
From Stochastic Mixability to Fast Rates

Nishant A. Mehta

Research School of Computer Science
Australian National University
nishant.mehta@anu.edu.au

Robert C. Williamson

Research School of Computer Science
Australian National University and NICTA
bob.williamson@anu.edu.au

Abstract

Empirical risk minimization (ERM) is a fundamental learning rule for statistical learning problems where the data is generated according to some unknown distribution P and returns a hypothesis f chosen from a fixed class \mathcal{F} with small loss ℓ . In the parametric setting, depending upon (ℓ, \mathcal{F}, P) ERM can have slow $(1/\sqrt{n})$ or fast $(1/n)$ rates of convergence of the excess risk as a function of the sample size n . There exist several results that give sufficient conditions for fast rates in terms of joint properties of ℓ , \mathcal{F} , and P , such as the margin condition and the Bernstein condition. In the non-statistical prediction with expert advice setting, there is an analogous slow and fast rate phenomenon, and it is entirely characterized in terms of the mixability of the loss ℓ (there being no role there for \mathcal{F} or P). The notion of stochastic mixability builds a bridge between these two models of learning, reducing to classical mixability in a special case. The present paper presents a direct proof of fast rates for ERM in terms of stochastic mixability of (ℓ, \mathcal{F}, P) , and in so doing provides new insight into the fast-rates phenomenon. The proof exploits an old result of Kemperman on the solution to the general moment problem. We also show a partial converse that suggests a characterization of fast rates for ERM in terms of stochastic mixability is possible.

1 Introduction

Recent years have unveiled central contact points between the areas of statistical and online learning. These include Abernethy et al.'s [1] unified Bregman-divergence based analysis of online convex optimization and statistical learning, the online-to-batch conversion of the exponentially weighted average forecaster (a special case of the aggregating algorithm for mixable losses) which yields the progressive mixture rule as can be seen e.g. from the work of Audibert [2], and most recently Van Erven et al.'s [21] injection of the concept of mixability into the statistical learning space in the form of stochastic mixability. It is this last connection that will be our departure point for this work.

Mixability is a fundamental property of a loss that characterizes when constant regret is possible in the online learning game of prediction with expert advice [23]. Stochastic mixability is a natural adaptation of mixability to the statistical learning setting; in fact, in the special case where the function class consists of all possible functions from the input space to the prediction space, stochastic mixability is equivalent to mixability [21]. Just as Vovk and coworkers (see e.g. [24, 8]) have developed a rich convex geometric understanding of mixability, stochastic mixability can be understood as a sort of effective convexity.

In this work, we study the $O(1/n)$ -fast rate phenomenon in statistical learning from the perspective of stochastic mixability. Our motivation is that stochastic mixability might characterize fast rates in statistical learning. As a first step, [Theorem 5](#) herein establishes via a rather direct argument that stochastic mixability implies an exact oracle inequality (i.e. with leading constant 1) with a fast rate for finite function classes, and [Theorem 7](#) extends this result to VC-type classes. This result can be understood as a new chapter in an evolving narrative that started with Lee et al.'s [13] seminal paper

showing fast rates for agnostic learning with squared loss over convex function classes, and that was continued by Mendelson [18] who showed that fast rates are possible for p -losses $(y, \hat{y}) \mapsto |y - \hat{y}|^p$ over effectively convex function classes by passing through a Bernstein condition (defined in (12)).

We also show that when stochastic mixability does not hold in a certain sense (described in Section 5), then the risk minimizer is not unique in a bad way. This is precisely the situation at the heart of the works of Mendelson [18] and Mendelson and Williamson [19], which show that having non-unique minimizers is symptomatic of bad geometry of the learning problem. In such situations, there are certain targets (i.e. output conditional distributions) close to the original target under which empirical risk minimization learns (ERM) at a slow rate, where the guilty target depends on the sample size and the target sequence approaches the original target asymptotically. Even the best known upper bounds have constants that blow up in the case of non-unique minimizers. Thus, whereas stochastic mixability implies fast rates, a sort of converse is also true, where learning is hard in a “neighborhood” of statistical learning problems for which stochastic mixability does not hold. In addition, since a stochastically mixable problem’s function class looks convex from the perspective of risk minimization, and since when stochastic mixability fails the function class looks non-convex from the same perspective (it has multiple well-separated minimizers), stochastic mixability *characterizes* the effective convexity of the learning problem from the perspective of risk minimization.

Much of the recent work in obtaining faster learning rates in agnostic learning has taken place in settings where a Bernstein condition holds, including results based on local Rademacher complexities [3, 10]. The Bernstein condition appears to have first been used by Bartlett and Mendelson [4] in their analysis of ERM; this condition is subtly different from the margin condition of Mammen and Tsybakov [15, 20], which has been used to obtain fast rates for classification. Lecué [12] pinpoints that the difference between the two conditions is that the margin condition applies to the excess loss relative to the best predictor (not necessarily in the model class) whereas the Bernstein condition applies to the excess loss relative to the best predictor in the model class. Our approach in this work is complementary to the approaches of previous works, coming from a different assumption that forms a bridge to the online learning setting. Yet this assumption is related; the Bernstein condition implies stochastic mixability under a bounded losses assumption [21]. Further understanding the connection between the Bernstein condition and stochastic mixability is an ongoing effort.

Contributions. The core contribution of this work is to show a new path to the $\tilde{O}(1/n)$ -fast rate in statistical learning. We are not aware of previous results that show fast rates from the stochastic mixability assumption. Secondly, we establish intermediate learning rates that interpolate between the fast and slow rate under a weaker notion of stochastic mixability. Finally, we show that in a certain sense stochastic mixability characterizes the effective convexity of the statistical problem.

In the next section we formally define the statistical problem, review stochastic mixability, and explain our high-level approach toward getting fast rates. This approach involves directly appealing to the Cramér-Chernoff method, from which nearly all known concentration inequalities arose in one way or another. In Section 3, we frame the problem of computing a particular moment of a certain excess loss random variable as a general moment problem. We sufficiently bound the optimal value of the moment, which allows for a direct application of the Cramér-Chernoff method. These results easily imply a fast rates bound for finite classes that can be extended to parametric (VC-type) classes, as shown in Section 4. We describe in Section 5 how stochastic mixability characterizes a certain notion of convexity of the statistical learning problem. In Section 6, we extend the fast rates results to classes that obey a notion we call weak stochastic mixability. Finally, Section 7 concludes this work with connections to related topics in statistical learning theory and a discussion of open problems.

2 Stochastic mixability, Cramér-Chernoff, and ERM

Let $(\ell, \mathcal{F}, \mathbb{P})$ be a statistical learning problem with $\ell : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}_+$ a nonnegative loss, $\mathcal{F} \subset \mathbb{R}^{\mathcal{X}}$ a compact function class, and \mathbb{P} a probability measure over $\mathcal{X} \times \mathcal{Y}$ for input space \mathcal{X} and output/target space \mathcal{Y} . Let Z be a random variable defined as $Z = (X, Y) \sim \mathbb{P}$. We assume for all $f \in \mathcal{F}$, $\ell(Y, f(X)) \leq V$ almost surely (a.s.) for some constant V .

A probability measure \mathbb{P} operates on functions and loss-composed functions as:

$$\mathbb{P} f = \mathbb{E}_{(X, Y) \sim \mathbb{P}} f(X) \qquad \mathbb{P} \ell(\cdot, f) = \mathbb{E}_{(X, Y) \sim \mathbb{P}} \ell(Y, f(X)).$$

Similarly, an empirical measure P_n associated with an n -sample \mathbf{z} , comprising n iid samples $(x_1, y_1), \dots, (x_n, y_n)$, operates on functions and loss-composed functions as:

$$P_n f = \frac{1}{n} \sum_{j=1}^n f(x_j) \quad P_n \ell(\cdot, f) = \frac{1}{n} \sum_{j=1}^n \ell(y_j, f(x_j)).$$

Let f^* be any function for which $P \ell(\cdot, f^*) = \inf_{f \in \mathcal{F}} P \ell(\cdot, f)$. For each $f \in \mathcal{F}$ define the excess risk random variable $Z_f := \ell(Y, f(X)) - \ell(Y, f^*(X))$.

We frequently work with the following two subclasses. For any $\varepsilon > 0$, define the subclasses

$$\mathcal{F}_{\leq \varepsilon} := \{f \in \mathcal{F} : P Z_f \leq \varepsilon\} \quad \mathcal{F}_{\geq \varepsilon} := \{f \in \mathcal{F} : P Z_f \geq \varepsilon\}.$$

2.1 Stochastic mixability

For $\eta > 0$, we say that (ℓ, \mathcal{F}, P) is η -stochastically mixable if for all $f \in \mathcal{F}$

$$\log E \exp(-\eta Z_f) \leq 0. \quad (1)$$

If η -stochastic mixability holds for some $\eta > 0$, then we say that (ℓ, \mathcal{F}, P) is *stochastically mixable*. Throughout this paper it is assumed that the stochastic mixability condition holds, and we take η^* to be the largest η such that η -stochastic mixability holds. Condition (1) has a rich history, beginning from the foundational thesis of Li [14] who studied the special case of $\eta^* = 1$ in density estimation with log loss from the perspective of information geometry. The connections that Li showed between this condition and convexity were strengthened by Grünwald [6, 7] and Van Erven et al. [21].

2.2 Cramér-Chernoff

The high-level strategy taken here is to show that with high probability ERM will not select a fixed hypothesis function f with excess risk above $\frac{a}{n}$ for some constant $a > 0$. For each hypothesis, this guarantee will flow from the Cramér-Chernoff method [5] by controlling the cumulant generating function (CGF) of $-Z_f$ in a particular way to yield exponential concentration. This control will be possible because the η^* -stochastic mixability condition implies that the CGF of $-Z_f$ takes the value 0 at some $\eta \geq \eta^*$, a fact later exploited by our key tool [Theorem 3](#).

Let Z be a real-valued random variable. Applying Markov's inequality to an exponentially transformed random variable yields that, for any $\eta \geq 0$ and $t \in \mathbb{R}$

$$\Pr(Z \geq t) \leq \exp(-\eta t + \log E \exp(\eta Z)); \quad (2)$$

the inequality is non-trivial only if $t > E Z$ and $\eta > 0$.

