
Flexible risk design using bi-directional dispersion

Matthew J. Holland
Osaka University

Abstract

Many novel notions of “risk” (e.g., CVaR, tilted risk, DRO risk) have been proposed and studied, but these risks are all at least as sensitive as the mean to loss tails on the upside, and tend to ignore deviations on the downside. We study a complementary new risk class that penalizes loss deviations in a bi-directional manner, while having more flexibility in terms of tail sensitivity than is offered by mean-variance. This class lets us derive high-probability learning guarantees without explicit gradient clipping, and empirical tests using both simulated and real data illustrate a high degree of control over key properties of the test loss distribution of gradient-based learners.

1 INTRODUCTION

What does it mean for a learner to successfully generalize? Broadly speaking, this is an ambiguous property of learning systems that can be defined, measured, and construed in countless ways. In the context of machine learning, however, the notion of “success” in off-sample generalization is almost without exception formalized as minimizing the expected value of a random loss $\mathbf{E}_\mu L(h)$, where h is a candidate parameter, model, or decision rule, and $L(h)$ is a random variable on a probability space $(\Omega, \mathcal{F}, \mu)$ (Mohri et al., 2012; Shalev-Shwartz and Ben-David, 2014). The idea of quantifying the *risk* of an unexpected outcome (here, a random loss) using the expected value dates back to the Bernoullis and Gabriel Cramer in the early 18th century (Bassett, 1987; Hacking, 2006). In a more modern context, the emphasis on *average* performance is the “general setting of the learning problem” of Vapnik (1999), and plays a central role in the decision-theoretic learning model of Haussler (1992). Use of the expected loss to quantify off-sample generalization has been essential to the development

of both the statistical and computational theories of learning (Devroye et al., 1996; Kearns and Vazirani, 1994).

While the expected loss still remains pervasive, important new lines of work on *risk-sensitive learning* have begun exploring novel feedback mechanisms for learning algorithms, in some cases derived directly from new risk functions that replace the expected loss. Learning algorithms designed using conditional value-at-risk (CVaR) (Curi et al., 2020) and tilted (or “entropic”) risk (Föllmer and Knispel, 2011; Li et al., 2021a,b) are well-known examples of location properties which emphasize loss tails in *one direction* more than the mean itself does. This is often used to increase sensitivity to “worst-case” events (Kashima, 2007; Takeda and Sugiyama, 2008), but in special cases where losses are bounded below, sensitivity to tails on the downside can be used to realize an insensitivity to tails on the upside (Lee et al., 2020). This strong asymmetry is not specific to the preceding two risk function classes, but rather is inherent in much broader classes such as optimized certainty equivalent (OCE) risk (Ben-Tal and Teboulle, 1986, 2007; Lee et al., 2020) and distributionally robust optimization (DRO) risk (Ben-Tal et al., 2013; Duchi and Namkoong, 2018, 2019; Gotoh et al., 2018). Unsurprisingly, naive empirical estimators of these risks are particularly fragile under outliers coming from the “sensitive direction,” as is evidenced by the plethora of attempts in the literature to design robust modifications (Holland and Haress, 2021; Prashanth et al., 2020; Zhai et al., 2021). In general, however, loss distributions can display long tails in either direction over the learning process (see Figure 3), particularly when losses are unbounded below (e.g., negative rewards (Sutton and Barto, 2018), unhinged loss (Van Rooyen et al., 2016)), and loss functions whose empirical mean has no minimum appear frequently (e.g., separable logistic regression (Albert and Anderson, 1984; Rousseeuw and Christmann, 2003)). Since the tail behavior of stochastic losses and gradients is well-known to play a critical role in the stability and robustness of learning systems (Şimşekli et al., 2019; Zhang et al., 2020), the inability to control tail sensitivity in both directions represents a genuine limitation to machine learning methodology.

A natural alternative class of risk functions that gives us control over tail sensitivity in *both* directions is that of the

“M-location” of the loss distribution, namely any value in

$$\arg \min_{\theta \in \mathbb{R}} \mathbf{E}_\mu \rho(L(h) - \theta) \subset \mathbb{R}, \quad (1)$$

where $\rho : \mathbb{R} \rightarrow \mathbb{R}_+$ is assumed to be such that this set of minimizers is non-empty. Here various special choices of ρ let us recover well-known locations, such as the mean (with $\rho(\cdot) = (\cdot)^2$), median ($\rho(\cdot) = |\cdot|$), arbitrary quantiles (via “pinball” function (Takeuchi et al., 2006)), and even further beyond to “expectiles” (using curved variants of the pinball function (Gneiting, 2011)). The obvious limitation here is that while computing (1) using empirical estimates is easy, minimization as a function of h is in general a difficult bi-level programming problem. As an alternative approach, in this paper we study the potential benefits and tradeoffs that arise in using performance criteria of the form

$$\min_{\theta \in \mathbb{R}} [\eta\theta + \mathbf{E}_\mu \rho(L(h) - \theta)] \quad (2)$$

where $\eta \in \mathbb{R}$. By making a sacrifice of fidelity to the M-location (1), we see that the criterion in (2) suggests a congenial objective function (joint in (h, θ)). Intuitively, one minimizes the sum of generalized “location” and “dispersion” properties, and the nature of this dispersion impacts the fidelity of the location term to the original M-location induced by ρ . These two locations align perfectly in the special case where we set $\rho(\cdot) = (\cdot)^2/2$ and $\eta = 1$, since (2) is equivalent to the mean-variance objective $\mathbf{E}_\mu L(h) + \text{var}_\mu L(h)/2$, but more generally, it is clear that allowing for more diverse choices of ρ gives us new freedom in terms of tail control with respect to both location and dispersion. We consider a concrete yet flexible class of risk functions that generalizes beyond (2), allows for easy implementation, and is analytically tractable from the standpoint of providing formal learning guarantees. Our main contributions are as follows:

- A new class of “threshold risks” (T-risks, §3) that provide a tractable alternative to M-locations, and a bi-directional complement to OCE/DRO risks (cf. §2).
- A stochastic learning algorithm for T-risks that enjoys high-probability guarantees of convergence to a stationary point under heavy-tailed losses/gradients, without manual clipping (details in §4, Theorem 3).
- Strong empirical evidence of the flexibility and utility inherent in T-risk learners. In particular: robustness to unbalanced noisy class labels without regularization (§5.1), sharp control over sensitivity to outliers in regression with convex base losses (§C.1), and smooth interpolation between mean and mean-variance minimizers on clean, normalized benchmark classification datasets (§C.2).

The overall flow of the paper is as follows. Background information on notation and related literature is given in §2,

and we introduce the new risk class of interest in §3. Formal aspects of the learning problem and algorithms centered around these risks are treated in §4. Empirical findings for the learning problem are explained and discussed in §C of the supplementary materials. Detailed notation, formal proofs and other supplementary results are organized in §A–§F of the appendix, and code for reproducing all the results in this paper is provided in an online repository.¹

2 BACKGROUND

OCE-type risks As a computationally convenient way to interpolate between the mean and the extreme values of $L(h)$, the *tilted risk* (Li et al., 2021a,b) is a natural choice, defined for $\gamma \neq 0$ as

$$R_{\text{tilt}}(h; \gamma) := \frac{1}{\gamma} \log \left(\mathbf{E}_\mu e^{\gamma L(h)} \right). \quad (3)$$

This is simply a re-scaling of the cumulant distribution function of $L(h)$, viewed as a function of h , where taking $\gamma \rightarrow \infty$ and $\gamma \rightarrow -\infty$ lets us approach the supremum and infimum of $L(h)$, respectively. Another important class of risk functions is based upon the *conditional value-at-risk* (CVaR) (Rockafellar and Uryasev, 2000), defined for $0 \leq \beta < 1$ as

$$R_{\text{CVaR}}(h; \beta) := \mathbf{E}_\mu [L(h) \mid L(h) \geq Q_\beta(h)]. \quad (4)$$

This is the expected loss at h , conditioned on the event that the loss exceeds the β -quantile of $L(h)$, denoted here by $Q_\beta(h) := \inf \{x \in \mathbb{R} : \mathbf{P}\{L(h) \leq x\} \geq \beta\}$. Both of these risk functions can be re-written in a form similar to that of (2), namely

$$h \mapsto \inf_{\theta \in \mathbb{R}} [\theta + \mathbf{E}_\mu \phi(L(h) - \theta)] \quad (5)$$

where $\phi(x) = (e^{\gamma x} - 1)/\gamma$ yields R_{tilt} (basic calculus), and $\phi(x) = \max\{0, x\}/(1 - \beta)$ yields R_{CVaR} (see Rockafellar and Uryasev (2000, 2002)). When $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is restricted to be a non-decreasing, closed, convex function which satisfies both $\phi(0) = 0$ and $1 \in \partial\phi(0)$, the mapping given in (5) is called an *optimized certainty equivalent (OCE) risk* (Ben-Tal and Teboulle, 1986, 2007; Lee et al., 2020). The class of OCE risks strictly generalizes the expected value (noting $\phi(x) = x$ is valid), and includes $R_{\text{tilt}}(\cdot; \gamma)$ when $\gamma > 0$, as well as R_{CVaR} . The mean-variance is sometimes stated to be an OCE risk (Lee et al., 2020, Table 1), but this fails to hold when losses are unbounded above and below.

