beta > 1/2$ and $0 < v < 4(\beta - 1/2)$,*

$$D^*(v, Q) = \text{KL}_2(\beta - v/2, \beta).$$

As a comparison of orders of magnitude, note that

$$\begin{aligned} \text{KL}_2(\beta - v/2, \beta) &= \frac{v^2}{8\beta(1-\beta)} - \frac{(2\beta-1)v^3}{48\beta^2(1-\beta)^2} + O(v^4), \\ \text{KL}_2(1/2 - v/2, 1/2) &= \frac{v^2}{2} + \frac{v^4}{12} + O(v^6), \\ L(v) &= \frac{v^2}{2} + \frac{v^4}{36} + O(v^6), \end{aligned}$$

¹Since we will not use this fact in the sequel, we only give a proof sketch. The case where $q_{\max} \geq 1/2$ is trivial, so assume $q_{\max} < 1/2$. Consider the following greedy algorithm: initialize A to be the empty set and repeatedly include the heaviest available atom such that A ’s total mass remains under $1/2$ (once an atom has been added to A , it is no longer “available”). If ω is the first atom whose inclusion will bring A ’s mass over $1/2$, either $A \cup \{\omega\}$ or $\Omega \setminus A$ establishes the bound in (3).

where the first two expansions are straightforward and the last one is well-known [14]. Combining the bound of Ordentlich and Weinberger (4) with Theorem 1, we get

$$\begin{aligned} \frac{1}{4(2\beta-1)} \ln \frac{\beta}{1-\beta} v^2 &\leq D^*(v, Q) \\ &\leq \text{KL}_2(\beta - v/2, \beta) \\ &= \frac{v^2}{8\beta(1-\beta)} - \mathcal{O}(v^3). \end{aligned}$$

As a consistency check, one may verify that

$$\frac{1}{4(2\beta-1)} \ln \frac{\beta}{1-\beta} \leq \frac{1}{8\beta(1-\beta)}$$

for $1/2 \leq \beta < 1$.

Theorem 2. *If $Q \in \mathcal{P}$ has full range, then*

$$D^*(v, Q) = L(v), \quad 0 < v < 2.$$

III. PROOFS

We will repeatedly invoke the standard fact that $D(\cdot|\cdot)$ is convex in both arguments [6], [29]. Our first lemma provides a structural result for extremal distributions. Suppose a distribution $Q \in \mathcal{P}$ is given, along with an $A \in \mathcal{F}$ and a $0 < v \leq 2(1 - Q(A))$. Denote by $\mathcal{P}(Q, A, v)$ the set of all distributions $P \in \mathcal{P}$ for which $V(P, Q) = v$ and $A = \{\omega \in \Omega : q(\omega) < p(\omega)\}$. The above “restriction” on the range of v derives from the fact that every $P \in \mathcal{P}(Q, A, v)$ must satisfy $V(P, Q) \leq 2(1 - Q(A))$.

Lemma 3. *For all $Q \in \mathcal{P}$, $A \in \mathcal{F}$ with $0 < Q(A) < 1$, and $v \in (0, 2(1 - Q(A))]$, let $P^* \in \mathcal{P}$ be the measure with density*

$$p^* = (a\mathbb{1}_A + b\mathbb{1}_{\Omega \setminus A})q,$$

where

$$a = 1 + \frac{v}{2Q(A)}, \quad b = 1 - \frac{v}{2(1 - Q(A))}.$$

Then P^* belongs to $\mathcal{P}(Q, A, v)$, and P^* is the unique minimizer of $D(P|Q)$ over $P \in \mathcal{P}(Q, A, v)$.

Proof: Obviously, P^* belongs to $\mathcal{P}(Q, A, v)$. We claim that

$$D(P|Q) = D(P|P^*) + D(P^*|Q) \tag{6}$$

holds for all $P \in \mathcal{P}(Q, A, v)$, whence the lemma follows. Indeed, putting $B = \Omega \setminus A$ and using the fact that

$$\begin{aligned} D(P|P^*) &= D(P|Q) - P(A) \ln a - P(B) \ln b, \\ D(P^*|Q) &= Q(A) a \ln a + Q(B) b \ln b, \end{aligned}$$

we see that (6) is equivalent to the identity

$$(Q(A) - P(A) + v/2) \ln a + (Q(B) - P(B) - v/2) \ln b = 0,$$

which follows immediately from the elementary fact that

$$P(A) - Q(A) = Q(B) - P(B) = v/2.$$

■

Our next result is that D^* actually has a somewhat simpler form than the original definition (2).