2.3 Analysis of ERM

We consider the *ERM estimator* $\hat{f}_{\mathbf{z}} := \arg \min_{f \in \mathcal{F}} P_n \ell(\cdot, f)$. That is, given an n -sample \mathbf{z} , ERM selects any $\hat{f}_{\mathbf{z}} \in \mathcal{F}$ minimizing the empirical risk $P_n \ell(\cdot, f)$. We say ERM is ε -good when $\hat{f}_{\mathbf{z}} \in \mathcal{F}_{\leq \varepsilon}$. In order to show that ERM is ε -good it is sufficient to show that for all $f \in \mathcal{F} \setminus \mathcal{F}_{\leq \varepsilon}$ we have $P Z_f > 0$. The goal is to show that with high probability ERM is ε -good, and we will do this by showing that with high probability uniformly for all $f \in \mathcal{F} \setminus \mathcal{F}_{\leq \varepsilon}$ we have $P_n Z_f > t$ for some slack $t > 0$ that will come in handy later.

For a real-valued random variable X , recall that the *cumulant generating function* of X is $\eta \mapsto \Lambda_X(\eta) := \log E e^{\eta X}$; we allow $\Lambda_X(\eta)$ to be infinite for some $\eta > 0$.

Theorem 1 (Cramér-Chernoff Control on ERM). *Let $a > 0$ and select f such that $E Z_f > 0$. Let $t < E Z_f$. If there exists $\eta > 0$ such that $\Lambda_{-Z_f}(\eta) \leq -\frac{a}{n}$, then*

$$\Pr\left\{P_n \ell(\cdot, f) \leq P_n \ell(\cdot, f^*) + t\right\} \leq \exp(-a + \eta t).$$

Proof. Let $Z_{f,1}, \dots, Z_{f,n}$ be iid copies of Z_f , and define the sum $S_{f,n} := \sum_{j=1}^n -Z_{f,j}$. Since $(-t) > E \frac{1}{n} S_{f,n}$, then from (2) we have

$$\begin{aligned} \Pr\left(\frac{1}{n} \sum_{j=1}^n Z_{f,j} \leq t\right) &= \Pr\left(\frac{1}{n} S_{f,n} \geq -t\right) \leq \exp(\eta t + \log E \exp(\eta S_{f,n})) \\ &= \exp(\eta t) (E \exp(-\eta Z_f))^n. \end{aligned}$$

Making the replacement $\Lambda_{-Z_f}(\eta) = \log \mathbb{E} \exp(-\eta Z_f)$ yields

$$\log \Pr \left(\frac{1}{n} S_{f,n} \geq -t \right) \leq \eta t + n \Lambda_{-Z_f}(\eta).$$

By assumption, $\Lambda_{-Z_f}(\eta) \leq -\frac{a}{n}$, and so $\Pr\{P_n Z_f \leq t\} \leq \exp(-a + \eta t)$ as desired. \square

This theorem will be applied by showing that for an excess loss random variable Z_f taking values in $[-1, 1]$, if for some $\eta > 0$ we have $\mathbb{E} \exp(-\eta Z_f) = 1$ and if $\mathbb{E} Z_f = \frac{a}{n}$ for some constant a (that can and must depend on n), then $\Lambda_{-Z_f}(\eta/2) \leq -\frac{c\eta a}{n}$ where $c > 0$ is a universal constant. This is the nature of the next section. We then extend this result to random variables taking values in $[-V, V]$.

3 Semi-infinite linear programming and the general moment problem

The key subproblem now is to find, for each excess loss random variable Z_f with mean $\frac{a}{n}$ and $\Lambda_{-Z_f}(\eta) = 0$ (for some $\eta \geq \eta^*$), a pair of constants $\eta_0 > 0$ and $c > 0$ for which $\Lambda_{-Z_f}(\eta_0) \leq -\frac{ca}{n}$. [Theorem 1](#) would then imply that ERM will prefer f^* over this particular f with high probability for ca large enough. This subproblem is in fact an instance of the general moment problem, a problem on which Kemperman [\[9\]](#) has conducted a very nice geometric study. We now describe this problem.

The general moment problem. Let $\mathcal{P}(\mathcal{A})$ be the space of probability measures over a measurable space $\mathcal{A} = (\mathcal{A}, \mathcal{S})$. For real-valued measurable functions h and $(g_j)_{j \in [m]}$ on a measurable space $\mathcal{A} = (\mathcal{A}, \mathcal{S})$, the general moment problem is

$$\begin{aligned} & \inf_{\mu \in \mathcal{P}(\mathcal{A})} \mathbb{E}_{X \sim \mu} h(X) \\ & \text{subject to } \mathbb{E}_{X \sim \mu} g_j(X) = y_j, \quad j \in \{1, \dots, m\}. \end{aligned} \quad (3)$$

Let the vector-valued map $g : \mathcal{A} \rightarrow \mathbb{R}^m$ be defined in terms of coordinate functions as $(g(x))_j = g_j(x)$, and let the vector $y \in \mathbb{R}^m$ be equal to (y_1, \dots, y_m) .

Let $D^* \subset \mathbb{R}^{m+1}$ be the set

$$D^* := \left\{ d^* = (d_0, d_1, \dots, d_m) \in \mathbb{R}^{m+1} : h(x) \geq d_0 + \sum_{j=1}^m d_j g_j(x) \quad \text{for all } x \in \mathcal{A} \right\}. \quad (4)$$

Theorem 3 of [\[9\]](#) states that if $y \in \text{int conv } g(\mathcal{A})$, the optimal value of problem (3) equals

$$\sup \left\{ d_0 + \sum_{j=1}^m d_j y_j : d^* = (d_0, d_1, \dots, d_m) \in D^* \right\}. \quad (5)$$

Our instantiation. We choose $\mathcal{A} = [-1, 1]$, set $m = 2$ and define $h, (g_j)_{j \in \{1,2\}}$, and $y \in \mathbb{R}^2$ as:

$$h(x) = -e^{(\eta/2)x}, \quad g_1(x) = x, \quad g_2(x) = e^{\eta x}, \quad y_1 = -\frac{a}{n}, \quad y_2 = 1,$$

for any $\eta > 0$, $a > 0$, and $n \in \mathbb{N}$. This yields the following instantiation of problem (3):

$$\inf_{\mu \in \mathcal{P}([-1,1])} \mathbb{E}_{X \sim \mu} -e^{(\eta/2)X} \quad (6a)$$

$$\text{subject to } \mathbb{E}_{X \sim \mu} X = -\frac{a}{n} \quad (6b)$$

$$\mathbb{E}_{X \sim \mu} e^{\eta X} = 1. \quad (6c)$$

Note that equation (5) from the general moment problem now instantiates to

$$\sup \left\{ d_0 - \frac{a}{n} d_1 + d_2 : d^* = (d_0, d_1, d_2) \in D^* \right\}, \quad (7)$$

with D^* equal to the set

$$\left\{ d^* = (d_0, d_1, d_2) \in \mathbb{R}^3 : -e^{(\eta/2)x} \geq d_0 + d_1 x + d_2 e^{\eta x} \quad \text{for all } x \in [-1, 1] \right\}. \quad (8)$$

Applying Theorem 3 of [\[9\]](#) requires the condition $y \in \text{int conv } g([-1, 1])$. We first characterize when $y \in \text{conv } g([-1, 1])$ holds and handle the $\text{int conv } g([-1, 1])$ version after [Theorem 3](#).

Lemma 2 (Feasible Moments). *The point $y = (-\frac{a}{n}, 1) \in \text{conv } g([-1, 1])$ if and only if*

$$\frac{a}{n} \leq \frac{e^\eta + e^{-\eta} - 2}{e^\eta - e^{-\eta}} = \frac{\cosh(\eta) - 1}{\sinh(\eta)}. \quad (9)$$

Proof. Let W denote the convex hull of $g([-1, 1])$. We need to see if $(-\frac{a}{n}, 1) \in W$. Note that W is the convex set formed by starting with the graph of $x \mapsto e^{\eta x}$ on the domain $[-1, 1]$, including the line segment connecting this curve's endpoints $(-1, e^{-\eta})$ to $(1, e^\eta)$, and including all of the points below this line segment but above the aforementioned graph. That is, W is precisely the set

$$W := \left\{ (x, y) \in \mathbb{R}^2 : e^{\eta x} \leq y \leq \frac{e^\eta + e^{-\eta}}{2} + \frac{e^\eta - e^{-\eta}}{2}x, \forall x \in [-1, 1] \right\}.$$

It remains to check that 1 is sandwiched between the lower and upper bounds at $x = -\frac{a}{n}$. Clearly the lower bound holds. Simple algebra shows that the upper bound is equivalent to condition (9). \square

Note that if (9) does not hold, then the semi-infinite linear program (6) is infeasible; infeasibility in turn implies that such an excess loss random variable cannot exist. Thus, we need not worry about whether (9) holds; it holds for *any* excess loss random variable satisfying constraints (6b) and (6c).

The following theorem is a key technical result for using stochastic mixability to control the CGF. The proof is long and can be found in [Appendix A](#).

Theorem 3 (Stochastic Mixability Concentration). *Let f be an element of \mathcal{F} with Z_f taking values in $[-1, 1]$, $n \in \mathbb{N}$, $\mathbb{E} Z_f = \frac{a}{n}$ for some $a > 0$, and $\Lambda_{-Z_f}(\eta) = 0$ for some $\eta > 0$. If*

$$\frac{a}{n} < \frac{e^\eta + e^{-\eta} - 2}{e^\eta - e^{-\eta}}, \quad (10)$$

then
$$\mathbb{E} e^{(\eta/2)(-Z_f)} \leq 1 - \frac{0.18(\eta \wedge 1)a}{n}.$$

Note that since $\log(1 - x) \leq -x$ when $x < 1$, we have $\Lambda_{-Z_f}(\eta/2) \leq -\frac{0.18(\eta \wedge 1)a}{n}$.

In order to apply [Theorem 3](#), we need (10) to hold, but only (9) is guaranteed to hold. The corner case is if (9) holds with equality. However, observe that one can always approximate the random variable X by a perturbed version X' which has nearly identical mean $a' \approx a$ and a nearly identical $\eta' \approx \eta$ for which $\mathbb{E}_{X' \sim \mu'} e^{\eta' X'} = 1$, and yet the inequality in (9) is strict. Later, in the proof of [Theorem 5](#), for any random variable that required perturbation to satisfy the interior condition (10), we implicitly apply the analysis to the perturbed version, show that ERM would not pick the (slightly different) function corresponding to the perturbed version, and use the closeness of the two functions to show that ERM also would not pick the original function.

We now present a necessary extension for the case of losses with range $[0, V]$, proved in [Appendix A](#).