DRO-type risks Another important class of risk functions is that of *robustly regularized* risks which are designed to ensure that risk minimizers are robust to a certain degree of divergence in the underlying data model. Making this concrete, it is typical to assume the random losses are the

¹<https://github.com/feedbackward/bdd>

outputs of a loss function ℓ depending on the candidate h and some random data Z , i.e., $L(h) = \ell(h; Z)$, with $Z \sim \mu$ as our reference model. To measure divergence from this reference model, it is convenient to use the Cressie-Read family of divergence functions (Duchi and Namkoong, 2018; Zhai et al., 2021), namely for any $c > 1$ and assuming $\nu \ll \mu$ (absolute continuity) holds, functions of the form

$$\text{Div}_c(\nu; \mu) := \mathbf{E}_\mu f_c \left(\frac{d\nu}{d\mu} \right), \quad (6)$$

where $f_c(x) := (x^c - cx + c - 1)/c(c - 1)$ and $d\nu/d\mu$ is the Radon-Nikodym density of ν with respect to μ .² The resulting risk, called the *DRO risk*, is defined as

$$R_{\text{DRO}}(h) := \sup \{ \mathbf{E} L(h) : L \in \mathcal{L} \} \quad (7)$$

where the constrained set of random losses \mathcal{L} , determined by $c > 1$ and $a > 0$, is defined as

$$\mathcal{L} := \{ L(\cdot) = \ell(\cdot; Z) : Z \sim \nu \text{ and } \text{Div}_c(\nu; \mu) \leq a \}. \quad (8)$$

For this particular family of divergences, the risk can be characterized as the optimal value of a simple optimization problem (Duchi and Namkoong, 2018), namely we have that $R_{\text{DRO}}(h)$ is equal to

$$\inf_{\theta \in \mathbb{R}} \left[\theta + (1 + c(c - 1)a)^{1/c} (\mathbf{E}_\mu (L(h) - \theta)_+^{c_*})^{1/c_*} \right] \quad (9)$$

where $c_* := c/(c - 1)$. While strictly speaking this is not an OCE risk, setting $\phi(x) = (1 + c(c - 1)a)^{c_*} x_+^{c_*}$, the DRO risk can then be written as

$$R_{\text{DRO}}(h) = \inf_{\theta \in \mathbb{R}} \left[\theta + [\mathbf{E}_\mu \phi(L(h) - \theta)]^{1/c_*} \right], \quad (10)$$

giving us an expression of this risk as the sum of a threshold and an asymmetric dispersion. When we set $c = 2$, this yields the well-known special case of χ^2 -DRO risk (Hashimoto et al., 2018; Zhai et al., 2021). In addition to the one-directional nature of the dispersion term in these risks, all of these risks are at least as sensitive to loss tails (on the upside) as the classical expected loss $\mathbf{E}_\mu L(h)$ is; this holds for CVaR (with $\beta > 0$), tilted risk (with $\gamma > 0$), and even robust variants of DRO risk (Zhai et al., 2021).

Key differences While it is clear that the form of the preceding risk classes given in (5) and (10) based on various choices of $\phi(\cdot)$ is the same as our $\rho(\cdot)$ -based risk of interest in (2), they are fundamentally different in that *none* of the choices of $\phi(\cdot)$ induce a meaningful M-location; since all these $\phi(\cdot)$ are monotonic on \mathbb{R} , both minimization and maximization of $\theta \mapsto \mathbf{E}_\mu \phi(L(h) - \theta)$ is trivially accomplished by taking $|\theta| \rightarrow \infty$. In stark contrast, $\rho(\cdot)$ is assumed to be such that the solution set in (1) is a subset of the real line. We will introduce a concrete and flexible class from which ρ will be taken in §3, and in Figure 1 give a side-by-side comparison with the ϕ functions discussed in the preceding paragraphs.

²For background on absolute continuity and density functions, see Ash and Doléans-Dade (2000, §2.2).

Closely related work Of all the papers cited above, the works of Lee et al. (2020) on OCE risks, and Li et al. (2021b,a) on tilted risks are of a similar nature to our paper, with the obvious difference being that the class of risks is fundamentally different, as described in the preceding paragraphs. Indeed, many of our empirical tests involve direct comparison with the risk classes studied in these works (e.g., Figures 2, 4, and 12), and so they provide critical context for our work. Previous work by Holland (2022) studies a rudimentary special case of what we call “minimal T-risk” here; the focus in that work was on obtaining learning guarantees (in expectation) when the risk is potentially non-convex and non-smooth in h , but with convex ρ , and no comparison was made with OCE/DRO risk classes. We build upon these results here, considering a broad class of dispersions ρ which are differentiable but need not be convex (see (15)); we show such risk classes can readily admit high-probability learning guarantees for stochastic gradient-based algorithms (Theorem 3), provide bounds on the average loss incurred by empirical risk minimizers using our risk (Proposition 7), and make detailed empirical comparisons with each of the key existing risk classes.

3 THRESHOLD RISK

To ground ourselves conceptually, let us refer to $L(h)$ as the *base loss* incurred by h . The exact nature of h is left completely abstract for the moment, as all that matters is the probability distribution of this base loss. By selecting an arbitrary *threshold* $\theta \in \mathbb{R}$, we define a broad class of properties as

$$R_\rho(h; \theta, \eta) := \eta\theta + \mathbf{E}_\mu \rho(L(h) - \theta). \quad (11)$$

Here $\eta \in \mathbb{R}$ is a weighting parameter allowed to be negative, and as a bare minimum, ρ is assumed to be such that the resulting M-location(s) are well-defined in the sense that the inclusion in (1) holds. We call $\rho(L(h) - \theta)$ the (random) *dispersion* of the base loss, taken with respect to threshold θ , and we refer to $R_\rho(h; \theta, \eta)$ in (11) as the *threshold risk* (or simply *T-risk*) under ρ .

3.1 Minimal T-risk and M-location

Arguably the most intuitive special case of T-risk is the *minimal T-risk*, where minimization is with respect to the threshold $\theta \in \mathbb{R}$. Let us denote this risk and the *optimal threshold set* as

$$\begin{aligned} \underline{R}_\rho(h; \eta) &:= \inf_{\theta \in \mathbb{R}} R_\rho(h; \theta, \eta) \\ \theta_\rho(h; \eta) &:= \arg \min_{\theta \in \mathbb{R}} R_\rho(h; \theta, \eta). \end{aligned} \quad (12)$$

Clearly, if ρ is bounded above or grows too slowly, we will have $\underline{R}_\rho(h; \eta) = -\infty$ and no real-valued minimizers, i.e., $\theta_\rho(h; \eta) = \emptyset$. Letting $M_\rho(h)$ denote the M-locations in (1),

for $\eta \neq 0$ we have

$$\theta_\rho(h; \eta) \neq \emptyset \implies M_\rho(h) \neq \emptyset, \quad (13)$$

although the converse does not hold in general.³ When $\eta = 0$, these two solution sets align, i.e., we have $\theta_\rho(h; 0) = M_\rho(h)$. More generally, depending on the sign of η , the optimal thresholds can be either larger or smaller than the corresponding M-locations. More precisely, for any $\theta' \in M_\rho(h)$, as long as $\theta_\rho(h; \eta)$ is non-empty, there exists $\theta \in \theta_\rho(h; \eta)$ such that $\theta \text{sign}(\eta) \leq \theta' \text{sign}(\eta)$.

Special case minimized by quantiles The form given in (11) is very general, but it can be understood as a straightforward generalization of the convex objective used to characterize quantiles. More precisely, taking $\beta \in (0, 1)$ and denoting the β -quantile of the base loss using $Q_\beta(h) := \inf \{x \in \mathbb{R} : \mathbf{P}\{\mathbf{L}(h) \leq x\} \geq \beta\}$, it is well-known that in the special case of $\rho(\cdot) = |\cdot|$, we have

$$Q_\beta(h) \in \theta_\rho(h; \theta, 1 - 2\beta) \quad (14)$$

for any choice of $0 < \beta < 1$, as long as $\mathbf{E}_\mu|\mathbf{L}(h)|$ is finite (Koltchinskii, 1997). The T-risk in (11) simply allows for a more flexible choice of ρ , and thus generalizes the dispersion term in this objective function.

3.2 T-risk with scaled Barron dispersion

In order to capture a range of sensitivities to loss tails in both directions, we would like to select ρ from a class of functions that gives us sufficient control over scale, boundedness, and growth rates. As a concrete choice, we propose to set ρ in (11) as $\rho(x) = \rho(x/\sigma; \alpha)$, where $\sigma > 0$ is a scaling parameter, and $\rho(\cdot; \alpha)$ with shape $\alpha \in [-\infty, 2]$ is a family of functions that ranges from bounded and logarithmic growth on the lower end to quadratic growth on the upper end, defined as:

$$\rho(x; \alpha) := \begin{cases} x^2/2, & \text{if } \alpha = 2 \\ \log(1 + x^2/2), & \text{if } \alpha = 0 \\ 1 - \exp(-x^2/2), & \text{if } \alpha = -\infty \\ \frac{|\alpha-2|}{\alpha} \left(\left(1 + \frac{x^2}{|\alpha-2|}\right)^{\alpha/2} - 1 \right), & \text{otherwise.} \end{cases} \quad (15)$$

At a high level, $\rho(\cdot; \alpha)$ is approximately quadratic near zero for any choice of shape α , but its growth as one deviates far from zero depends greatly on α . We refer to (15) as the *Barron class* of functions for computing dispersion.⁴

³For example, consider choices of ρ that are “re-descending” (Huber and Ronchetti, 2009).