Lemma 4. *For all distributions Q and all $v > 0$,*

$$D^*(v, Q) = \inf_{P: V(P, Q) = v} D(P|Q).$$

Proof: For any $\varepsilon > 0$, let $P_\varepsilon \in \mathcal{P}$ be such that $V(P_\varepsilon, Q) \geq v$ and

$$D(P_\varepsilon|Q) < D^*(v, Q) + \varepsilon \tag{7}$$

and define, for $0 \leq \delta \leq 1$,

$$P_{\varepsilon, \delta} = \delta P_\varepsilon + (1 - \delta)Q.$$

Since $V(P_{\varepsilon, \bar{\delta}}, Q) = \delta V(P_\varepsilon, Q)$, we may always choose $\bar{\delta} = \bar{\delta}(P_\varepsilon)$ so that $V(P_{\varepsilon, \bar{\delta}}, Q) = v$. By convexity of $D(\cdot \| \cdot)$, we have

$$\begin{aligned} D(P_{\varepsilon, \bar{\delta}} \| Q) &\leq \bar{\delta} D(P_\varepsilon \| Q) + (1 - \bar{\delta}) D(Q \| Q) \\ &\leq D(P_\varepsilon \| Q) \\ &< D^*(v, Q) + \varepsilon, \end{aligned}$$

and hence

$$\inf_{P: V(P, Q) \geq v} D(P \| Q) = \inf_{P: V(P, Q) = v} D(P \| Q).$$

Proof of Theorem 2: Below we take the infimum over \mathcal{P} in two steps: first over $\mathcal{P}(Q, A, v)$ and then over $A \in \mathcal{F}$ satisfying $Q(A) \leq 1 - v/2$. It follows from Lemmas 3 and 4 that

$$\begin{aligned} D^*(v, Q) &= \inf_{P \in \mathcal{P}: V(P, Q) = v} D(P \| Q) \\ &= \inf_A \inf_{P \in \mathcal{P}(Q, A, v)} D(P \| Q) \\ &= \inf_A D(P^* \| Q) \\ &= \inf_A [Q(A)a \ln a + Q(\Omega \setminus A)b \ln b] \\ &= \inf_A \text{KL}_2(Q(A) + v/2, Q(A)), \end{aligned} \tag{8}$$

where P^* , a and b are as defined in Lemma 3. Using the fact [14], [16] that

$$L(v) = \inf_{0 < x < 1 - v/2} \text{KL}_2(x + v/2, x),$$

we have that $D^*(v, Q) = L(v)$ for Q with full range, which proves the claim. \blacksquare

For the proof of Theorem 1, we will need additional lemmata, the first of which will allow us to restrict our attention to distributions with binary support. Now it is well known [14] that for each pair of distributions P, Q , there is a pair of binary distributions P', Q' such that $V(P', Q') = V(P, Q)$ and $D(P' \| Q') = D(P \| Q)$ (this fact is generalized to general f -divergences in [16]). However, in our case Q is fixed whereas only P is allowed to vary, and so this result is not directly applicable. Still, an analogue of this phenomenon also holds in our case. We will consistently use π to denote a map $\Omega \rightarrow \{1, 2\}$ and for $Q \in \mathcal{P}$, the notation $\pi(Q)$ refers to the distribution $(Q(\pi^{-1}(1)), Q(\pi^{-1}(2)))$ on $\{1, 2\}$. For measurable $A \subseteq \Omega$, the map $\pi_A : \Omega \rightarrow \{1, 2\}$ is defined by $\pi_A^{-1}(1) = A = \Omega \setminus \pi_A^{-1}(2)$.