Lemma 4 (Bounded Losses). *Let $g_1(x) = x$ and $y_2 = 1$ be common settings for the following two problems. The instantiation of problem (3) with $\mathcal{A} = [-V, V]$, $h(x) = -e^{(\eta/2)x}$, $g_2(x) = e^{\eta x}$, and $y_1 = -\frac{a}{n}$ has the same optimal value as the instantiation of problem (3) with $\mathcal{A} = [-1, 1]$, $h(x) = -e^{(V\eta/2)x}$, $g_2(x) = e^{(V\eta)x}$, and $y_1 = -\frac{a/V}{n}$.*

4 Fast rates

We now show how the above results can be used to obtain an exact oracle inequality with a fast rate. We first present a result for finite classes and then present a result for VC-type classes (classes with logarithmic universal metric entropy).

Theorem 5 (Finite Classes Exact Oracle Inequality). *Let $(\ell, \mathcal{F}, \mathbb{P})$ be η^* -stochastically mixable, where $|\mathcal{F}| = N$, ℓ is a nonnegative loss, and $\sup_{f \in \mathcal{F}} \ell(Y, f(X)) \leq V$ a.s. for a constant V . Then for all $n \geq 1$, with probability at least $1 - \delta$*

$$\mathbb{P} \ell(\cdot, \hat{f}_{\mathbf{z}}) \leq \mathbb{P} \ell(\cdot, f^*) + \frac{6 \max \left\{ V, \frac{1}{\eta^*} \right\} (\log \frac{1}{\delta} + \log N)}{n}.$$

Proof. Let $\gamma_n = \frac{a}{n}$ for a constant a to be fixed later. For each $\eta > 0$, let $\mathcal{F}_{\geq \gamma_n}^{(\eta)} \subset \mathcal{F}_{\geq \gamma_n}$ correspond to those functions in $\mathcal{F}_{\geq \gamma_n}$ for which η is the largest constant such that $\mathbb{E} \exp(-\eta Z_f) = 1$. Let $\mathcal{F}_{\geq \gamma_n}^{\text{hyper}} \subset \mathcal{F}_{\geq \gamma_n}$ correspond to functions f in $\mathcal{F}_{\geq \gamma_n}$ for which $\lim_{\eta \rightarrow \infty} \mathbb{E} \exp(-\eta Z_f) < 1$. Clearly, $\mathcal{F}_{\geq \gamma_n} = \left(\bigcup_{\eta \in [\eta^*, \infty)} \mathcal{F}_{\geq \gamma_n}^{(\eta)} \right) \cup \mathcal{F}_{\geq \gamma_n}^{\text{hyper}}$. The excess loss random variables corresponding to elements $f \in \mathcal{F}_{\geq \gamma_n}^{\text{hyper}}$ are “hyper-concentrated” in the sense that they are infinitely stochastically mixable. However, [Lemma 10](#) in [Appendix B](#) shows that for each hyper-concentrated Z_f , there exists another excess loss random variable Z'_f with mean arbitrarily close to that of Z_f , with $\mathbb{E} \exp(-\eta Z'_f) = 1$ for some arbitrarily large but finite η , and with $Z'_f \leq Z_f$ with probability 1. The last property implies that the empirical risk of Z'_f is no greater than that of Z_f ; hence for each hyper-concentrated Z_f it is sufficient (from the perspective of ERM) to study a corresponding Z'_f . From now on, we implicitly make this replacement in $\mathcal{F}_{\geq \gamma_n}$ itself, so that we now have $\mathcal{F}_{\geq \gamma_n} = \bigcup_{\eta \in [\eta^*, \infty)} \mathcal{F}_{\geq \gamma_n}^{(\eta)}$.

Consider an arbitrary $a > 0$. For some fixed $\eta \in [\eta^*, \infty)$ for which $|\mathcal{F}_{\geq \gamma_n}^{(\eta)}| > 0$, consider the subclass $\mathcal{F}_{\geq \gamma_n}^{(\eta)}$. Individually for each such function, we will apply [Theorem 1](#) as follows. From [Lemma 4](#), we have $\Lambda_{-Z_f}(\eta/2) = \Lambda_{-\frac{1}{V}Z_f}(V\eta/2)$. From [Theorem 3](#), the latter is at most $-\frac{0.18(V\eta \wedge 1)(a/V)}{n} = -\frac{0.18\eta a}{(V\eta \vee 1)n}$. Hence, [Theorem 1](#) with $t = 0$ and the η from the Theorem taken to be $\eta/2$ implies that the probability of the event $\mathbb{P}_n \ell(\cdot, f) \leq \mathbb{P}_n \ell(\cdot, f^*)$ is at most $\exp\left(-0.18 \frac{\eta}{V\eta \vee 1} a\right)$. Applying the union bound over all of $\mathcal{F}_{\geq \gamma_n}$, we conclude that

$$\Pr \{ \exists f \in \mathcal{F}_{\geq \gamma_n} : \mathbb{P}_n \ell(\cdot, f) \leq \mathbb{P}_n \ell(\cdot, f^*) \} \leq N \exp\left(-\eta^* \left(\frac{0.18a}{V\eta^* \vee 1}\right)\right).$$

Since ERM selects hypotheses on their empirical risk, from inversion it holds that with probability at least $1 - \delta$ ERM will not select any hypothesis with excess risk at least $\frac{6 \max\{V, \frac{1}{\eta^*}\} (\log \frac{1}{\delta} + \log N)}{n}$. \square

Before presenting the result for VC-type classes, we require some definitions. For a pseudometric space (\mathcal{G}, d) , for any $\varepsilon > 0$, let $\mathcal{N}(\varepsilon, \mathcal{G}, d)$ be the ε -covering number of (\mathcal{G}, d) ; that is, $\mathcal{N}(\varepsilon, \mathcal{G}, d)$ is the minimal number of balls of radius ε needed to cover \mathcal{G} . We will further constrain the cover (the set of centers of the balls) to be a subset of \mathcal{G} (i.e. to be proper), thus ensuring that the stochastic mixability assumption transfers to any (proper) cover of \mathcal{F} . Note that the “proper” requirement at most doubles the constant K below, as shown by [Vidyasagar \[22, Lemma 2.1\]](#).

We now state a localization-based result that allows us to extend the result for finite classes to VC-type classes. Although the localization result can be obtained by combining standard techniques,¹ we could not find this particular result in the literature. Below, an ε -net \mathcal{F}_ε of a set \mathcal{F} is a subset of \mathcal{F} such that \mathcal{F} is contained in the union of the balls of radius ε with centers in \mathcal{F}_ε .

Theorem 6. *Let \mathcal{F} be a separable function class whose functions have range bounded in $[0, V]$ and for which, for a constant $K \geq 1$, for each $u \in (0, K]$ the $L_2(\mathbb{P})$ covering numbers are bounded as*

$$\mathcal{N}(u, \mathcal{F}, L_2(\mathbb{P})) \leq \left(\frac{K}{u}\right)^c. \quad (11)$$

Suppose \mathcal{F}_ε is a minimal ε -net for \mathcal{F} in the $L_2(\mathbb{P})$ norm, with $\varepsilon = \frac{1}{n}$. Denote by $\pi : \mathcal{F} \rightarrow \mathcal{F}_\varepsilon$ an $L_2(\mathbb{P})$ -metric projection from \mathcal{F} to \mathcal{F}_ε . Then, provided that $\delta \leq \frac{1}{2}$, with probability at most δ can there exist $f \in \mathcal{F}$ such that

$$\mathbb{P}_n f < \mathbb{P}_n(\pi(f)) - \frac{V}{n} \left(1080C \log(2Kn) + 90 \sqrt{\left(\log \frac{1}{\delta}\right) C \log(2Kn) + \log \frac{e}{\delta}} \right).$$

The proof is presented in [Appendix C](#). We now present the fast rates result for VC-type classes. The proof (in [Appendix C](#)) uses [Theorem 6](#) and the proof of the [Theorem 5](#). Below, we denote the loss-composed version of a function class \mathcal{F} as $\ell \circ \mathcal{F} := \{\ell(\cdot, f) : f \in \mathcal{F}\}$.

¹See e.g. the techniques of [Massart and Nédélec \[16\]](#) and equation (3.17) of [Koltchinskii \[11\]](#).

Theorem 7 (VC-Type Classes Exact Oracle Inequality). *Let $(\ell, \mathcal{F}, \mathbb{P})$ be η^* -stochastically mixable with $\ell \circ \mathcal{F}$ separable, where, for a constant $K \geq 1$, for each $\varepsilon \in (0, K]$ we have $\mathcal{N}(\ell \circ \mathcal{F}, L_2(\mathbb{P}), \varepsilon) \leq \left(\frac{K}{\varepsilon}\right)^C$, and $\sup_{f \in \mathcal{F}} \ell(Y, f(X)) \leq V$ a.s. for a constant $V \geq 1$. Then for all $n \geq 5$ and $\delta \leq \frac{1}{2}$, with probability at least $1 - \delta$*

$$\mathbb{P} \ell(\cdot, \hat{f}_z) \leq \mathbb{P} \ell(\cdot, f^*) + \frac{1}{n} \max \left\{ \frac{8 \max \left\{ V, \frac{1}{\eta^*} \right\} \left(C \log(Kn) + \log \frac{2}{\delta} \right)}{2V \left(1080C \log(2Kn) + 90 \sqrt{\left(\log \frac{2}{\delta} \right) C \log(2Kn) + \log \frac{2e}{\delta}} \right)} \right\} + \frac{1}{n}.$$

5 Characterizing convexity from the perspective of risk minimization

In the following, when we say $(\ell, \mathcal{F}, \mathbb{P})$ has a *unique minimizer* we mean that any two minimizers f_1^*, f_2^* of $\mathbb{P} \ell(\cdot, f)$ over \mathcal{F} satisfy $\ell(Y, f_1^*(X)) = \ell(Y, f_2^*(X))$ a.s. We say the excess loss class $\{\ell(\cdot, f) - \ell(\cdot, f^*) : f \in \mathcal{F}\}$ satisfies a (β, B) -Bernstein condition with respect to \mathbb{P} for some $B > 0$ and $0 < \beta \leq 1$ if, for all $f \in \mathcal{F}$:

$$\mathbb{P}(\ell(\cdot, f) - \ell(\cdot, f^*))^2 \leq B \left(\mathbb{P}(\ell(\cdot, f) - \ell(\cdot, f^*)) \right)^\beta. \quad (12)$$

It already is known that the stochastic mixability condition guarantees that there is a unique minimizer [21]; this is a simple consequence of Jensen's inequality. This leaves open the question: if stochastic mixability does not hold, are there necessarily non-unique minimizers? We show that in a certain sense this is indeed the case, in bad way: the set of minimizers will be a disconnected set.