⁴The reason for this naming is that Barron (2019) recently studied this class in the context of designing *loss functions* for computer vision applications. We remark that this differs considerably from our usage in computing the *dispersion of random losses*, where the loss function underlying the base loss is left completely arbitrary.

Recalling the risks reviewed in §2, since $\rho(\cdot; \alpha)$ is flat at zero and symmetric about zero, the Barron class clearly takes us beyond the functions $\phi(\cdot)$ allowed by OCE risks (5) and used in typical DRO risk definitions (10); see Figure 1 for a visual comparison.

As mentioned in §B, we will often use the generic shorthand notation $\rho_\sigma(x) := \rho(x/\sigma)$, and drop the dependence on α when clear from context. The shape parameter α gives us direct control over the conditions needed for a finite T-risk $R_\rho(h; \theta, \eta)$, as the following lemma shows.

Lemma 1 (Finiteness and shape). *Let ρ be from the Barron class (15). Then in order to ensure $\mathbf{E}_\mu \rho_\sigma(\mathbf{L}(h) - \theta) < \infty$ holds for all $\theta \in \mathbb{R}$, each of the following conditions (depending on the value of α) is sufficient. For $0 < \alpha \leq 2$, let $\mathbf{E}_\mu |\mathbf{L}(h)|^\alpha < \infty$. For $\alpha = 0$, let $\mathbf{E}_\mu |\mathbf{L}(h)|^c < \infty$ for some $c > 0$. For $-\infty \leq \alpha < 0$, let $\mathbf{L}(h)$ be \mathcal{F} -measurable. Furthermore, for the cases where $\alpha \neq 0$, the above conditions are also necessary.*

Assuming μ -integrability as in Lemma 1, the Barron class furnishes a non-empty set of M-locations $M_\rho(h)$ for any choice of α , and when restricted to $\alpha \geq 1$ with appropriate settings of η and σ , the optimal threshold set $\theta_\rho(h; \eta)$ contains a single unique solution (see Lemma 10). For any valid choice of α , the function $\rho(\cdot; \alpha)$ is twice continuously differentiable on \mathbb{R} (see §F.3 for exact expressions). All the limits in α behave as we would expect: $\rho(x; \alpha) \rightarrow \rho(x; c)$ as $\alpha \rightarrow c$ for $c \in \{-\infty, 0, 2\}$ (see §D.2 for details). For $\alpha \geq 0$, the dispersion function is unbounded, with growth ranging from logarithmic to quadratic depending on the choice of α . For $\alpha < 0$, the dispersion function is bounded. The mapping $x \mapsto \rho_\sigma(x; \alpha)$ is convex on \mathbb{R} for $\alpha \geq 1$, and for $\alpha < 1$ it is only convex between $\pm\sigma\sqrt{|\alpha-2|/(1-\alpha)}$, and concave elsewhere (see Lemma 8). The class $R_\rho(h; \theta, \eta)$ of T-risks (11) under the scaled Barron dispersion $\rho_\sigma(x; \alpha)$ is the central focus of this paper.

3.3 Sensitivity to outliers and tail direction

Before we consider the learning problem, which typically involves the evaluation of many different loss distributions over the course of training, here we consider a fixed distribution, and numerically compare the T-risk (11) and M-location (1) induced by ρ from the Barron class in §3.2, along with the key OCE and DRO risks discussed in §2.

Experiment setup We generate random values to simulate loss distributions, and evaluate how the values returned by each risk function change as we modify their respective parameters. Letting \mathbf{L} denote the random base loss we are simulating, we specify a parametric distribution for \mathbf{L} , from which we take an independent sample $\{\mathbf{L}_1, \dots, \mathbf{L}_m\}$. In all cases, we center the true distribution such that $\mathbf{E}_\mu \mathbf{L} = 0$. We use this common sample to compare the values returned by each of the aforementioned risks, as well as the optimal

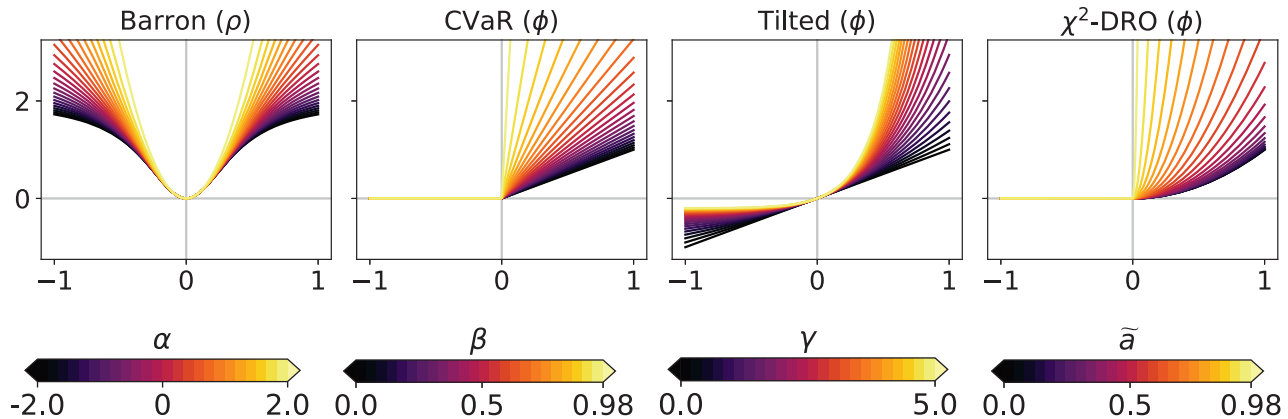


Figure 1: Left-most plot: the graph of $x \mapsto \rho(x/\sigma; \alpha)$ from §3.2 for varying choices of α , with σ fixed to $\sigma = 0.2$ for visual ease. Middle two plots: graphs of $\phi(x)$ in (5) for CVaR and tilted risk, respectively over different choices of β and γ . Right-most plot: graph of $\phi(x)$ in (10) for the χ^2 -DRO risk, where $a \geq 0$ is re-parametrized using $0 \leq \tilde{a} < 1$ using the relation $a = ((1 - \tilde{a})^{-1} - 1)^2/2$.

choice of threshold parameter θ . To ensure that key trends are consistent across samples, we take a large sample size of $m = 10^4$. For T-risk, we adjust α and η , and for M-location, just α . In both cases, we leave $\sigma = 0.5$ fixed. For CVaR, we modify the quantile level $0 < \beta < 1$. For tilted risk, we modify the parameter $\gamma \in \mathbb{R}$. For χ^2 -DRO risk, we modify $0 < \tilde{a} < 1$, having re-parameterized a in (8) by $a := ((1 - \tilde{a})^{-1} - 1)^2/2$, as is common practice (Zhai et al., 2021).

Representative results An illustrative example is given in Figure 2, where we look at how each risk class behaves under a centered asymmetric distribution, before and after flipping it (i.e., under L and $-L$). Starting from the two left-most plots, we show $\underline{R}_\rho(L; \eta)$ (dashed curves) and $\theta_\rho(L; \eta)$ (solid curves) from (12) as a function of α , coloring the area between these graphs in gray. The first plot corresponds to $\eta = 1$, the second to $\eta = -1$. Similarly for the M-location we plot $M_\rho(L)$ (solid) and $M_\rho(L) + \mathbf{E}_\mu \rho(L - M_\rho(L))$ (dashed). Analogous values are plotted for each of the other classes; note that for the tilted risk (3) with $\gamma > 0$, the optimal threshold and the risk value are in fact the same value (see §D.1). The right-most plot is a histogram of the random sample $\{L_1, \dots, L_m\}$, here from a centered log-Normal distribution. All plots share a common vertical axis, and horizontal rules are drawn at the median (red, solid) and at the mean (gray, dotted; always zero due to centering). The critical point to emphasize here is how all the OCE and DRO risks here are highly asymmetric in terms of their tail sensitivity, in stark contrast with both the M-location and the T-risk. Turning tail sensitivity high enough in each of these classes (e.g., $\alpha > 1.5$, $\beta > 0.5$, $\gamma > 1.0$, $\tilde{a} > 0.5$), note how flipping the distribution tails from the upside (top row) to the downside (bottom row) leads to a dramatic decrease in all risks but the T-risk and

M-location. Finally, note how the T-risk thresholds $\theta_\rho(L; \eta)$ close in on the M-location $M_\rho(L)$ from above/below depending on whether η is negative/positive. Results for numerous distributions are available in our online repository (cf. §1).

4 LEARNING ALGORITHM ANALYSIS

We now proceed to consider the learning problem, in which the goal is ultimately to select a candidate h such that the distribution of $L(h)$ is “optimal” in the sense of achieving the smallest possible value of T-risk $R_\rho(h; \theta, \eta)$ in (11), with ρ taken from the Barron class (15), and h taken from some set \mathcal{H} . An obvious take-away of the integrability conditions (Lemma 1) is that even when the base loss is heavy-tailed in the sense of having infinite higher-order moments, we can always adjust the dispersion function $\rho(\cdot; \alpha)$ in such a way that transforming the base loss to obtain new feedback

$$L_\rho(h; \theta, \eta) := \eta\theta + \rho_\sigma(L(h) - \theta; \alpha) \quad (16)$$

gives us an unbiased estimator of the finite T-risk, i.e., $\mathbf{E}_\mu L_\rho(h; \theta, \eta) = R_\rho(h; \theta, \eta) \in \mathbb{R}$. Intuitively, one expects that a similar property can be leveraged to control heavy-tailed stochastic gradients used in an iterative learning algorithm. We explore this point in detail in §4.1. We will then complement this analysis by considering in §4.2 the basic properties of T-risk at the minimal threshold $\underline{R}_\rho(h; \eta)$ given in (12), viewed from the perspectives of axiomatic risk design and empirical risk minimization.