Lemma 5. *Let $Q \in \mathcal{P}$ be a distribution whose support contains at least two points. Then*

- (i) *For any measurable map $\pi : \Omega \rightarrow \{1, 2\}$ and any distribution $P' = (p'_1, p'_2)$ on $\{1, 2\}$, there exists a $P \in \mathcal{P}$ such that $V(P, Q) = V(P', \pi(Q))$ and $D(P \| Q) = D(P' \| \pi(Q))$. In particular,*

$$D^*(v, \pi(Q)) \geq D^*(v, Q), \quad v > 0.$$

- (ii) *For all $v > 0$, there is a measurable $\pi : \Omega \rightarrow \{1, 2\}$ such that $D^*(v, Q) = D^*(v, \pi(Q))$.*

Proof: Let $P' = (p'_1, p'_2)$ be a distribution on $\{1, 2\}$ and define the distribution $P \in \mathcal{P}$ as the mixture

$$P = p'_1 Q(\cdot | \pi^{-1}(1)) + p'_2 Q(\cdot | \pi^{-1}(2)).$$

Then

$$\begin{aligned} V(P, Q) &= \int_{\Omega} |p(\omega) - q(\omega)| \, d\mu(\omega) \\ &= \int_{\pi^{-1}(1)} \left| p'_1 \frac{q(\omega)}{Q(\pi^{-1}(1))} - q(\omega) \right| \, d\mu(\omega) + \int_{\pi^{-1}(2)} \left| p'_2 \frac{q(\omega)}{Q(\pi^{-1}(2))} - q(\omega) \right| \, d\mu(\omega) \\ &= Q(\pi^{-1}(1)) \left| \frac{p'_1}{Q(\pi^{-1}(1))} - 1 \right| + Q(\pi^{-1}(2)) \left| \frac{p'_2}{Q(\pi^{-1}(2))} - 1 \right| \\ &= |p'_1 - Q(\pi^{-1}(1))| + |p'_2 - Q(\pi^{-1}(2))| \\ &= V(P', \pi(Q)) \end{aligned}$$

and

$$\begin{aligned}
D(P\|Q) &= \int_{\Omega} p(\omega) \log \frac{p(\omega)}{q(\omega)} d\mu(\omega) \\
&= \int_{\pi^{-1}(1)} p'_1 \frac{q(\omega)}{Q(\pi^{-1}(1))} \log \frac{p'_1 q(\omega)/Q(\pi^{-1}(1))}{q(\omega)} d\mu(\omega) + \int_{\pi^{-1}(2)} p'_2 \frac{q(\omega)}{Q(\pi^{-1}(2))} \log \frac{p'_2 q(\omega)/Q(\pi^{-1}(2))}{q(\omega)} d\mu(\omega) \\
&= p'_1 \log \frac{P'_1}{Q(\pi^{-1}(1))} + p'_2 \log \frac{P'_2}{Q(\pi^{-1}(2))} \\
&= D(P'\|\pi(Q)).
\end{aligned}$$

Hence,

$$\begin{aligned}
D^*(v, \pi(Q)) &= \inf_{P': V(P', \pi(Q))=v} D(P'\|\pi(Q)) \\
&= \inf_{P=P'_1 Q(\cdot|\pi^{-1}(1)) + P'_2 Q(\cdot|\pi^{-1}(2)): V(P', \pi(Q))=v} D(P\|Q) \\
&\geq \inf_{P: V(P, Q)=v} D(P\|Q) \\
&= D^*(v, Q),
\end{aligned}$$

where the first and last identities follow from Lemma 4. This proves (i). For any $\varepsilon > 0$, the proof of Lemma 4 furnishes a $P_\varepsilon \in \mathcal{P}$ such that $V(P_\varepsilon, Q) = v$ and $D(P_\varepsilon\|Q) < D^*(v, Q) + \varepsilon$. Define π by

$$\pi(\omega) = \begin{cases} 1, & p_\varepsilon(\omega) < q(\omega), \\ 2, & \text{else.} \end{cases}$$

Then

$$\begin{aligned}
v = V(P_\varepsilon, Q) &= \int_{p_\varepsilon < q} |p_\varepsilon(\omega) - q(\omega)| d\mu(\omega) + \int_{p_\varepsilon \geq q} |p_\varepsilon(\omega) - q(\omega)| d\mu(\omega) \\
&= V(\pi(P_\varepsilon), \pi(Q))
\end{aligned}$$

and

$$D(\pi(P_\varepsilon)\|\pi(Q)) \leq D(P_\varepsilon\|Q) < D^*(v, Q) + \varepsilon,$$

where the first inequality follows from the data processing inequality [29, Theorem 9]. Since $\varepsilon > 0$ is arbitrary, we have that $D^*(v, \pi(Q)) \leq D^*(v, Q)$. Taking $P' = \pi(P_\varepsilon)$, it follows from (i) that $D(P'\|\pi(Q)) = D(P_\varepsilon\|Q)$, which proves (ii). ■

Next, we characterize the extremal P^* satisfying $D(P^*\|Q) = D^*(v, Q)$ in the binary case.