For any $\varepsilon > 0$, define \mathcal{G}_ε as the class $\mathcal{G}_\varepsilon := \{f^*\} \cup \{f \in \mathcal{F} : \|f - f^*\|_{L_1(\mathbb{P})} \geq \varepsilon\}$, where in case there are multiple minimizers in \mathcal{F} we arbitrarily select one of them as f^* . Since we assume that \mathcal{F} is compact and $\mathcal{G}_\varepsilon \setminus \{f^*\}$ is equal to \mathcal{F} minus an open set homeomorphic to the unit $L_1(\mathbb{P})$ ball, $\mathcal{G}_\varepsilon \setminus \{f^*\}$ is also compact.

Theorem 8 (Non-Unique Minimizers). *Suppose there exists some $\varepsilon > 0$ such that \mathcal{G}_ε is not stochastically mixable. Then there are minimizers $f_1^*, f_2^* \in \mathcal{F}$ of $\mathbb{P} \ell(\cdot, f)$ over \mathcal{F} such that it is not the case that $\ell(Y, f_1^*(X)) = \ell(Y, f_2^*(X))$ a.s.*

Proof. Select $\varepsilon > 0$ as in the theorem and some fixed $\eta > 0$. Since \mathcal{G}_ε is not η -stochastically mixable, there exists $f_\eta \in \mathcal{G}_\varepsilon$ such that $\Lambda_{-Z_{f_\eta}}(\eta) > 0$. Note that there exists $\eta' \in (0, \eta)$ with $\Lambda_{-Z_{f_\eta}}(\eta') = 0$; if not, $\lim_{\eta \downarrow 0} \frac{\Lambda_{-Z_{f_\eta}}(\eta) - \Lambda_{-Z_{f_\eta}}(0)}{\eta} > 0 \Rightarrow \Lambda'_{-Z_{f_\eta}}(0) > 0$, so $\Lambda'_{-Z_{f_\eta}}(0) = \mathbb{E}(-Z_{f_\eta})$ implies that $\mathbb{E} Z_{f_\eta} < 0$, a contradiction! From Lemma 2, $\mathbb{E} Z_{f_\eta} \leq \frac{\cosh(\eta') - 1}{\sinh(\eta')}$; for $\eta' \geq 0$ the RHS has upper bound $\frac{\eta'}{2}$ since the derivative of $\frac{\eta'}{2} - \frac{\cosh(\eta') - 1}{\sinh(\eta')}$ is the nonnegative function $\frac{1}{2} \tanh^2(\eta'/2)$ and $\left(\frac{\eta'}{2} - \frac{\cosh(\eta') - 1}{\sinh(\eta')}\right) \Big|_{\eta'=0} = 0$. Thus, $\mathbb{E} Z_{f_\eta} \rightarrow 0$ as $\eta \rightarrow 0$. As $\mathcal{G}_\varepsilon \setminus \{f^*\}$ is compact, we can take a positive decreasing sequence $(\eta_j)_j$ approaching 0, corresponding to a sequence $(f_{\eta_j})_j \subset \mathcal{G}_\varepsilon \setminus \{f^*\}$ with limit point $g^* \in \mathcal{G}_\varepsilon \setminus \{f^*\}$ for which $\mathbb{E} Z_{g^*} = 0$, and so there is a risk minimizer in $\mathcal{G}_\varepsilon \setminus \{f^*\}$. \square

The implications of having non-unique risk minimizers. In the case of non-unique risk minimizers, Mendelson [17] showed that for p -losses $(y, \hat{y}) \mapsto |y - \hat{y}|^p$ with $p \in [2, \infty)$ there is an n -indexed sequence of probability measures $(\mathbb{P}^{(n)})_n$ approaching the true probability measure as $n \rightarrow \infty$ such that, for each n , ERM learns at a slow rate under sample size n when the true distribution is $\mathbb{P}^{(n)}$. This behavior is a consequence of the statistical learning problem's poor geometry: there are multiple minimizers and the set of minimizers is not even connected. Furthermore, in this case, the best known fast rate upper bounds (see [18] and [19]) have a multiplicative constant that approaches ∞ as the target probability measure approaches a probability measure for which there are non-unique minimizers. The reason for the poor upper bounds in this case is that the constant B in the Bernstein condition explodes, and the upper bounds rely upon the Bernstein condition.

6 Weak stochastic mixability

For some $\kappa \in [0, 1]$, we say $(\ell, \mathcal{F}, \mathbb{P})$ is (κ, η_0) -weakly stochastically mixable if, for every $\varepsilon > 0$, for all $f \in \{f^*\} \cup \mathcal{F}_{\geq \varepsilon}$, the inequality $\log \mathbb{E} \exp(-\eta_\varepsilon Z_f) \leq 0$ holds with $\eta_\varepsilon := \eta_0 \varepsilon^{1-\kappa}$. This concept was introduced by Van Erven et al. [21] without a name.

Suppose that some fixed function has excess risk $a = \varepsilon$. Then, roughly, with high probability ERM does not make a mistake provided that $a\eta_n = \frac{1}{n}$, i.e. when $\varepsilon \cdot \eta_0 \varepsilon^{1-\kappa} = \frac{1}{n}$ and hence when $\varepsilon = (\eta_0 n)^{-1/(2-\kappa)}$. Modifying the proof of the finite classes result ([Theorem 5](#)) to consider all functions in the subclass $\mathcal{F}_{\geq \gamma_n}$ for $\gamma_n = (\eta_0 n)^{-1/(2-\kappa)}$ yields the following corollary of [Theorem 5](#).

Corollary 9. *Let $(\ell, \mathcal{F}, \mathbb{P})$ be (κ, η_0) -weakly stochastically mixable for some $\kappa \in [0, 1]$, where $|\mathcal{F}| = N$, ℓ is a nonnegative loss, and $\sup_{f \in \mathcal{F}} \ell(Y, f(X)) \leq V$ a.s. for a constant V . Then for any $n \geq \frac{1}{\eta_0} V^{(1-\kappa)/(2-\kappa)}$, with probability at least $1 - \delta$*

$$\mathbb{P} \ell(\cdot, \hat{f}_{\mathbf{z}}) \leq \mathbb{P} \ell(\cdot, f^*) + \frac{6 \left(\log \frac{1}{\delta} + \log N \right)}{(\eta_0 n)^{1/(2-\kappa)}}.$$

It is simple to show a similar result for VC-type classes; the ε -net can still be taken at the resolution $\frac{1}{n}$, but we need only apply the analysis to the subclass of \mathcal{F} with excess risk at least $(\eta_0 n)^{-1/(2-\kappa)}$.

7 Discussion

We have shown that stochastic mixability implies fast rates for VC-type classes, using a direct argument based on the Cramér-Chernoff method and sufficient control of the optimal value of a certain instance of the general moment problem. The approach is amenable to localization in that the analysis separately controls the probability of large deviations for individual elements of \mathcal{F} . An important open problem is to extend the results presented here for VC-type classes to results for nonparametric classes with polynomial metric entropy, and moreover, to achieve rates similar to those obtained for these classes under the Bernstein condition.

There are still some unanswered questions with regards to the connection between the Bernstein condition and stochastic mixability. Van Erven et al. [[21](#)] showed that for bounded losses the Bernstein condition implies stochastic mixability. Therefore, when starting from a Bernstein condition, [Theorem 5](#) offers a different path to fast rates. An open problem is to settle the question of whether the Bernstein condition and stochastic mixability are equivalent. Previous results [[21](#)] suggest that the stochastic mixability does imply a Bernstein condition, but the proof was non-constructive, and it relied upon a bounded losses assumption. It is well known (and easy to see) that both stochastic mixability and the Bernstein condition hold only if there is a unique minimizer. [Theorem 8](#) shows in a certain sense that if stochastic mixability does not hold, then there cannot be a unique minimizer. Is the same true when the Bernstein condition fails to hold? Regardless of whether stochastic mixability is equivalent to the Bernstein condition, the direct argument presented here and the connection to classical mixability, which does characterize constant regret in the simpler non-stochastic setting, motivates further study of stochastic mixability.

Finally, it would be of great interest to discard the bounded losses assumption. Ignoring the dependence of the metric entropy on the maximum possible loss, the upper bound on the loss V enters the final bound through the difficulty of controlling the minimum value of $u_\eta(-1)$ when η is large (see the proof of [Theorem 3](#)). From extensive experiments with a grid-approximation linear program, we have observed that the worst (CGF-wise) random variables for fixed negative mean and fixed optimal stochastic mixability constant are those which place very little probability mass at $-V$ and most of the probability mass at a small positive number that scales with the mean. These random variables correspond to functions that with low probability beat f^* by a large (loss) margin but with high probability have slightly higher loss than f^* . It would be useful to understand if this exotic behavior is a real concern and, if not, find a simple, mild condition on the moments that rules it out.

Acknowledgments

RCW thanks Tim van Erven for the initial discussions around the Cramér-Chernoff method during his visit to Canberra in 2013 and for his gracious permission to proceed with the present paper without him as an author, and both authors thank him for the further enormously helpful spotting of a serious error in our original proof for fast rates for VC-type classes. This work was supported by the Australian Research Council (NAM and RCW) and NICTA (RCW). NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence program.

References

- [1] Jacob Abernethy, Alekh Agarwal, Peter L. Bartlett, and Alexander Rakhlin. A stochastic view of optimal regret through minimax duality. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT 2009)*, 2009.
- [2] Jean-Yves Audibert. Fast learning rates in statistical inference through aggregation. *The Annals of Statistics*, 37(4):1591–1646, 2009.
- [3] Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.
- [4] Peter L. Bartlett and Shahar Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.
- [5] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press, 2013.
- [6] Peter Grünwald. Safe learning: bridging the gap between Bayes, MDL and statistical learning theory via empirical convexity. In *Proceedings of the 24th International Conference on Learning Theory (COLT 2011)*, pages 397–419, 2011.
- [7] Peter Grünwald. The safe Bayesian. In *Proceedings of the 23rd International Conference on Algorithmic Learning Theory (ALT 2012)*, pages 169–183. Springer, 2012.
- [8] Yuri Kalnishkan and Michael V. Vyugin. The weak aggregating algorithm and weak mixability. In *Proceedings of the 18th Annual Conference on Learning Theory (COLT 2005)*, pages 188–203. Springer, 2005.
- [9] Johannes H.B. Kemperman. The general moment problem, a geometric approach. *The Annals of Mathematical Statistics*, 39(1):93–122, 1968.
- [10] Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006.
- [11] Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems: Ecole d'Été de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer, 2011.
- [12] Guillaume Lecué. *Interplay between concentration, complexity and geometry in learning theory with applications to high dimensional data analysis*. Habilitation à diriger des recherches, Université Paris-Est, 2011.
- [13] Wee Sun Lee, Peter L. Bartlett, and Robert C. Williamson. The importance of convexity in learning with squared loss. *IEEE Transactions on Information Theory*, 44(5):1974–1980, 1998.
- [14] Jonathan Qiang Li. *Estimation of mixture models*. PhD thesis, Yale University, 1999.
- [15] Enno Mammen and Alexandre B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.
- [16] Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- [17] Shahar Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Transactions on Information Theory*, 54(8):3797–3803, 2008.
- [18] Shahar Mendelson. Obtaining fast error rates in nonconvex situations. *Journal of Complexity*, 24(3):380–397, 2008.
- [19] Shahar Mendelson and Robert C. Williamson. Agnostic learning nonconvex function classes. In *Proceedings of the 15th Annual Conference on Computational Learning Theory (COLT 2002)*, pages 1–13. Springer, 2002.
- [20] Alexander B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- [21] Tim Van Erven, Peter D. Grünwald, Mark D. Reid, and Robert C. Williamson. Mixability in statistical learning. In *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, pages 1700–1708, 2012.
- [22] Mathukumalli Vidyasagar. *Learning and Generalization with Applications to Neural Networks*. Springer, 2002.
- [23] Volodya Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56(2):153–173, 1998.
- [24] Volodya Vovk. Competitive on-line statistics. *International Statistical Review*, 69(2):213–248, 2001.