4.1 T-risk and stochastic gradients

For the time being let us fix an arbitrary threshold $\theta \in \mathbb{R}$, and assuming the gradient $L'(h)$ is μ -almost surely finite, denote the partial derivative of the transformed losses (16)




Figure 2: Evaluation of risk class behavior after flipping an asymmetric distribution.

with respect to h by

$$\partial_h L_\rho(h; \theta, \eta) := \rho'_\sigma(L(h) - \theta)L'(h). \quad (17)$$

Writing $\partial_h R_\rho$ for the gradient of $h \mapsto R_\rho(h; \theta, \eta)$, an analogue of Lemma 1 for gradients holds.

Lemma 2 (Unbiased gradients). *Let \mathcal{U} be an open subset of any metric space such that $\mathcal{H} \subseteq \mathcal{U}$. Let the base loss map $h \mapsto L(h)$ be Fréchet differentiable on \mathcal{U} (μ -almost surely), with gradient denoted by $L'(h)$ for each $h \in \mathcal{U}$. Fixing any choice of $-\infty \leq \alpha \leq 2$, we have that*

$$\mathbf{E}_\mu \left[\sup_{h \in \mathcal{H}} \|\rho'_\sigma(L(h))L'(h)\| \right] < \infty \implies \mathbf{E}_\mu [\partial_h L_\rho] = \partial_h R_\rho \quad (18)$$

with the implied equality valid on all of $\mathcal{H} \times \mathbb{R}^2$.

Consider a setting in which the gradients are heavy-tailed, i.e., where $\mathbf{E}_\mu \|L'(h)\|^p = \infty$ for $p > 2$. If the ultimate goal of learning is minimization of $h \mapsto \mathbf{E}_\mu L(h)$, then in order to obtain high-probability guarantees of finding a nearly-stationary point with rates matching the in-expectation case, one cannot naively use the raw gradients $L'(h)$, but must rather carry out a delicate truncation which accounts for the bias incurred (Cutkosky and Mehta, 2021; Gorbunov et al., 2020, 2021; Nazin et al., 2019). On the other hand, if the ultimate objective is $h \mapsto R_\rho(h; \theta, \eta)$, then using (17) there is zero bias by design (Lemma 2), and when we take the shape parameter of our dispersion function such that $\alpha \leq 0$, we have

$$\mathbf{P} \{ \|\partial_h L_\rho(h; \theta, \eta)\| \leq \Gamma \} = 1 \quad (19)$$

for an appropriate choice of $0 < \Gamma < \infty$ under standard loss functions such as quadratic and logistic losses, even when the random losses and gradients are heavy-tailed (see Corollary 4).

To see how this plays out for the analysis of learning algorithms, let us consider plugging the raw stochastic gradients $\partial_h L_\rho$ into a simple update procedure. Given an independent sequence of random losses (L_1, L_2, \dots) , let us denote by $(L_{\rho,1}, L_{\rho,2}, \dots)$ the transformed losses computed via (16), and for a sequence (h_1, h_2, \dots) let $G_t := \partial_h L_{\rho,t}(h_t; \theta, \eta)$ denote the resulting stochastic gradients for any integer $t \geq 1$. Fixing $\theta \in \mathbb{R}$ and letting $h_1 \in \mathcal{H}$ denote an arbitrary initial value, we consider a particular sequence generated using the following update rule:

$$h_{t+1} = h_t - a_t \widetilde{M}_t, \quad (20)$$

where a_t is a non-negative step-size we control, and the update direction satisfies

$$\begin{aligned} \|\widetilde{M}_t\| &= 1, \langle \widetilde{M}_t, M_t \rangle = \|M_t\|, \\ \text{where } M_t &:= bM_{t-1} + (1-b)G_t \end{aligned} \quad (21)$$

with $0 < b < 1$ also being a controllable parameter. This is an unconstrained, normalized stochastic gradient descent routine using momentum; it modifies the procedure of Cutkosky and Mehta (2021) in that we do not truncate G_t . Note that if \mathcal{H} is a Banach space and \mathcal{H}^* its dual, in general we have that G_t and M_t are elements of \mathcal{H}^* . When \mathcal{H} is reflexive (e.g., any Hilbert space), it is simple to construct \widetilde{M}_t from M_t .⁵ The following theorem shows how the gradient norms incurred by this algorithm can be bounded with high probability.

Theorem 3 (Stationary points with high probability). *Let \mathcal{H} be a reflexive Banach space, with a Fréchet differentiable norm satisfying $\|h_1 + h_2\|^2 \leq \|h_1\|^2 + \langle \nabla \|h_1\|^2, h_2 \rangle + \|h_2\|^2$ for any $h_1, h_2 \in \mathcal{H}$. In addition, assume the losses are such that (19) holds on \mathcal{H} , $\mathbf{E}_\mu \|L'(h_1) - L'(h_2)\| \leq$*

⁵Given any $h^* \in \mathcal{H}^*$, we can always find $h \in \mathcal{H}$ such that $\langle h^*, h \rangle = \|h\| \|h^*\|$ (Luenberger, 1969, §5.6).

$\lambda_1 \|h_1 - h_2\|$ for any $h_1, h_2 \in \mathcal{H}$, and $\mathbf{E}_\mu \|L'\|_{\mathcal{H}}^2 < \infty$, where $\|L'\|_{\mathcal{H}} := \sup_{h \in \mathcal{H}} \|L'(h)\|$. Run the learning algorithm in (20)–(21) for T iterations, with $b = 1 - 1/\sqrt{T}$ and $a_t = (1/T)^{3/4}$ for all steps $t = 1, 2, \dots, T$, assuming each $h_t \in \mathcal{H}$. Taking any $0 < \delta < 1$, it then follows that

$$\frac{1}{T} \sum_{t=1}^T \|\partial_h R_\rho(h_t; \theta, \eta)\| \leq \frac{c_1}{T^{1/4}} + \frac{c_2}{\sqrt{T}} + \frac{\lambda}{2T^{3/4}}$$

with probability no less than $1 - \delta$, with coefficients

$$\begin{aligned} c_1 &:= R_\rho(h_1; \theta, \eta) - R_\rho(h_{T+1}; \theta, \eta) \\ &\quad + 16\Gamma \sqrt{\log(3T\delta^{-1})} + 2\lambda \\ c_2 &:= 20\Gamma \log(3T\delta^{-1}) + 2\Gamma \\ \lambda &:= \frac{\lambda_3 + \lambda_4}{\sigma} + \frac{\lambda_2}{\sigma^2} \mathbf{E}_\mu \|L'\|_{\mathcal{H}}, \quad \lambda_2 := \|\rho''\|_\infty, \\ \lambda_3 &:= \left(\frac{2\lambda_2}{\sigma}\right) \mathbf{E}_\mu \|L'\|_{\mathcal{H}}^2, \quad \lambda_4 := \lambda_1 \|\rho'\|_\infty \end{aligned}$$

for any choice of $\sigma > 0$.

These high-probability $\mathcal{O}(T^{-1/4})$ rates match standard guarantees in the stochastic optimization literature under non-convex objectives in expectation (Davis and Drusvyatskiy, 2019; Ghadimi et al., 2016). The main take-away of Theorem 3 is that even if the random losses and gradients are unbounded and heavy-tailed, as long as the dispersion function ρ is chosen to modulate extreme values (such that (19) holds), we can obtain confidence intervals for the T-risk gradient norms incurred by stochastic gradient updates.⁶ Distribution control is implied by the risk design, and thus there is no additional need for truncation or bias control. The following corollary illustrates how the bounded gradient condition in Theorem 3 is satisfied under very weak assumptions on the data.

Corollary 4. *Assume the random losses $L(h)$ are driven by random data (X, Y) , where X takes values in a Banach space \mathcal{X} , and $\mathcal{H} \subset \mathcal{X}^*$ has finite diameter. Consider the following losses:*

- E1. *Quadratic loss:* $L(h) = (h(X) - Y)^2/2$, with $h(X) = \langle h, X \rangle$ and $Y = \langle h^*, X \rangle + \varepsilon$, where ε is zero-mean, has finite variance, and is independent of X .
- E2. *Logistic loss:* $L(h) = \log(\sum_{j=1}^k \exp(\langle h_j, X \rangle)) - \sum_{j=1}^k \tilde{Y}_j \langle h_j, X \rangle$, where we have $k \geq 2$ classes, $h = (h_1, \dots, h_k)$ with each $h_j \in \mathcal{H}$, and $(\tilde{Y}_1, \dots, \tilde{Y}_k)$ is a one-hot representation of the class label Y assigned to X .

If we set $\rho(\cdot) = \rho(\cdot; \alpha)$ with $\alpha \leq 0$, then under the examples E1–E2, as long as $\mathbf{E}_\mu \|X\|^2 < \infty$, we have that the bounds assumed by Theorem 3, including (19), are satisfied on \mathcal{H} .

⁶One can also modify the argument slightly to obtain guarantees for each individual iterate in the style of Cutkosky and Mehta (2021, Thm. 3).