Lemma 6. *Let $Q = (q_0, 1 - q_0)$ be a binary distribution with $q_0 > 1/2$ and $v \in (0, 2q_0]$. Then the unique P^* satisfying $V(P^*, Q) = v$ and $D(P^*\|Q) = D^*(v, Q)$ is*

$$P^* = \left(q_0 - \frac{v}{2}, 1 - q_0 + \frac{v}{2} \right).$$

Proof: By Lemma 4, there are at most two possibilities for P^* , namely,

$$P^* = P_1 = \left(q_0 - \frac{v}{2}, 1 - q_0 + \frac{v}{2} \right)$$

and

$$P^* = P_2 = \left(q_0 + \frac{v}{2}, 1 - q_0 - \frac{v}{2} \right).$$

(Actually, if $v > 2(1 - q_0)$ then only P_1 is a valid distribution.) A second-order Taylor expansion yields

$$\text{KL}_2(q_0 + x, q_0) = \frac{1}{2} \frac{x^2}{(q_0 + \theta)(1 - q_0 - \theta)}$$

for some θ between 0 and x . Hence

$$\text{KL}_2\left(q_0 - \frac{v}{2}, q_0\right) < \frac{1}{2} \frac{(v/2)^2}{q_0(1 - q_0)} < \text{KL}_2\left(q_0 + \frac{v}{2}, q_0\right)$$

for all $v \in (0, 2(1 - q_0)]$, which implies that $P^* = P_1$. ■

Proof of Theorem 1 (a): The first inequality is an immediate consequence of Lemma 4. To prove the second one, let $Q \in \mathcal{P}$ be a distribution with balance coefficient β , and $0 < v < 1$. By definition of β , for all $\varepsilon > 0$ there is a measurable $A_\varepsilon \subseteq \Omega$ such that $\beta \leq Q(A_\varepsilon) \leq \beta + \varepsilon$. Then Lemma 5(i) implies that

$$D^*(v, Q) \leq D^*(v, \pi_{A_\varepsilon}(Q))$$

and by taking ε arbitrarily small,

$$D^*(v, Q) \leq D^*(v, Q'),$$

where $Q' = (\beta, 1 - \beta)$. Finally, Lemma 6 implies that $D^*(v, Q') = \text{KL}_2(\beta - v/2, \beta)$. ■

Lemma 7. For every fixed $0 < \delta < 1/2$, the binary divergence $\text{KL}_2(x - \delta, x)$ is strictly increasing in x on $[1/2 + \delta/2, 1]$.

Proof: Define the function

$$F(x) = \text{KL}_2(x - \delta, x).$$

Since KL-divergence is jointly convex in the distributions, F is a convex function. Thus, it is sufficient to prove that $F'(x)$ is positive for $x = 1/2 + \delta/2$. We have

$$F'(x) = \frac{\delta + (1-x)x \ln(1 - \frac{\delta}{x}) + (1-x)x \ln \frac{1-x}{1-x+\delta}}{(1-x)x}$$

and

$$F'(1/2 + \delta/2) = \frac{4\delta}{1 - \delta^2} + 2 \log \frac{1 - \delta}{1 + \delta} =: G(\delta).$$

Now $G(0) = 0$ and

$$G'(\delta) = 8 \left(\frac{\delta}{1 - \delta^2} \right)^2 > 0,$$

which proves the lemma. ■

Proof of Theorem 1 (b): Consider a $Q \in \mathcal{P}$ with balance coefficient $\beta > 1/2$ and $0 < v < 4(\beta - 1/2)$. Then Lemma 5 implies that

$$\begin{aligned} D^*(v, Q) &= \inf_{A \in \mathcal{F}} D^*(v, \pi_A(Q)) \\ &= \inf_{A \in \mathcal{F}: Q(A) > 1/2} D^*(v, (Q(A), (1 - Q(A)))) \end{aligned}$$

where the second identity holds because $D^*(v, (q_0, q_1)) = D^*(v, (q_1, q_0))$. Invoking Lemma 6, we have that for $Q(A) > 1/2$,

$$D^*(v, \pi_A(Q)) = \text{KL}_2\left(Q(A) - \frac{v}{2}, Q(A)\right)$$

and hence

$$D^*(v, Q) = \inf_{A: Q(A) > 1/2} \text{KL}_2\left(Q(A) - \frac{v}{2}, Q(A)\right).$$