A Proof of Theorem 3

Proof of Theorem 3. From Theorem 3 of Kemperman [9], if the moment values vector $(-\frac{a}{n}, 1)$ is in $\text{int conv } g([-1, 1])$, the optimal objective value of problem (6) is equal to

$$\sup \left\{ d_0 - \frac{a}{n}d_1 + d_2 : d^* = (d_0, d_1, d_2) \in D^* \right\}. \quad (13)$$

From the Feasible Moments Lemma (Lemma 2 in the paper), we see that (10) corresponds to the interior point condition.

Since we assume the interior point condition is satisfied, any $d^* \in D^*$ provides a lower bound on the optimal value of (6), and hence after negation provides an upper bound on the problem with same moment constraints and the objective $\sup E e^{(\eta/2)X}$; this is precisely what we are after.

We therefore focus on picking a good $d^* = (d_0, d_1, d_2) \in \mathbb{R}^3$. The inequality condition in (8) is now

$$-e^{(\eta/2)x} \geq d_0 + d_1x + d_2e^{\eta x} \quad \text{for all } x \in [-1, 1].$$

In particular, this inequality must hold at $x = 0$, yielding the constraint $-1 \geq d_0 + d_2$. We now change variables to $c_0 = -d_0$, $c_1 = -d_1/\eta$,² and $c_2 = -d_2$, yielding the inequality condition

$$u_\eta(x) := -e^{(\eta/2)x} + c_0 + c_2e^{\eta x} + \eta c_1x \geq 0.$$

Now the condition at $x = 0$ implies that $c_0 + c_2 = 1$, and so we make the replacement

$$c_0 = 1 - c_2, \quad (14)$$

in the definition of $u_\eta(x)$, yielding the inequality

$$u_\eta(x) = 1 + c_2(e^{\eta x} - 1) - e^{(\eta/2)x} + \eta c_1x \geq 0.$$

Constraints from the local minimum at 0

Since $u_\eta(0) = 0$, we need $x = 0$ to be a local minimum of u , and so we require both conditions

- (a) $u'(0) = 0$
- (b) $u''(0) \geq 0$

to hold since otherwise there exists some small $\varepsilon > 0$ such that either $u_\eta(\varepsilon) < 0$ or $u_\eta(-\varepsilon) < 0$.

For (a), we compute

$$u'(x) = \eta c_2 e^{\eta x} - \frac{\eta}{2} e^{(\eta/2)x} + \eta c_1.$$

Since we require $u'(0) = 0$, we pick up the constraint

$$\eta \left(c_2 - \frac{1}{2} + c_1 \right) = 0,$$

and since $\eta > 0$ by assumption, we have

$$c_1 = \frac{1}{2} - c_2. \quad (15)$$

Thus, we can eliminate c_1 from $u_\eta(x)$:

$$u_\eta(x) = 1 + c_2(e^{\eta x} - 1) - e^{(\eta/2)x} + \eta \left(\frac{1}{2} - c_2 \right) x \geq 0.$$

For (b), it is sufficient to have $u''(0) > 0$. Observe that

$$u''(x) = \eta^2 c_2 e^{\eta x} - \frac{\eta^2}{4} e^{(\eta/2)x},$$

so that $u''(0) = \eta^2 (c_2 - \frac{1}{4})$, and hence for

$$c_2 > \frac{1}{4} \quad (16)$$

we have $u''(0) > 0$.

Thus far, we have picked up the constraints (14), (15), and (16).

²We scale by η here because we are chasing a certain η -dependent rate.

The other minima of $u_\eta(x)$

Now, observe that $u'(x)$ has at most two roots, because with the substitution $y = e^{(\eta/2)x}$, we have

$$u'(x) = \eta c_2 y^2 - \frac{\eta}{2} y + \eta \left(\frac{1}{2} - c_2 \right),$$

which is a quadratic equation in y with two roots:

$$y = \left\{ 1, \frac{1-2c_2}{2c_2} \right\} \Rightarrow x = \left\{ 0, \frac{2}{\eta} \log \frac{1-2c_2}{2c_2} \right\}.$$

Now, since we take $c_2 > \frac{1}{4}$ and since the second root is negative, we know that u is increasing on $[0, 1]$ (and we already knew that $u_\eta(0) = 0$). It remains to find conditions on c_2 such that $u_\eta(-1) \geq 0$ because that implies that $u_\eta(x) \geq 0$ for all $x \in [-1, 0]$. We consider the case $\eta \leq 1$ and $\eta > 1$ separately.

In either case, we need to check the nonnegativity of

$$\begin{aligned} u_\eta(-1) &= 1 + c_2(e^{-\eta} - 1) - e^{-(\eta/2)} - \eta \left(\frac{1}{2} - c_2 \right) \\ &= \left(1 - \frac{\eta}{2} \right) - e^{-(\eta/2)} + c_2 (e^{-\eta} - (1 - \eta)). \end{aligned}$$

Case $\eta \leq 1$: We observe that $u_\eta(-1) = 0$ when $\eta = 0$. Now, we will see what constraints on c_2 guarantee that $\frac{d}{d\eta} u_\eta(-1) \geq 0$ for $\eta \in [0, 1]$. We want

$$\frac{d}{d\eta} u_\eta(-1) = -c_2 e^{-\eta} + \frac{1}{2} e^{-\eta/2} - \frac{1}{2} + c_2 \geq 0$$

which is equivalent to the condition

$$c_2 \geq \frac{1}{2} \left(\frac{1 - e^{-\eta/2}}{1 - e^{-\eta}} \right).$$

The RHS is increasing in η , and so we need only consider $\eta = 1$, yielding the bound

$$c_2 \geq \frac{1}{2} \frac{e - \sqrt{e}}{e - 1} = 0.3112\dots,$$

and so if $c_2 \geq 0.32$, then $u_\eta(-1) \geq 0$ as desired.

Case $\eta > 1$: Let $c_2 = \frac{1}{2} - \frac{\alpha}{\eta}$ for some $\alpha \geq 0$. With this substitution, we have

$$\begin{aligned} u_\eta(-1) &= 1 + c_2(e^{-\eta} - 1) - e^{-(\eta/2)} - \eta \left(\frac{1}{2} - c_2 \right) \\ &= 1 + \left(\frac{1}{2} - \frac{\alpha}{\eta} \right) (e^{-\eta} - 1) - e^{-(\eta/2)} - \alpha \\ &= \left(\frac{1 + e^{-\eta}}{2} - e^{-\eta/2} \right) + \alpha \left(-1 + \frac{1}{\eta} (1 - e^{-\eta}) \right) \end{aligned}$$

Since we want the above to be nonnegative for all $\eta > 1$, we arrive at the condition

$$\alpha \leq \inf_{\eta \geq 1} \left\{ \frac{\frac{1+e^{-\eta}}{2} - e^{-\eta/2}}{1 - \frac{1}{\eta} (1 - e^{-\eta})} \right\} \quad (17)$$

Plotting suggests that the minimum is attained at $\eta = 1$, with the value $\frac{1}{2}(\sqrt{e} - 1)^2$. We will fix α to this value and verify that

$$\left(\frac{1 + e^{-\eta}}{2} - e^{-\eta/2} \right) + \left(\frac{1}{2}(\sqrt{e} - 1)^2 \right) \left(-1 + \frac{1}{\eta} (1 - e^{-\eta}) \right) \geq 0. \quad (18)$$

This is true with equality at $\eta = 0$. The derivative of the LHS with respect to η is

$$\frac{1}{2}e^{-\eta} \left(e^{\eta/2} - 1 - \frac{(\sqrt{e}-1)^2(e^\eta - \eta - 1)}{\eta^2} \right).$$

The derivative is positive at $\eta = 1$, so 0 is a candidate minimum. Eventually, $\frac{(\sqrt{e}-1)^2(e^\eta - \eta - 1)}{\eta^2}$ grows more quickly than $e^{\eta/2} - 1$ and surpasses the latter in value. The derivative is therefore negative for all sufficiently large η , and so we need only take the minimum of the LHS of (18) evaluated at $\eta = 1$ and the limiting value as $\eta \rightarrow \infty$. We have

$$\lim_{\eta \rightarrow \infty} \left(\frac{1 + e^{-\eta}}{2} - e^{-\eta/2} \right) + \left(\frac{1}{2}(\sqrt{e}-1)^2 \right) \left(-1 + \frac{1}{\eta} (1 - e^{-\eta}) \right) = \sqrt{e} - \frac{e}{2} \geq 0$$

Hence, (18) indeed holds for $\alpha = \frac{1}{2}(\sqrt{e}-1)^2$. We conclude that $u_\eta(-1) \geq 0$ when $\alpha \leq \frac{1}{2}(\sqrt{e}-1)^2$.