In practice, the threshold θ will not typically be fixed arbitrarily, but rather selected in a data-dependent fashion, potentially optimized alongside h ; the impact of such algorithmic choices will be evaluated in our empirical tests in §C. In the following sub-section, we consider some key properties of the special case in which threshold θ is always taken to yield the smallest overall T-risk value.

Limitations An obvious limitation of our analysis is that we have assumed that each $h_t \in \mathcal{H}$; this matches the setup of (Cutkosky and Mehta, 2021), but the condition that $\mathbf{E}_\mu \|L'\|_{\mathcal{H}}^2 < \infty$ will sometimes require \mathcal{H} to have a finite diameter. Modifying the procedure to allow for projection of the iterates to \mathcal{H} is a point of technical interest, but is out of this paper’s scope. Another limitation is that our current approach using smoothness properties (Lemma 12) necessitates an assumption of gradients with finite second-order moments. In the special case where $\alpha = 1$, arguments based on smoothness can be replaced with arguments based on weak convexity (Davis and Drusvyatskiy, 2019; Holland, 2022); but this fails for more general α since $\rho(\cdot; \alpha)$ is not convex for $\alpha < 1$, and not Lipschitz for $\alpha > 1$. Another potential option is to split the sample, leverage stronger guarantees available in expectation for gradient descent run on the subsets, and robustly choose the best candidate based on a validation set of data (Holland, 2021).

4.2 T-risk with minimizing thresholds

Recall the minimal T-risk $\underline{R}_\rho(h; \eta)$ and optimal thresholds $\theta_\rho(h; \eta)$ defined earlier in (12), here restricted to scaled $\rho(\cdot) = \rho_\sigma(\cdot; \alpha)$ from the Barron class (15). This is arguably the most natural subset of T-risks, with a functional form aligned with OCE/DRO risks discussed in §2. For readability, we denote the dispersion of $L(h)$ measured about $\theta \in \mathbb{R}$ using ρ as

$$D_\rho(h; \theta) := \mathbf{E}_\mu \rho(L(h) - \theta). \quad (22)$$

Here we will overload our notation, writing $\underline{R}_\rho(L; \eta)$, $\theta_\rho(L; \eta)$, and $D_\rho(L; \theta)$ when we want to leave h abstract and focus on random losses in a set $L \in \mathcal{L}$. For this risk to be finite, we require $\alpha \geq 1$, and in addition need $|\eta| < 1/\sigma$ for the special case of $\alpha = 1$; otherwise, the dispersion term grows too slowly and we have $\underline{R}_\rho(L; \eta) = -\infty$ (Lemma 10). When finite, the optimal threshold is unique, and overloading our notation once more we use $\theta_\rho(L; \eta)$ to denote it. The following lemma summarizes some basic properties of the minimal T-risk.

Lemma 5. *Let \mathcal{L} be such that for each $L \in \mathcal{L}$ we have $\underline{R}_\rho(L; \eta) < \infty$, and let α , σ , and η be such that $\underline{R}_\rho(L; \eta) > -\infty$. Under these assumptions, the dispersion part of the minimal T-risk is translation invariant, i.e., for any $c \in \mathbb{R}$, we have $D_\rho(L + c; \theta_\rho(L + c; \eta)) = D_\rho(L; \theta_\rho(L; \eta))$. This dispersion term is always non-negative, and if $\eta \neq 0$, then we have $D_\rho(L; \theta_\rho(L; \eta)) > 0$ for any $L \in \mathcal{L}$, even if L is*

constant. If \mathcal{L} is convex, then so is the map $L \mapsto \underline{R}_\rho(L; \eta)$. The optimal threshold is translation equivariant in that we have $\theta_\rho(L+c; \eta) = \theta_\rho(L; \eta) + c$ for any choice of $c \in \mathbb{R}$, and it is monotonic in that whenever $L_1, L_2 \in \mathcal{L}$ and $L_1 \leq L_2$ almost surely, we have $\theta_\rho(L_1; \eta) \leq \theta_\rho(L_2; \eta)$.

Let us briefly discuss the properties described in Lemma 5 with a bit more context. One of the best-known classes of risk functions is that of *coherent risks* (Artzner et al., 1999), typically characterized by properties of convexity, monotonicity, translation equivariance, and positive homogeneity (Ruszczyński and Shapiro, 2006). Our general notion of “dispersion” is often referred to as “deviation” in the risk literature, and the properties of translation invariance, sub-linearity (implying convexity), non-negativity, and definiteness (i.e., zero only for constants) allow one to establish links between deviations and coherent risks (Rockafellar et al., 2006). In general, the risk $L \mapsto \underline{R}_\rho(L; \eta)$ takes us outside of this traditional class, while still maintaining lucid connections as summarized in the preceding lemma.

Remark 6 (Risk quadrangle). It should also be noted that our T-risks are not what would typically be called “risks” in the context of the “expectation quadrangle” in the framework developed by Rockafellar and Uryasev (2013). With their Example 7 as a clear reference for comparison, our function $\rho(\cdot)$ corresponds to their “error integrand” $e(\cdot)$, and the risk derived from their quadrangle would be

$$L \mapsto \mathbf{E}_\mu L + \mathbf{E}_\mu \rho(L - M_\rho(L)) = \inf_{\theta \in \mathbb{R}} [\theta + \mathbf{E}_\mu v(L - \theta)]$$

where $v(x) := \rho(x) + x$ and $M_\rho(L)$ denotes the M-location (1) under ρ . Under appropriate choice of ρ , this is an OCE-type risk, and evidently our optimal thresholds $\theta_\rho(L; \eta)$ with $\eta \neq 0$ do not make an appearance in risks of this form, highlighting the distinct nature of $\underline{R}_\rho(L; \eta)$.

Next, we briefly consider the question of how algorithms designed to minimize $\underline{R}_\rho(h; \eta)$ perform in terms of the classical risk, namely the expected loss $R(h) := \mathbf{E}_\mu L(h)$. This is a big topic, but as an initial look, we consider how the expected loss incurred by minimizers of the empirical T-risk can be controlled given sufficiently good concentration of the M-estimators induced by ρ , and the empirical mean. Let us denote any empirical T-risk minimizer by

$$\hat{h}_\rho \in \arg \min_{h \in \mathcal{H}} \left[\inf_{\theta \in \mathbb{R}} \left(\eta \theta + \frac{1}{n} \sum_{i=1}^n \rho_\sigma(L_i(h) - \theta) \right) \right] \quad (23)$$

where for simplicity (L_1, \dots, L_n) is an iid sample of random losses.

Proposition 7. *Take any T-risk parameters $\alpha \geq 1$, $\eta > 0$, and $\sigma > 0$, denoting the resulting empirical (minimal) T-risk minimizer \hat{h}_ρ in (23). Let $\hat{R}(h)$, $\hat{M}_\rho(h)$, $\hat{\theta}_\rho(h; \eta)$, and $\hat{D}_\rho(h; \theta)$ denote the empirical analogues of $R(h)$, $M_\rho(h)$,*

$\theta_\rho(h; \eta)$, and $D_\rho(h; \theta)$ for each $h \in \mathcal{H}$. Then, with $\|\cdot\|_{\mathcal{H}}$ as in Theorem 3, we have

$$\begin{aligned} R(\hat{h}_\rho) &\leq R(h^*) + \|\hat{M}_\rho - \hat{\theta}_\rho\|_{\mathcal{H}} + 2\|\hat{M}_\rho - R\|_{\mathcal{H}} \\ &\quad + \frac{1}{\eta} \hat{D}_\rho(h^*; \hat{M}_\rho(h^*)) + 4\|R - \hat{R}\|_{\mathcal{H}} \end{aligned}$$

where we can freely choose h^* to be optimal in the expected loss, i.e., we can take $h^* \in \arg \min_{h \in \mathcal{H}} R(h)$.

We have left the upper bound in Proposition 7 rather abstract to emphasize the key factors that can be used to control expected loss bounds for empirical T-risk minimizers and keep the overall narrative clear. Note that in the special case of $\alpha = 2$ one has $\hat{M}_\rho(h) - \hat{\theta}_\rho(h; \eta) = \eta$, and more generally taking $\eta \rightarrow 0$ sends this difference to zero. There exists tension due to the $1/\eta$ coefficient on the dispersion term. All other terms, including the dispersion itself, are free of η . Note also that the remaining term $\|\hat{M}_\rho - R\|_{\mathcal{H}}$ is the (uniform) difference between the M-estimator induced by ρ and the classical risk; this can be modulated directly using the scale parameter σ , and sharp bounds for a broad classes of M-estimators have been established recently (Minsker, 2019). Our Proposition 7 is analogous to Theorem 7 of Lee et al. (2020), with the key difference being that we study (minimal) T-risks instead of OCE risks, without assuming that losses are bounded (below or above).

5 LEARNING APPLICATIONS

To complement the “static” empirical analysis in §3.3 and the formal insights for learning algorithms in §4, we have also empirically investigated how risk function design and data properties impact the behaviour of stochastic gradient-based learners in both classification and regression tasks, using both simulated data and well-established standard benchmark datasets. We consider three settings:

- Classification with noisy and unbalanced labels
- Control of outlier sensitivity in regression
- Distribution control under SGD on benchmark datasets

Due to limited space, we only give details for the first of these three experiments here; the rest are relegated to §C and §D.5–§D.6 in the supplementary materials. The key take-away here is that even when we introduce a random data-driven learning algorithm into the picture, and the dynamics of the loss/gradient distributions become very complicated, the T-risk class offers a surprisingly high degree of control over the test loss distribution incurred by final algorithm output. From control of outlier sensitivity even under high-leverage points, to a smooth interpolation between ERM and mean-variance solutions using stochastic gradient descent to update thousands of parameters, our initial results are suggestive of a useful class of performance criteria that can be used alongside OCE/DRO risks.