Since $1/2 + v/4 \leq \beta \leq Q(A)$, we may invoke Lemma 7 with $x = Q(A)$ and $\delta = v/2$ to conclude that $D^*(v, Q) = \text{KL}_2(\beta - v/2, \beta)$. ■

IV. APPLICATION: CONVERGENCE OF THE EMPIRICAL DISTRIBUTION

The results in [4] have bearing on the convergence of the empirical distribution to the true one in the total variation norm. More precisely, the paper considers a sequence of i.i.d. \mathbb{N} -valued random variables X_1, X_2, \dots , distributed according to $Q = (q_1, q_2, \dots)$ and denotes

$$J_n = V(Q, \hat{Q}_n), \quad n \in \mathbb{N},$$

where \hat{Q}_n is the empirical distribution induced by the first n observations. Let us recall Sanov's Theorem [6], [9], which yields

$$-\lim_{n \rightarrow \infty} \frac{1}{n} \ln Q(J_n - \mathbb{E}J_n > \varepsilon) = D^*(\varepsilon, Q).$$

Since the map $(X_1, \dots, X_n) \mapsto J_n$ is $2/n$ -Lipschitz continuous with respect to the Hamming distance, McDiarmid's inequality [22] implies

$$Q(|J_n - \mathbb{E}J_n| > \varepsilon) \leq 2 \exp\left(-\frac{n\varepsilon^2}{2}\right), \quad n \in \mathbb{N}, \varepsilon > 0. \quad (9)$$

Being a rather general-purpose tool, in many cases McDiarmid's bound does not yield optimal estimates. Since for balanced distributions $D^*(\varepsilon, Q) \leq \varepsilon^2/2 + O(\varepsilon^4)$, we see that the estimate in (9) actually has the optimal constant $1/2$ in the exponent. (See [3, Theorem 1] for other instances where the quantity $\varepsilon^2/2$ emerges in the exponent.) We also see that the McDiarmid's bound must be suboptimal for unbalanced distributions. The exponential decrease of $Q(|J_n - \mathbb{E}J_n| > \varepsilon)$ implies that $J_n - \mathbb{E}J_n$

tends to zero almost surely. We should note that $\mathbb{E}J_n$ will tend to zero but the rate of convergence may be arbitrarily slow. In [4] it was shown that

$$\mathbb{E}J_n \leq n^{-1/2} \sum_{j \in \mathbb{N}} q_j^{1/2}$$

and that for Q with finite support of size k ,

$$\mathbb{E}J_n \leq \left(\frac{k}{n}\right)^{1/2}.$$

In greater generality, it was shown that

$$\frac{1}{4}(\Lambda_n - n^{-1/2}) \leq \mathbb{E}J_n \leq \Lambda_n, \quad n \geq 2,$$

where

$$\Lambda_n(Q) = n^{-1/2} \sum_{q_j \geq 1/n} q_j^{1/2} + 2 \sum_{q_j < 1/n} q_j$$

tends to zero for n tending to infinity, although the rate at which $\Lambda_n(Q)$ decays may be arbitrarily slow, depending on Q .

ACKNOWLEDGEMENTS

We thank László Györfi and Robert Williamson for helpful correspondence, and in particular for bringing [3] to our attention. We are grateful to Sergio Verdú for a careful reading of the paper and useful suggestions. The comments of the anonymous referees have greatly contributed to the quality of this paper. In particular the proof of Theorem 2 has been considerably simplified due to their comments.