Putting it all together

In the regime $\eta \leq 1$, we have the constraints $c_2 > \frac{1}{4}$ and $c_2 \geq \frac{1}{2} \frac{e-\sqrt{e}}{e-1}$ (which exceeds $\frac{1}{4}$), so we can choose

$$c_1 = \frac{1}{2} - c_2 = \frac{1}{2} - \frac{1}{2} \frac{e-\sqrt{e}}{e-1} = \frac{1}{2} \frac{\sqrt{e}-1}{e-1} = 0.1877 \dots$$

In the regime $\eta > 1$, we have the constraints $c_2 > \frac{1}{4}$ and $\alpha \leq \frac{1}{2}(\sqrt{e}-1)^2 \Rightarrow c_2 \geq \frac{1}{2} - \frac{1}{2\eta}(\sqrt{e}-1)^2$ (which always exceeds $\frac{1}{4}$ for $\eta \geq 1$), so we can choose

$$c_1 = \frac{1}{2} - c_2 = \frac{\alpha}{\eta} \leq \frac{(\sqrt{e}-1)^2}{2\eta} = \frac{0.2104 \dots}{\eta}.$$

The result follows by observing that in the case of $\eta \leq 1$, the supremum in (13) is lower bounded by $-1 + \frac{0.18a\eta}{n}$, and hence the optimal objective value of (6) is lower bounded by the same quantity. Therefore, the problem with the same constraints and the objective $\sup_{\mu \in [-1,1]} \mathbb{E} e^{(\eta/2)X}$ has its optimal objective value upper bounded by $1 - \frac{0.18a\eta}{n}$. Repeat the same argument for the case of $\eta > 1$. \square

Proof of Lemma 4. Let X be a random variable taking values in $[-V, V]$ with mean $-\frac{a}{n}$ and $\mathbb{E} e^{\eta X} = 1$, and let Y be a random variable taking values in $[-1, 1]$ with mean $-\frac{a/V}{n}$ and $\mathbb{E} e^{(V\eta)Y} = 1$. Consider a random variable \tilde{X} that is a $\frac{1}{V}$ -scaled independent copy of X ; observe that $\mathbb{E} \tilde{X} = -\frac{a/V}{n}$ and $\mathbb{E} e^{(V\eta)\tilde{X}} = 1$. Let the maximal possible value of $\mathbb{E} e^{(\eta/2)X}$ be b_X , and let the maximal possible value of $\mathbb{E} e^{(V\eta/2)Y}$ be b_Y . We claim that $b_X = b_Y$. Let X be a random variable with a distribution that maximizes $\mathbb{E} e^{(\eta/2)X}$ subject to the previously stated constraints on X . Since \tilde{X} satisfies $\mathbb{E} e^{(V\eta/2)\tilde{X}} = b_X$, setting $Y = \tilde{X}$ shows that in fact $b_Y \geq b_X$. A symmetric argument (starting with Y and passing to some $\tilde{Y} = VY$) implies that $b_X \geq b_Y$. \square

B Hyper-concentrated excess losses

Lemma 10. *Let Z be a random variable with probability measure P supported on $[-V, V]$. Suppose that $\lim_{\eta \rightarrow \infty} \mathbb{E} \exp(-\eta Z) < 1$ and $\mathbb{E} Z = \mu > 0$. Then there is a suitable modification of Z' for which $Z' \leq Z$ with probability 1, the mean of Z' is arbitrarily close to μ , and $\mathbb{E} \exp(-\eta Z') = 1$ for arbitrarily large η .*

Proof. First, observe that $Z \geq 0$ a.s. If not, then there must be some finite $\eta > 0$ for which $\mathbb{E} \exp(-\eta Z) = 1$. Now, consider a random variable Z' with probability measure Q_ϵ , a modification of Z (with probability measure P) constructed in the following way. Define $A := [\mu, V]$ and $A^- := [-V, -\mu]$. Then for any $\epsilon > 0$ we define Q_ϵ as

$$dQ_\epsilon(z) = \begin{cases} (1-\epsilon)dP(z) & \text{if } z \in A \\ \epsilon dP(-z) & \text{if } z \in A^- \\ dP(z) & \text{otherwise.} \end{cases}$$

Additionally, we couple P and Q_ϵ such that the couple (Z, Z') is a coupling of (P, Q_ϵ) satisfying

$$\mathbb{E}_{(Z, Z') \sim (P, Q_\epsilon)} \mathbf{1}_{\{Z \neq Z'\}} = \min_{(P', Q'_\epsilon)} \mathbb{E}_{(Z, Z') \sim (P', Q'_\epsilon)} \mathbf{1}_{Z \neq Z'},$$

where the min is over all couplings of P and Q_ϵ . This coupling ensures that $Z' \leq Z$ with probability 1; i.e. Z' is dominated by Z .

Now,

$$\begin{aligned} \mathbb{E} \exp(-\eta Z') &= \int_{-V}^V e^{-\eta z} dQ_\epsilon(z) \\ &= \int_{A^-} e^{-\eta z} dQ_\epsilon(z) + \int_A e^{-\eta z} dQ_\epsilon(z) + \int_{[0, V] \setminus A} e^{-\eta z} dQ_\epsilon(z) \\ &= \epsilon \int_{A^-} e^{-\eta z} dP(-z) + (1 - \epsilon) \int_A e^{-\eta z} dP(z) + \int_{[0, V] \setminus A} e^{-\eta z} dP(z) \\ &= \epsilon \int_A e^{\eta z} dP(z) + (1 - \epsilon) \int_A e^{-\eta z} dP(z) + \int_{[0, V] \setminus A} e^{-\eta z} dP(z) \\ &\geq \epsilon e^{\mu \eta} P(A) + (1 - \epsilon) \int_A e^{-\eta z} dP(z) + \int_{[0, V] \setminus A} e^{-\eta z} dP(z). \end{aligned} \quad (19)$$

Now, on the one hand, for any $\eta > 0$, the sum of the two right-most terms in (19) is strictly less than 1 by assumption. On the other hand, $\eta \rightarrow \epsilon P(A) e^{\mu \eta}$ is exponentially increasing since $\epsilon > 0$ and $\mu > 0$ (and hence $P(A) > 0$ as well) by assumption; thus, the first term in (19) can be made arbitrarily large for large enough η . Consequently, we can choose $\epsilon > 0$ as small as desired and then choose $\eta < \infty$ as large as desired such that the mean of Z' is arbitrarily close to μ and $\mathbb{E} \exp(-\eta Z') = 1$ respectively. \square

C Proof of VC-type results

C.1 Proof of Theorem 6

The localization result (Theorem 6) is a simple consequence of the following theorem.

Theorem 11 (Local Analysis). *Let $\mathcal{F} \subset \mathbb{R}^X$ be a separable function class for which:*

- *the constant zero function is an element of \mathcal{F} ;*
- *every function $f \in \mathcal{F}$ satisfies $0 \leq f \leq 1$;*
- $\sup_{f \in \mathcal{F}} \|f\|_{L_2(P)} \leq \varepsilon := \frac{1}{n}$.

Further assume that for some $C \geq 1$, for a constant $K \geq 1$, for each $u \in (0, K]$ the $L_2(P)$ covering numbers of \mathcal{F} are bounded as

$$\mathcal{N}(u, \mathcal{F}, L_2(P)) \leq \left(\frac{K}{u}\right)^C.$$

Then provided that $n \geq 4$ and $y > 0$, with probability at least $1 - e^{-y}$

$$\sup_{f \in \mathcal{F}} P_n f \leq \frac{1}{n} \left(990C \log(2Kn) + \sqrt{2y(1 + 3960C \log(2Kn))} + \frac{2y}{3} + 1 \right).$$

REMARKS

- (i) The class \mathcal{F} is contained in an $L_2(P)$ -ball of radius ε , and if interpreted as a loss class it is assumed that the losses are bounded.

(ii) Suppose the function class \mathcal{F} is constructed by selecting from a larger class an ε -ball in the $L_2(\mathbb{P})$ pseudometric around some function f_0 from the same larger class and taking for each function the absolute difference with f_0 . Then the zero function trivially is in \mathcal{F} since $|f_0 - f_0|$ is in the class. In this setup, the theorem states that with high probability there is no function in the class whose empirical risk will be “much” smaller/larger than the empirical risk of f_0 .

Proof of Theorem 11. For the proof, we introduce the random variables $Z = \sup_{f \in \mathcal{F}} P_n f$ and $\bar{Z} = \sup_{f \in \mathcal{F}} (P_n - P)f$. The proof is in three steps.

Step 1: Centering approximation

It is easy to see that $Z \leq \bar{Z} + \varepsilon$, since

$$\begin{aligned} Z &= \sup_{f \in \mathcal{F}} P_n f = \sup_{f \in \mathcal{F}} (P_n - P)f + P f \\ &\leq \sup_{f \in \mathcal{F}} (P_n - P)f + \sup_{f \in \mathcal{F}} P f \\ &\leq \bar{Z} + \sup_{f \in \mathcal{F}} \|f_0 - f\|_{L_1(\mathbb{P})} \\ &\leq \bar{Z} + \sup_{f \in \mathcal{F}} \|f_0 - f\|_{L_2(\mathbb{P})} \\ &\leq \bar{Z} + \varepsilon, \end{aligned}$$

where the penultimate inequality follows from Jensen’s inequality.

Step 2: Concentration of \bar{Z} arounds its expectation

We will apply Bousquet’s version of Talagrand’s inequality, appearing as equation (18) of [16] and reproduced below for convenience:

If \mathcal{G} is a countable family of measurable functions such that, for some positive constants v and b , one has, for every $g \in \mathcal{G}$, $\mathbb{P} g^2 \leq v$ and $\|g\|_\infty \leq b$, then, for every positive y , the following inequality holds for $W = \sup_{g \in \mathcal{G}} (P_n - P)g$:

$$\Pr \left\{ W - \mathbb{E} W \geq \sqrt{\frac{2(v + 4b \mathbb{E} W)y}{n}} + \frac{2by}{3n} \right\} \leq e^{-y}.$$

We take \mathcal{G} to be \mathcal{F} itself; since \mathcal{F} is separable and hence admits a countable dense subset, the countability assumption in Talagrand’s inequality is not an issue. Observe that for every $f \in \mathcal{F}$ we have

- $\|f\|_\infty \leq 1$ (from the range constraints on f)
- $\mathbb{P} f^2 \leq \|f\|_{L_2(\mathbb{P})} \leq \varepsilon$ (by the small $L_2(\mathbb{P})$ -ball assumption on \mathcal{F}).

Thus, taking $b = 1$ and $v = \varepsilon$, we have

$$\Pr \left\{ \bar{Z} - \mathbb{E} \bar{Z} \geq \sqrt{\frac{2(\varepsilon + 4 \mathbb{E} \bar{Z})y}{n}} + \frac{2y}{3n} \right\} \leq e^{-y}. \quad (20)$$

It remains to bound $\mathbb{E} \bar{Z}$. If it can be shown to be $\tilde{O}(\frac{1}{n})$ then we will have the desired result after taking $\varepsilon = O(\frac{1}{n})$.