Figure 3: Left-most plot: simulated binary classification data, with unbalanced classes and randomly flipped labels (indicated by red diamonds). Remaining plots: loss distributions for three types of loss functions, evaluated at two different candidates, corresponding to the dashed pink and dotted purple lines in the left-most plot.




Figure 4: Classification error and norm trajectories (over iteration number, \log_{10} scale) for gradient descent implementations of each risk class, using unhinged base loss and data given in Figure 3.

5.1 Classification with noisy and unbalanced labels

We start with what is essentially a toy example that lets us explore how the tails and orientation of the loss distribution impacts typical gradient-based classifiers.

Experiment setup As an initial example using simulated data, we design a binary classification problem by generating 500 labeled data points as shown in Figure 3 (left-most plot). The majority class comprises 95% of the sample, and we flip 5% of the labels uniformly at random. Aside from the flipped labels, the data is linearly separable. We consider two candidate classifiers to initialize an iterative learning algorithm: one candidate that does well on the majority class (dotted purple), and one candidate that does well on the minority class (dashed pink). Due to class imbalance, the loss distributions incurred by these two candidates are highly asymmetric and differ in their long tail direction (Figure 3, remaining plots). We run empirical risk minimization for each risk class described in §2 (joint in (h, θ)), implemented by 15,000 iterations of full-batch gradient descent, using a step size of 0.01, and three different choices of base loss function (logistic, hinge, unhinged). We run this procedure on the same data for a wide range of risk parameters α (T-

risk), β (CVaR), γ (tilted), and a (χ^2 -DRO), and choose a representative setting for each risk class as the one achieving the best final training classification error.

Representative results The full-batch (training) classification error and norm trajectories for these representatives initialized at each of the two candidates just mentioned (“mostly correct” and “mostly incorrect”) and trained using the unhinged loss as a base loss are shown in Figure 4. Regardless of the direction of the initial loss distribution tails, we see that using the Barron-type T-risk has the flexibility to achieve both a stable and superior long-term error rate, while at the same time penalizing exceedingly overconfident correct examples (via bi-directionality), and thereby keeping the norm of the linear classifier small. An almost identical trend is also observed under the (binary) logistic loss, whereas classification error rates under the hinge loss tend to be less stable (see Figures 10–11 in §D.6).

Acknowledgements

This work was supported by JST ACT-X Grant Number JPMJAX2000 and JST PRESTO Grant Number JPMJPR21C6.

References

- Albert, A. and Anderson, J. A. (1984). On the existence of maximum likelihood estimates in logistic regression models. *Biometrika*, 71(1):1–10.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.
- Ash, R. B. and Doléans-Dade, C. A. (2000). *Probability and Measure Theory*. Academic Press, 2nd edition.
- Barron, J. T. (2019). A general and adaptive robust loss function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4331–4339.
- Bassett, Jr, G. W. (1987). The St. Petersburg paradox and bounded utility. *History of Political Economy*, 19(4):517–523.
- Ben-Tal, A., Den Hertog, D., De Waegenaere, A., Melenberg, B., and Rennen, G. (2013). Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357.
- Ben-Tal, A. and Teboulle, M. (1986). Expected utility, penalty functions, and duality in stochastic nonlinear programming. *Management Science*, 32(11):1445–1466.
- Ben-Tal, A. and Teboulle, M. (2007). An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance*, 17(3):449–476.
- Blackard, J. A. and Dean, D. J. (1999). Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and Electronics in Agriculture*, 24(3):131–151.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Cohen, G., Afshar, S., Tapson, J., and Van Schaik, A. (2017). EMNIST: an extension of MNIST to handwritten letters. *arXiv preprint arXiv:1702.05373v2*.
- Şimşekli, U., Sagun, L., and Gürbüzbalaban, M. (2019). A tail-index analysis of stochastic gradient noise in deep neural networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 5827–5837.
- Curi, S., Levy, K. Y., Jegelka, S., and Krause, A. (2020). Adaptive sampling for stochastic risk-averse learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, pages 1036–1047.
- Cutkosky, A. and Mehta, H. (2021). High-probability bounds for non-convex stochastic optimization with heavy tails. *arXiv preprint arXiv:2106.14343v2*.
- Davis, D. and Drusvyatskiy, D. (2019). Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239.
- Devroye, L., Györfi, L., and Lugosi, G. (1996). *A Probabilistic Theory of Pattern Recognition*. Springer.
- Duchi, J. and Namkoong, H. (2018). Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750v6*.
- Duchi, J. and Namkoong, H. (2019). Variance-based regularization with convex objectives. *Journal of Machine Learning Research*, 20(68):1–55.
- Föllmer, H. and Knispel, T. (2011). Entropic risk measures: Coherence vs. convexity, model ambiguity and robust large deviations. *Stochastics and Dynamics*, 11(02n03):333–351.
- Ghadimi, S., Lan, G., and Zhang, H. (2016). Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2):267–305.
- Gneiting, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106(494):746–762.
- Gorbunov, E., Danilova, M., and Gasnikov, A. (2020). Stochastic optimization with heavy-tailed noise via accelerated gradient clipping. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*.
- Gorbunov, E., Danilova, M., Shibaev, I., Dvurechensky, P., and Gasnikov, A. (2021). Near-optimal high probability complexity bounds for non-smooth stochastic optimization with heavy-tailed noise. *arXiv preprint arXiv:2106.05958*.
- Gotoh, J.-y., Kim, M. J., and Lim, A. E. (2018). Robust empirical optimization is almost the same as mean–variance optimization. *Operations Research Letters*, 46(4):448–452.
- Hacking, I. (2006). *The emergence of probability: a philosophical study of early ideas about probability, induction and statistical inference*. Cambridge University Press, 2nd edition.
- Hashimoto, T. B., Srivastava, M., Namkoong, H., and Liang, P. (2018). Fairness without demographics in repeated loss minimization. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 1929–1938.
- Hausler, D. (1992). Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100(1):78–150.
- Holland, M. J. (2021). Robustness and scalability under heavy tails, without strong convexity. In *Proceedings*

- of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS), volume 130 of *Proceedings of Machine Learning Research*.
- Holland, M. J. (2022). Learning with risks based on M-location. *Machine Learning*.
- Holland, M. J. and Haress, E. M. (2021). Learning with risk-averse feedback under potentially heavy tails. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 130 of *Proceedings of Machine Learning Research*.
- Huber, P. J. and Ronchetti, E. M. (2009). *Robust Statistics*. John Wiley & Sons, 2nd edition.
- Kashima, H. (2007). Risk-sensitive learning via minimization of empirical conditional value-at-risk. *IEICE Transactions on Information and Systems*, 90(12):2043–2052.
- Kearns, M. J. and Vazirani, U. V. (1994). *An Introduction to Computational Learning Theory*. MIT Press.
- Koltchinskii, V. I. (1997). M -estimation, convexity and quantiles. *The Annals of Statistics*, pages 435–477.
- Lee, J., Park, S., and Shin, J. (2020). Learning bounds for risk-sensitive learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, pages 13867–13879.
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. (2021a). On tilted losses in machine learning: Theory and applications. *arXiv preprint arXiv:2109.06141v1*.
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. (2021b). Tilted empirical risk minimization. In *The 9th International Conference on Learning Representations (ICLR)*.
- Luenberger, D. G. (1969). *Optimization by Vector Space Methods*. John Wiley & Sons.
- Minsker, S. (2019). Uniform bounds for robust mean estimators. *arXiv preprint arXiv:1812.03523v4*.
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2012). *Foundations of Machine Learning*. MIT Press.
- Nazin, A. V., Nemirovsky, A. S., Tsybakov, A. B., and Juditsky, A. B. (2019). Algorithms of robust stochastic optimization based on mirror descent method. *Automation and Remote Control*, 80(9):1607–1627.
- Nesterov, Y. (2004). *Introductory Lectures on Convex Optimization: A Basic Course*. Springer.
- Penot, J.-P. (2012). *Calculus Without Derivatives*, volume 266 of *Graduate Texts in Mathematics*. Springer.
- Prashanth, L. A., Jagannathan, K., and Kolla, R. K. (2020). Concentration bounds for CVaR estimation: The cases of light-tailed and heavy-tailed distributions. In *37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pages 5577–5586.
- Rockafellar, R. T. and Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42.
- Rockafellar, R. T. and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471.
- Rockafellar, R. T. and Uryasev, S. (2013). The fundamental risk quadrangle in risk management, optimization and statistical estimation. *Surveys in Operations Research and Management Science*, 18(1-2):33–53.
- Rockafellar, R. T., Uryasev, S., and Zabarankin, M. (2006). Generalized deviations in risk analysis. *Finance and Stochastics*, 10(1):51–74.
- Rousseeuw, P. J. and Christmann, A. (2003). Robustness against separation and outliers in logistic regression. *Computational Statistics & Data Analysis*, 43(3):315–332.
- Rousseeuw, P. J. and Leroy, A. M. (1987). *Robust Regression and Outlier Detection*. John Wiley & Sons.
- Ruszczynski, A. and Shapiro, A. (2006). Optimization of convex risk functions. *Mathematics of Operations Research*, 31(3):433–452.
- Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition.
- Takeda, A. and Sugiyama, M. (2008). ν -support vector machine as conditional value-at-risk minimization. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, pages 1056–1063.
- Takeuchi, I., Le, Q. V., Sears, T. D., and Smola, A. J. (2006). Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7:1231–1264.
- Van Rooyen, B., Menon, A., and Williamson, R. C. (2016). Learning with symmetric label noise: The importance of being unhinged. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pages 10–18.
- Vapnik, V. N. (1999). *The Nature of Statistical Learning Theory*. Statistics for Engineering and Information Science. Springer, 2nd edition.
- Venables, W. N. and Ripley, B. D. (2002). *Modern Applied Statistics with S*. Springer, 4th edition.
- Zhai, R., Dan, C., Kolter, J. Z., and Ravikumar, P. (2021). DORO: Distributional and outlier robust optimization. In *38th International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 12345–12355.
- Zhang, J., Karimireddy, S. P., Veit, A., Kim, S., Reddi, S., Kumar, S., and Sra, S. (2020). Why are adaptive methods good for attention models? In *Advances in Neural Information Processing Systems*, volume 33, pages 15383–15393.