REFERENCES

- [1] Alexander M. Barg, Leonid A. Bassalygo, Vladimir M. Blinovskii, et al. In memory of Mark Semënovich Pinsker. *Problemy Peredachi Informatsii*, 40(1):3–5, 2004.
- [2] Andrew R. Barron. Entropy and the central limit theorem. *Ann. Probab.*, 14(1):336–342, 1986.
- [3] Jan Beirlant, Luc Devroye, László Györfi, and Igor Vajda. Large deviations of divergence measures on partitions. *J. Statist. Plann. Inference*, 93(1-2):1–16, 2001.
- [4] Daniel Berend and Aryeh Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics & Probability Letters*, 83(4):1254–1259, 2013.
- [5] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, Cambridge, 2006.
- [6] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, Hoboken, NJ, second edition, 2006.
- [7] Imre Csiszár. A note on Jensen’s inequality. *Studia Sci. Math. Hungar.*, 1:185–188, 1966.
- [8] Imre Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.
- [9] Frank den Hollander. *Large deviations*, volume 14 of *Fields Institute Monographs*. American Mathematical Society, Providence, RI, 2000.
- [10] Sever S. Dragomir. Upper bounds for the Kullback-Leibler distance and applications. *Bull. Math. Soc. Sci. Math. Roumanie (N.S.)*, 43(91)(1):25–37, 2000.
- [11] Sever S. Dragomir and Vido Gluščević. New estimates of the Kullback-Leibler distance and applications. In *Inequality theory and applications. Vol. I*, pages 123–137. Nova Sci. Publ., Huntington, NY, 2001.
- [12] Sever S. Dragomir and Vido Gluščević. Some inequalities for the Kullback-Leibler and χ^2 -distances in information theory and applications. *Tamsui Oxf. J. Math. Sci.*, 17(2):97–111, 2001.
- [13] Eyal Even-Dar, Sham M. Kakade, and Yishay Mansour. The value of observation for monitoring dynamic systems. In *IJCAI*, pages 2474–2479, 2007.
- [14] Alexei A. Fedotov, Peter Harremoës, and Flemming Topsøe. Refinements of Pinsker’s inequality. *IEEE Trans. Inform. Theory*, 49(6):1491–1498, 2003.
- [15] Gustavo L. Gilardoni. On Pinsker’s and Vajda’s type inequalities for Csiszár’s f -divergences. *IEEE Trans. Inform. Theory*, 56(11):5377–5386, 2010.
- [16] Peter Harremoës and Igor Vajda. On pairs of f -divergences and their joint range. *IEEE Trans. Inform. Theory*, 57(6):3230–3235, 2011.
- [17] Johannes H. B. Kemperman. On the optimum rate of transmitting information. In *Probability and Information Theory (Proc. Internat. Sympos., McMaster Univ., Hamilton, Ont., 1968)*, pages 126–169. Springer, Berlin, 1969.
- [18] Johannes H. B. Kemperman. On the optimum rate of transmitting information. *Ann. Math. Statist.*, 40:2156–2177, 1969.
- [19] Olaf Krafft. A note on exponential bounds for binomial probabilities. *Ann. Inst. Stat. Math.*, 21:219–220, 1969.
- [20] Solomon Kullback. A lower bound for discrimination information in terms of variation. *IEEE Trans. Inform. Theory*, 13:126–127, 1967. Correction, volume 16, p. 652, 1970.
- [21] Katalin Marton. Bounding \bar{d} -distance by informational divergence: a method to prove measure concentration. *Ann. Probab.*, 24(2):857–866, 1996.
- [22] Colin McDiarmid. On the method of bounded differences. In J. Siemons, editor, *Surveys in Combinatorics, volume 141 of LMS Lecture Notes Series*, pages 148–188. Morgan Kaufmann Publishers, San Mateo, CA, 1989.
- [23] Henry P. McKean, Jr. Speed of approach to equilibrium for Kac’s caricature of a Maxwellian gas. *Arch. Rational Mech. Anal.*, 21:343–367, 1966.
- [24] Erik Ordentlich and Marcelo J. Weinberger. A distribution dependent refinement of pinsker’s inequality. *IEEE Transactions on Information Theory*, 51(5):1836–1840, 2005.
- [25] Mark D. Reid and Robert C. Williamson. Generalised Pinsker inequalities. In *COLT*, 2009.
- [26] Flemming Topsøe. Bounds for entropy and divergence for distributions over a two-element set. *JIPAM. J. Inequal. Pure Appl. Math.*, 2(2):Article 25, 13 pp. (electronic), 2001.
- [27] Godfried T. Toussaint. Probability of error, expected divergence, and the affinity of several distributions. *IEEE Trans. Systems Man Cybernet.*, SMC-8(6):482–485, 1978.
- [28] Igor Vajda. Note on discrimination information and variation. *IEEE Trans. Information Theory*, IT-16:771–773, 1970.
- [29] Tim van Erven and Peter Harremoës. Rényi divergence and Kullback-Leibler divergence. *IEEE Trans. Inform. Theory*. Accepted for publication.