Step 3: Controlling the size of $\mathbb{E} \bar{Z}$

Controlling $\mathbb{E} \bar{Z}$ can be done through chaining after passing to a symmetrized empirical process. This control is shown in Lemma 12, stated after the current proof, yielding the result

$$\mathbb{E} \bar{Z} \leq \frac{990C \log(2Kn)}{n}. \quad (21)$$

Putting it all together

The desired result follows by plugging (21) into the concentration result (20), incorporating the ε approximation term from Step 1, and setting $\varepsilon = \frac{1}{n}$. \square

Lemma 12. *Take the same conditions as Theorem 11 (Local Analysis Theorem), but instead allow that all $f \in \mathcal{F}$ need only satisfy $0 \leq f \leq V$ for some $V \geq 1$. Then provided that $n \geq 4$,*

$$\mathbb{E} \sup_{f \in \mathcal{F}} (P_n - P)f \leq \frac{990CV \log(2Kn)}{n}$$

Proof. To avoid measurability issues, we operate under the assumption that \mathcal{F} is countable and in the final step of the proof apply an approximation argument.

Let $\epsilon_1, \dots, \epsilon_n$ be iid Rademacher random variables. We write \mathbb{E}_ϵ for the expected value with respect to the random variables $\epsilon_1, \dots, \epsilon_n$. That is, if A is a random variable depending only on $\epsilon_1, \dots, \epsilon_n, X_1, \dots, X_n$, then $\mathbb{E}_\epsilon A = \mathbb{E}[A \mid X_1, \dots, X_n]$. Also, let $\mathcal{F}_{|\mathbf{x}}$ be the coordinate projection of \mathcal{F} onto the sample $\mathbf{X} = (X_1, \dots, X_n)$:

$$\mathcal{F}_{|\mathbf{x}} := \{(f(X_1), \dots, f(X_n)) : f \in \mathcal{F}\}.$$

Finally, for a set $\mathcal{G} \subset \mathbb{R}^n$ let $D(\mathcal{G})$ be half of the ℓ_2 -radius of \mathcal{G} , defined as

$$D(\mathcal{G}) := \frac{1}{2} \sup_{g \in \mathcal{G}} \|g\|_2;$$

it makes sense to refer to this as a (half) radius since we will consider $D(\mathcal{F})$ and the zero function is in \mathcal{F} . Our life will be made easier if we use the lower bounded quantity $D(\mathcal{G}) \vee \sigma$, for some deterministic $\sigma \leq 1$ to be chosen later.

The first step is symmetrization. The second step is based on a chaining argument, the result of which is Corollary 13.2 of Boucheron et al. [5], restated here³ we state the result in terms of in a specialization to Rademacher processes for convenience:

Let (\mathcal{T}, d) be a finite pseudometric space and $(X_t)_{t \in \mathcal{T}}$ be a collection of sub-Gaussian random variables. Then for any $t_0 \in \mathcal{T}$,

$$\mathbb{E} \sup_{t \in \mathcal{T}} X_t - X_{t_0} \leq 12 \int_0^{\delta/2} \sqrt{\log \mathcal{N}(u, \mathcal{T}, d)} du,$$

where $\delta = \sup_{t \in \mathcal{T}} d(t, t_0)$.

By pushing the cardinality of \mathcal{T} to infinity, the above result also applies to countable classes. As noted by Boucheron et al. [5] in the paragraph concluding the statement of their Corollary 13.2, this result applies to Rademacher processes. In our case, t_0 will correspond to the zero function element of \mathcal{F} .

Define $\bar{P}_n := P_n - P$. Now, from symmetrization and the above chaining-based result applied to the resulting Rademacher process, we have

$$\begin{aligned} n \mathbb{E} \sup_{f \in \mathcal{F}} \bar{P}_n f &\leq 2 \mathbb{E} \left(\mathbb{E}_\epsilon \sup_{f \in \mathcal{F}} \sum_{j=1}^n \epsilon_j f(X_j) \right) \\ &\leq 24 \mathbb{E} \int_0^{D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma} \sqrt{\log \mathcal{N}(u, \mathcal{F}_{|\mathbf{x}}, \|\cdot\|_2)} du, \end{aligned}$$

which (since if $f_{|\mathbf{x}}$ is the obvious coordinate projection of $f \in \mathcal{F}$, then $\|f_{|\mathbf{x}}\|_2 = \sqrt{n} \|f\|_{L_2(P_n)}$) is at most

$$\begin{aligned} &24 \mathbb{E} \int_0^{D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma} \sqrt{\mathcal{H}_2 \left(\frac{u}{\sqrt{n}}, \mathcal{F}_{|\mathbf{x}} \right)} du \\ &\leq 24\sqrt{C} \mathbb{E} \int_0^{D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma} \sqrt{\log \frac{K\sqrt{n}}{u}} du, \end{aligned}$$

³Boucheron et al. [5] stated this result in terms of packing numbers, but careful inspection of their proof reveals that the argument works for covering numbers as well. Moreover, other proofs generally use covering numbers.

where $\mathcal{H}_2(u, \mathcal{T}) := \sup_Q \log \mathcal{N}(u, \mathcal{T}, L_2(Q))$ is the *universal metric entropy* of \mathcal{T} , and in the above display we have $\mathcal{H}_2\left(\frac{u}{\sqrt{n}}, \mathcal{F}_{|\mathbf{x}}\right)$ rather than $\mathcal{H}_2(u, \mathcal{F}_{|\mathbf{x}})$ because we work with the $L_2(\mathbb{P}_n)$ -norm scaled by \sqrt{n} .

Making the substitution $t = u/(D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma)$, the above is equal to

$$\begin{aligned} & 24\sqrt{\mathcal{C}} \mathbb{E} \left((D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma) \int_0^1 \sqrt{\log \frac{K\sqrt{n}}{t(D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma)}} dt \right) \\ & \leq 24\sqrt{\mathcal{C}} \mathbb{E} \left((D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma) \left(\sqrt{\log \frac{K\sqrt{n}}{D(\mathcal{F}_{|\mathbf{x}}) \vee \sigma}} + \frac{\sqrt{\pi}}{2} \right) \right) \\ & \leq 24\sqrt{\mathcal{C}} \left(\sqrt{\log \frac{K\sqrt{n}}{\sigma}} + \frac{\sqrt{\pi}}{2} \right) (\sigma + \mathbb{E} D(\mathcal{F}_{|\mathbf{x}})). \end{aligned}$$

Now, we focus on $\mathbb{E} D(\mathcal{F}_{|\mathbf{x}})$. Observe that

$$\begin{aligned} \mathbb{E} D(\mathcal{F}_{|\mathbf{x}}) &= \frac{\sqrt{n}}{2} \mathbb{E} \sup_{f \in \mathcal{F}} \sqrt{\mathbb{P}_n f^2} \\ &= \frac{\sqrt{n}}{2} \mathbb{E} \sup_{f \in \mathcal{F}} \sqrt{\bar{\mathbb{P}}_n f^2 + \mathbb{P} f^2} \\ &\leq \frac{\sqrt{n}}{2} \left[\mathbb{E} \sup_{f \in \mathcal{F}} \sqrt{\bar{\mathbb{P}}_n f^2} + \sup_{f \in \mathcal{F}} \sqrt{\mathbb{P} f^2} \right] \\ &\leq \frac{\sqrt{n}}{2} \left[\sqrt{\mathbb{E} \sup_{f \in \mathcal{F}} \bar{\mathbb{P}}_n f^2} + \varepsilon \right]. \end{aligned}$$

where the first part of the last step follows from Jensen's inequality and the second part follows from the small $L_2(\mathbb{P})$ -ball assumption on \mathcal{F} . The above is at most

$$\frac{\sqrt{Vn}}{2} \left[\sqrt{\mathbb{E} \sup_{f \in \mathcal{F}} \bar{\mathbb{P}}_n f^2} + \varepsilon \right].$$

Thus, putting everything together and making the replacement $\varepsilon = \frac{1}{n}$, we have

$$\mathbb{E} \sup_{f \in \mathcal{F}} \bar{\mathbb{P}}_n f \leq \frac{24\sqrt{\mathcal{C}}}{n} \left(\sqrt{\log \frac{K\sqrt{n}}{\sigma}} + \frac{\sqrt{\pi}}{2} \right) \left(\sigma + \frac{\sqrt{Vn}}{2} \left[\sqrt{\mathbb{E} \sup_{f \in \mathcal{F}} \bar{\mathbb{P}}_n f^2} + \frac{1}{n} \right] \right).$$

Finding the minimal value of $\mathbb{E} \sup_{f \in \mathcal{F}} \bar{\mathbb{P}}_n f$ just amounts to solving a quadratic equation, yielding the solution set

$$\sqrt{\mathbb{E} \sup_{f \in \mathcal{F}} \bar{\mathbb{P}}_n f^2} \leq \frac{\psi\sqrt{Vn}}{2} + \sqrt{\psi \left(\sigma + \frac{\sqrt{V/n}}{2} \right)}$$

for $\psi = \frac{24\sqrt{\mathcal{C}}}{n} \left(\sqrt{\log \frac{K\sqrt{n}}{\sigma}} + \frac{\sqrt{\pi}}{2} \right)$.

Making the replacement $\sigma = \frac{1}{n}$, squaring, and some coarse bounding yields

$$\mathbb{E} \sup_{f \in \mathcal{F}} \bar{\mathbb{P}}_n f \leq \frac{990\mathcal{C}V \log(Kn)}{n}$$

for $n \geq 4$, $V \geq 1$, and $\mathcal{C} \geq 1$.

We now present the approximation argument to handle separable \mathcal{F} . Since \mathcal{F} is separable, it suffices to consider a countable dense subset $\mathcal{F}' \subset \mathcal{F}$; however, a little more work is required as the covering

numbers of \mathcal{F}' may differ slightly from the covering numbers of \mathcal{F} . We now control the covering numbers of \mathcal{F}' in terms of the covering numbers of \mathcal{F} . Observe that if there is an ε -net of \mathcal{F} of cardinality N , then there is a (2ε) -net of some $\mathcal{F}' \subset \mathcal{F}$ of cardinality N . Hence, if there is an optimal ε -net of \mathcal{F} of cardinality N , then an optimal (2ε) -net of \mathcal{F}' has cardinality at most N . That is, for any probability measure Q on \mathcal{X} , for any $u > 0$ we have $\mathcal{N}(2u, \mathcal{F}', L_2(Q)) \leq \mathcal{N}(u, \mathcal{F}, L_2(Q))$. Thus, the result for separable \mathcal{F} holds by replacing the constant K with $2K$. \square

We now prove the localization result.