A Appendix summary

For ease of reference, we start the appendix with a concise tables of contents.

- §B: Notation
- §C: Learning applications
 - §C.1: Control of outlier sensitivity
 - §C.2: Distribution control under SGD on benchmark datasets
- §D: Additional information
 - §B: Notation
 - §D.1: Details for tilted risk
 - §D.2: Details for Barron class limits
 - §D.3: Additional lemmas for Barron class and T-risk
 - §D.4: Smoothness of the T-risk
 - §D.5: Experimental details
 - §D.6: Additional figures
- §E: Detailed proofs
 - §E.1: Proofs of results in the main text
 - §E.2: Smoothness computations (proof of Lemma 12)
- §F: Additional technical facts
 - §F.1: Lipschitz properties
 - §F.2: Convexity
 - §F.3: Derivatives for the Barron class
 - §F.4: Elementary inequalities
 - §F.5: Expected dispersion is coercive

B Notation

Random quantities To start, let us clarify the nature of the random losses we consider. With the context of the underlying probability space $(\Omega, \mathcal{F}, \mu)$, we write $L(h) := L(h; \cdot) : \Omega \rightarrow \mathbb{R}$ to refer to a random variable (i.e., a \mathcal{F} -measurable function) on Ω , though we only use the form $L(h)$ in the body of this paper. When we talk about “sampling” losses or a “random draw” of the losses, this amounts to computing a realization $L(h; \omega) \in \mathbb{R}$. We use standard notation for taking expectation, e.g., $\mathbf{E}_\mu L(h) := \int_\Omega L(h; \omega) \mu(d\omega)$. These conventions extend to random quantities based on the losses (e.g., the gradient $L'(h)$ considered in §4.1). Similarly, we will use \mathbf{P} as a general-purpose probability function, representing both μ and product measures; when the source of randomness is not immediate from the context, it will be stated explicitly. We will use $\mathbf{R}(\cdot)$ as a generic symbol for risk functions (often modified with subscripts), with the understanding that $\mathbf{R}(\cdot)$ maps random losses $L(h)$ to real values. We will overload this notation, writing $\mathbf{R}(L)$ when the role of h is unimportant, and writing $\mathbf{R}(h) := \mathbf{R}(L(h))$ when we want to emphasize the dependence on h . This convention will be applied to other quantities as well, such as writing $D_\rho(h) := D_\rho(L(h)) := \mathbf{E}_\mu \rho(L(h) - \theta)$ for the expected dispersion induced by ρ , first defined in (22).

Norms We will use $\|\cdot\|$ as a general-purpose notation for all norms that appear in this paper. That is, we do not use different notation to distinguish different norm spaces. The reason for this is that we will never consider two distinct norms on the same set; each norm is associated with a distinct set, and thus as long as it is clear which set a particular element belongs to, there should be no confusion. The only exception to this rule is the special case of \mathbb{R} , in which we write $|\cdot|$ for the absolute value, as is traditional. For a function $f : \mathbb{R} \rightarrow \mathbb{R}$ in one variable, we use f' to denote the usual derivative. More general notions (e.g., Gateaux or Fréchet differentials) only make an appearance in §4, and the generality they afford us is not crucial to the main narrative, so the details can be easily skipped over if the reader is unfamiliar with such concepts. All the other undefined notation we use is essentially standard, and can be found in most introductory analysis textbooks.

Miscellaneous For a function $f : \mathbb{R} \rightarrow \mathbb{R}$ in one variable, we use f' to denote the usual derivative. More general notions (e.g., Gateaux or Fréchet differentials) only make an appearance in §4, and the generality they afford us is not crucial to the main narrative, so the details can be easily skipped over if the reader is unfamiliar with such concepts. All the other undefined notation we use is essentially standard, and can be found in most introductory analysis textbooks. Particularly in formal proofs, we will frequently make use of the notation $\rho_\sigma(x) := \rho(x/\sigma)$, where $\sigma > 0$ is a scaling parameter.

C Learning applications

To complement the “static” empirical analysis in §3.3 and the formal insights for learning algorithms in §4, here we empirically investigate how risk function design and data properties impact the behaviour of stochastic gradient-based learners.

C.1 Control of outlier sensitivity

Experiment setup For a crystal-clear example of real data that induces losses with heavy tails, we consider the famous Belgian phone call dataset (Rousseeuw and Leroy, 1987; Venables and Ripley, 2002) with input features normalized to the unit interval, and raw outputs used as-is. We conduct two regression tasks: the first uses the original data just stated, and the second has us modify such that it is an outlier on both the horizontal and vertical axes (we just multiply the original data point by 5); such points are said to have “high leverage” (Huber and Ronchetti, 2009, Ch. 7). For these two tasks, once again we run empirical risk minimization under each of the risk classes of interest, implemented using 15,000 iterations of full-batch gradient descent, with fixed step size 0.005. To illustrate the flexibility of the T-risk, here instead of minimizing (11) jointly in (h, θ) , we fix θ at the start of the learning process along with σ and η , iteratively optimizing h only (i.e., θ is not updated at any point). For simplicity, we set θ and σ both to be the median of the losses incurred at initialization. All other risk classes are precisely as in the previous experiment.

Representative results The final regression lines obtained under each risk setting by running the learning procedure just described are plotted along with the data in Figure 5. Colors correspond to the individual risk function choices within each risk family, as denoted by color bars under each plot. The gray regression line denotes the common initial value used by all methods. Algorithm outputs using CVaR and χ^2 -DRO are always at least as sensitive to outliers as the ordinary least-squares (OLS) solution. While the tilted risk does let us interpolate between lower and upper quantiles, this transition is not smooth; even trying 20 values within the small window of $\gamma \in [-0.025, 0.025]$, the algorithm outputs essentially jump between two extremes. This difference is particularly lucid in the bottom plots of Figure 5, where we have a high-leverage point. In contrast, with all other parameters of T-risk fixed, note that just tweaking α gives us a remarkable degree of flexibility to control the final output; while the base loss is fixed to the squared error, the regression lines range from those that ignore outliers to those that are very sensitive to outliers. In the high-leverage case, it is well-known that this cannot be achieved by simply changing to a different convex base loss (e.g., MAE instead of OLS), giving a concise illustration of the flexibility inherent in the T-risk class. Algorithm behavior can be sensitive to initial value; see §D.5 and Figure 13 in the appendix for an example where the naive median-based procedure described here causes gradient-based algorithms to stall out.

C.2 Distribution control under SGD on benchmark datasets

Experiment setup Finally we consider tests using datasets and models which are orders of magnitude larger than the previous two experimental setups. We use several well-known benchmark datasets for multi-class classification; details are given in §D.5. All features are normalized to the unit interval (categorical variables are one-hot), and scores for each class are computed by a linear combination of features. As a base loss, we use the multi-class logistic regression loss. Here we investigate how a stochastic gradient descent implementation (with averaging) of the empirical T-risk minimization (joint in (h, θ)) behaves as we control the shape parameter α of the Barron class, with $\sigma = 0.99$ and $\eta = 1$ fixed throughout. We record the base loss (logistic loss) and zero-one loss (misclassification error) at the start and end of each epoch. We also record the full base loss distribution after the last epoch concludes. We run 5 independent trials, in which the data is shuffled and initial parameters are determined randomly. In each trial, for all datasets and methods, we use a mini-batch size of 32, and we run 30 epochs. As reference algorithms, we consider “vanilla” ERM, using the traditional risk $\mathbf{E}_\mu \mathbf{L}(h)$, and mean-variance implemented as a special case of T-risk (with $\rho(\cdot) = (\cdot)^2/2$ and $\eta = 1$). 80% of the data is used for training, 10% for validation, and 10% for testing. For each risk class, we try five different step sizes. Validation data is used to evaluate different step sizes and choose the best one for each risk setting. All the results we present here are based on loss values computed on the test set: solid lines represent averages taken over trials, and shaded areas denote \pm standard