Proof of Theorem 6. First, so that we can just handle the case of functions with range $[0, 1]$, we (crudely) apply our analysis to the function class after scaling all functions by the factor $\frac{1}{V}$, and scale the approximation term in the last step by the factor V .⁴

For any $f_0 \in \mathcal{F}_\varepsilon$, observe that $\pi^{-1}(f_0)$ is the set of those functions that are covered by f_0 in the $L_2(\mathbb{P})$ -norm. We apply the Local Analysis Theorem (Theorem 11) to each element of the set of localized *absolute difference* function classes

$$\{\mathcal{G}_{f_0} : f_0 \in \mathcal{F}_\varepsilon\},$$

for $\mathcal{G}_{f_0} := \{|f_0 - f| : f \in \pi^{-1}(f_0)\}$. Consider an arbitrary $f_0 \in \mathcal{F}_\varepsilon$ and its corresponding class \mathcal{G}_{f_0} . Since $\tilde{\mathcal{G}}_{f_0} := \{f_0 - f : f \in \pi^{-1}(f_0)\}$ is isomorphic to a subset of \mathcal{F} , and since any ε -net for $\tilde{\mathcal{G}}_{f_0}$ trivially gives rise to an ε -net for \mathcal{G}_{f_0} by taking the absolute value of each function from the original ε -net, it follows the $L_2(\mathbb{P})$ covering numbers of \mathcal{G}_{f_0} are bounded just as in (11).

Taking the union bound over \mathcal{F}_ε with Theorem 11 implies that with probability at least $1 - \delta$

$$\begin{aligned} & \max_{f_0 \in \mathcal{F}_\varepsilon} \sup_{f \in \pi^{-1}(f_0)} \mathbb{P}_n |f_0 - f| \\ & \leq \frac{1}{n} \left(990\mathcal{C} \log(2Kn) + \sqrt{2 \left(\log \frac{1}{\delta} + \mathcal{C} \log(Kn) \right) (1 + 3960\mathcal{C} \log(2Kn))} \right) \\ & \quad + \frac{1}{n} \left(\frac{2 \left(\log \frac{1}{\delta} + \mathcal{C} \log(Kn) \right)}{3} + 1 \right). \end{aligned}$$

Ignoring the $\frac{1}{n}$ factor, the RHS is at most

$$990\mathcal{C} \log(2Kn) + \sqrt{2 \left(\log \frac{1}{\delta} + \mathcal{C} \log(Kn) \right) (1 + 3960\mathcal{C} \log(2Kn))} + \log \frac{e}{\delta} + \mathcal{C} \log(Kn),$$

which is at most

$$\begin{aligned} & 991\mathcal{C} \log(2Kn) \\ & + \sqrt{2 \log \frac{1}{\delta} + 2\mathcal{C} \log(Kn) + 7920 \left(\log \frac{1}{\delta} \right) \mathcal{C} \log(2Kn) + 7920(\mathcal{C} \log(2Kn))^2} + \log \frac{e}{\delta} \\ & \leq 1080\mathcal{C} \log(2Kn) + \sqrt{2 \log \frac{1}{\delta} + 2\mathcal{C} \log(Kn) + 7920 \left(\log \frac{1}{\delta} \right) \mathcal{C} \log(2Kn) + \log \frac{e}{\delta}} \\ & \leq 1080\mathcal{C} \log(2Kn) + 90 \sqrt{\left(\log \frac{1}{\delta} \right) \mathcal{C} \log(2Kn) + \log \frac{e}{\delta}}, \end{aligned}$$

where the last inequality holds provided that δ is not too large; it suffices to assume $\delta \leq \frac{1}{2}$. \square

Finally, we prove the fast rates exact oracle inequality for VC-type classes.

⁴It may be possible to get a weaker dependence on V with a more careful argument that depends on V throughout; in particular, Talagrand's inequality can handle the parameter V .

C.2 Proof of Theorem 7

Proof of Theorem 7. For convenience, we begin by abusing notation and redefining \mathcal{F} as $\mathcal{F} := \ell \circ \mathcal{F}$; the abuse includes f^* being redefined as $\ell(\cdot, f^*)$. With these abuses, for any $f \in \mathcal{F}$ we redefine Z_f as $Z_f := f - f^*$.

Next, we introduce a few subclasses that we will use. Recall that for any $\gamma_n > 0$, $\mathcal{F}_{\geq \gamma_n}$ is the subclass of \mathcal{F} for which the excess risk is at least γ_n . Also, for any $\gamma_n > 0$, let $\mathcal{F}_{\geq \gamma_n, \varepsilon}$ be a proper cover of $\mathcal{F}_{\geq \gamma_n}$ with respect to the $L_2(\mathbb{P}_n)$ norm, with $\varepsilon = \frac{1}{n}$. For each $\eta > 0$ and $\mathcal{F}_{\geq \gamma_n, \varepsilon}$, let $\mathcal{F}_{\geq \gamma_n, \varepsilon}^{(\eta)} \subset \mathcal{F}_{\geq \gamma_n, \varepsilon}$ correspond to those functions for which η is the largest constant such that $\mathbb{E} \exp(-\eta Z_f) = 1$. After making the same implicit change to $\mathcal{F}_{\geq \gamma_n, \varepsilon}$ for “hyper-concentrated” excess loss random variables (i.e. those Z_f for which $\lim_{\eta \rightarrow \infty} \mathbb{E} \exp(-\eta Z_f) < 1$) as was made to $\mathcal{F}_{\geq \gamma_n}$ in the proof of Theorem 5, we have $\mathcal{F}_{\geq \gamma_n, \varepsilon} = \bigcup_{\eta \in [\eta^*, \infty)} \mathcal{F}_{\geq \gamma_n, \varepsilon}^{(\eta)}$.

Let $\gamma_n = \frac{a}{n}$ for some constant $a > 1$ to be fixed later. Consider an arbitrary $\eta \in [\eta^*, \infty)$ for which $|\mathcal{F}_{\geq \gamma_n, \varepsilon}^{(\eta)}| > 0$, and recall that all functions f in this class satisfy $\mathbb{E} Z_f \geq \frac{a}{n}$. Individually for each such function f , we will apply the Cramér-Chernoff Theorem (Theorem 1) as follows. From the Bounded Losses Lemma (Lemma 4), we have $\Lambda_{-Z_f}(\eta/2) = \Lambda_{-\frac{1}{\eta} Z_f}(V\eta/2)$. From the Stochastic Mixability Concentration Theorem (Theorem 3), the latter is at most

$$-\frac{0.18(V\eta \wedge 1)(a/V)}{n} = -\frac{0.18\eta a}{(V\eta \vee 1)n}.$$

Hence, the Cramér-Chernoff Theorem (Theorem 1) with $t = \frac{a}{2n}$ and the η from that Theorem taken to be $\eta/2$ implies that the probability of the event $\mathbb{P}_n f \leq \mathbb{P}_n f^* + \frac{a}{2n}$ is at most

$$\exp\left(-0.18 \frac{\eta}{V\eta \vee 1} a + \frac{a\eta}{4n}\right) = \exp\left(-\eta a \left(\frac{0.18}{V\eta \vee 1} - \frac{1}{4n}\right)\right).$$

Applying the union bound over all of $\mathcal{F}_{\geq \gamma_n, \varepsilon}$, we conclude that

$$\Pr\left\{\exists f \in \mathcal{F}_{\geq \gamma_n, \varepsilon} : \mathbb{P}_n f \leq \mathbb{P}_n f^* + \frac{a}{2n}\right\} \leq \left(\frac{K}{\varepsilon}\right)^{\mathcal{C}} \exp\left(-\eta^* a \left(\frac{0.18}{V\eta^* \vee 1} - \frac{1}{4n}\right)\right).$$

Now, observe that if we consider some fixed failure probability $\frac{\delta}{2}$ and invert to obtain the corresponding a , we have

$$\begin{aligned} a &= \frac{\mathcal{C} \log \frac{K}{\varepsilon} + \log \frac{2}{\delta}}{\eta^* \left(\frac{0.18}{V\eta^* \vee 1} - \frac{1}{4n}\right)} = \frac{\mathcal{C} \log \frac{K}{\varepsilon} + \log \frac{2}{\delta}}{\eta^* \left(\frac{0.18 - (V\eta^* \vee 1)/(4n)}{V\eta^* \vee 1}\right)} \\ &\leq \frac{(V\eta^* \vee 1) (\mathcal{C} \log \frac{K}{\varepsilon} + \log \frac{2}{\delta})}{\eta^* (0.18 - \frac{1}{4n})} \\ &\leq 8 \left(V \vee \left(\frac{1}{\eta^*}\right)\right) \left(\mathcal{C} \log \frac{K}{\varepsilon} + \log \frac{2}{\delta}\right) =: \lambda, \end{aligned} \quad (22)$$

for $\gamma_n^{(1)} := \frac{\lambda}{n}$, where the last inequality holds since $n \geq 5$. Note that by instead setting $\frac{a}{n} = \gamma_n^{(1)}$ (defined in (22)) the failure probability can only decrease. Thus, for any $\gamma_n \geq \gamma_n^{(1)}$, we have

$$\Pr\left\{\exists f \in \mathcal{F}_{\geq \gamma_n, \varepsilon} : \mathbb{P}_n f \leq \mathbb{P}_n f^* + \frac{\gamma_n}{2}\right\} \leq \frac{\delta}{2}.$$

Next, we control the behavior of the subclass $\mathcal{F}_{\geq \gamma_n} \setminus \mathcal{F}_{\geq \gamma_n, \varepsilon}$. From Theorem 6, if $\delta \leq \frac{1}{2}$

$$\Pr\left\{\exists f \in \mathcal{F}_{\geq \gamma_n} : \mathbb{P}_n f < \mathbb{P}_n \pi(f) - \gamma_n^{(2)}\right\} \leq \frac{\delta}{2}.$$

for $\gamma_n^{(2)} = \frac{V}{n} \left(1080\mathcal{C} \log(2Kn) + 90\sqrt{(\log \frac{2}{\delta}) \mathcal{C} \log(2Kn)} + \log \frac{2\varepsilon}{\delta}\right)$.

Now, combining the above two high probability guarantees, with probability at least $1 - \delta$ both statements below hold for all $f \in \mathcal{F}_{\geq \gamma_n}$:

$$\begin{aligned} P_n f &\geq P_n \pi(f) - \gamma_n^{(2)} \\ P_n \pi(f) &\geq P_n f^* + \frac{\gamma_n}{2}. \end{aligned}$$

Thus, with the same probability, for all $f \in \mathcal{F}_{\geq \gamma_n}$:

$$P_n f \geq P_n f^* + \frac{\gamma_n}{2} - \gamma_n^{(2)}.$$

Setting $\gamma_n = (\gamma_n^{(1)} \vee 2\gamma_n^{(2)}) + \frac{1}{n}$, and recalling that ERM selects hypotheses purely based on their empirical risk, we see that with probability at least $1 - \delta$, ERM will not select any hypothesis whose excess risk is at least

$$(\gamma_n^{(1)} \vee (2\gamma_n^{(2)})) + \frac{1}{n}. \quad \square$$