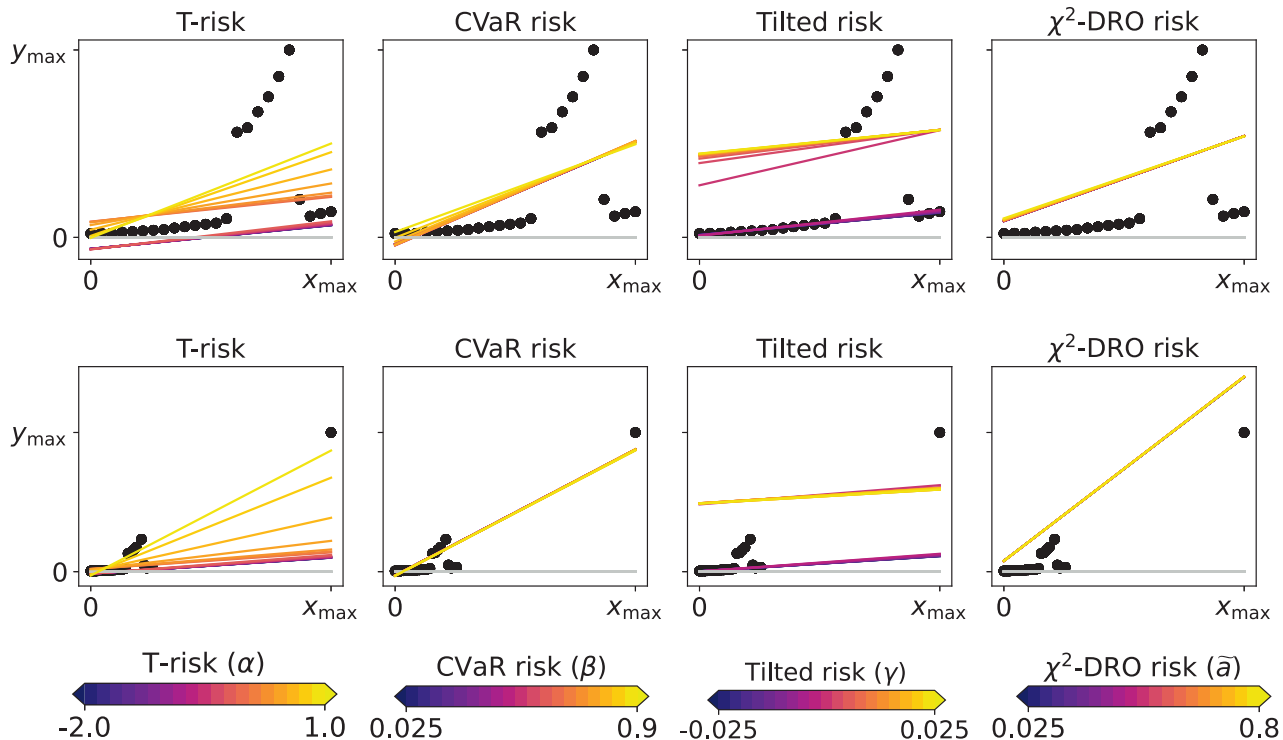


Figure 5: Using T-risks with thresholds set to median initialization error, we can achieve a smooth transition between outlier-robust and outlier-sensitive solutions, even under high-leverage points. Top: original “phones” dataset. Bottom: modified data including a single high-leverage point.

deviation over trials.

Representative results In Figures 6–7, we give results based on the following two datasets: “extended MNIST” (47 classes, balanced) (Cohen et al., 2017), and “cover type” (7 classes, imbalanced) (Blackard and Dean, 1999). We plot the average and standard deviation of the base and zero-one losses as a function of epoch number, plus give a histogram of test (base) losses for a single trial, compared with the loss distribution incurred by a random initialization of the same model (left-most plot; gray is test, black is training). Colors are analogous to previous experiments, here evenly spaced over the allowable range ($1 \leq \alpha \leq 2$). It is clear how modifying the Barron dispersion function shape across this range lets us flexibly interpolate between the test (base) loss distribution achieved by a traditional ERM solution and a mean-variance solution, in terms of both the mean and standard deviation. This monotonicity (as a function of α) is salient in the base loss, but this does not always appear for the zero-one loss. Since these datasets have normalized features with negligible label noise, egregious outliers are rare, and thus the trends observed for the mean and standard deviation here also hold for outlier-resistant location-dispersion pairs such as the median and the median-absolute-deviations about the median. Similar results for several other datasets are provided in §D.6, and we remark that the key trends hold across all datasets tested. We have seen in our previous experiments how under heavy-tailed losses/gradients the T-risk solution can differ greatly from that of the vanilla ERM solution, so it is interesting to observe how under large, normalized, clean classification datasets, the T-risk allows us to very smoothly control a tradeoff between average test loss and variance on the test set.

D Additional information

Here we provide further details that complement the body of this paper.

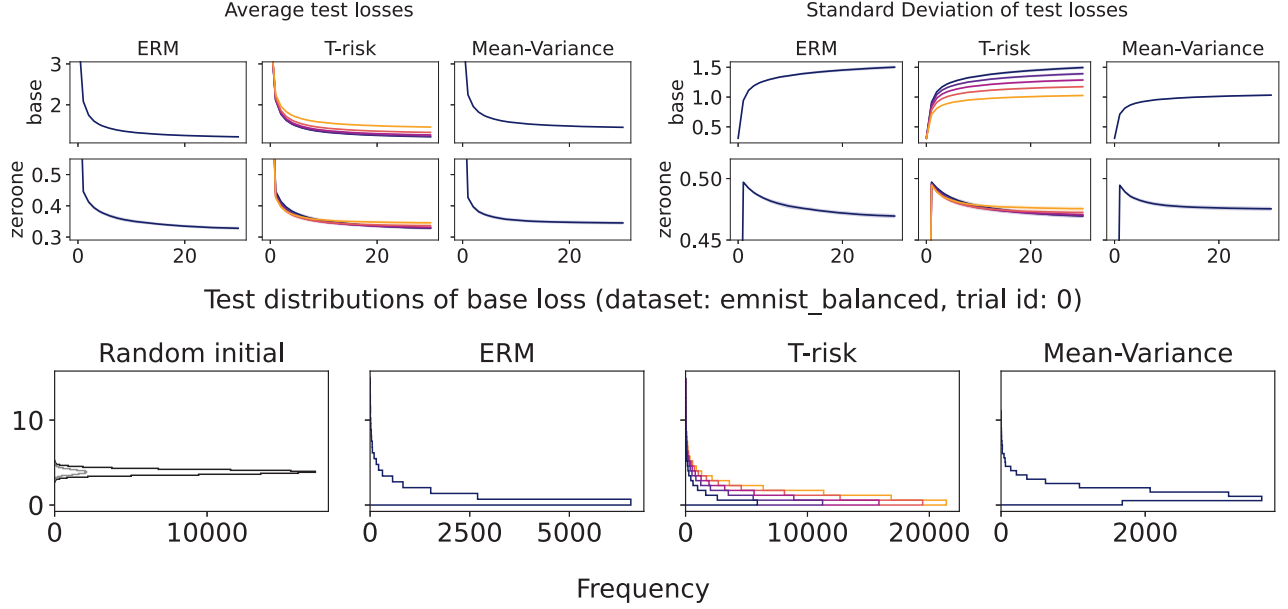


Figure 6: Using T-risk to interpolate between test loss distributions (dataset: emnist_balanced).

D.1 Details for tilted risk

Let $X \sim \mu$ be an arbitrary random variable. Assuming the distribution is such that we can differentiate through the integral, we have

$$\frac{d}{d\theta} \left[\theta + \frac{1}{\gamma} \left(\mathbf{E}_{\mu} e^{\gamma(X-\theta)} - 1 \right) \right] = 0 \iff \mathbf{E}_{\mu} e^{\gamma(X-\theta)} = 1. \quad (24)$$

Let θ^* be any value that satisfies the first-order optimality condition in (24). It follows that

$$\theta^* + \frac{1}{\gamma} \left(\mathbf{E}_{\mu} e^{\gamma(X-\theta^*)} - 1 \right) = \theta^*.$$

It is easy to confirm that setting $\theta^* = (1/\gamma) \log(\mathbf{E}_{\mu} e^{\gamma X})$ gives us a valid solution. For more background, see the recent works of Li et al. (2021b,a) and the references therein. Note also that this log-exponential criterion appears (with $\gamma = 1$) in Rockafellar and Uryasev (2013, Ex. 8).

D.2 Details for Barron class limits

For the limit as $\alpha \rightarrow 0$, use the fact that for any $a > 0$, we have

$$\lim_{x \rightarrow 0^+} \frac{(a^x - 1)}{x} = \log(a). \quad (25)$$

This equality is sometimes known as Halley's formula. For the limit as $\alpha \rightarrow -\infty$, first note that for any $\alpha < 0$ we can write $|2 - \alpha| = 2 + |\alpha|$, and thus $|\alpha|/2 = (|2 - \alpha|/2) - 1$. With this in mind, for any $\alpha < 0$ we can observe

$$\begin{aligned} \left(1 + \frac{x^2}{|\alpha - 2|} \right)^{\alpha/2} &= \left(1 + \frac{x^2}{|\alpha - 2|} \right)^{-|\alpha|/2} = \left(1 + \frac{x^2}{|\alpha - 2|} \right)^{1 - (|2 - \alpha|/2)} \\ &= \frac{\left(1 + \frac{x^2}{|\alpha - 2|} \right)}{\sqrt{\left(1 + \frac{x^2}{|\alpha - 2|} \right)^{|\alpha - 2|}}} \\ &\rightarrow \frac{1}{\sqrt{\exp(x^2)}} = \exp(-x^2/2), \end{aligned}$$

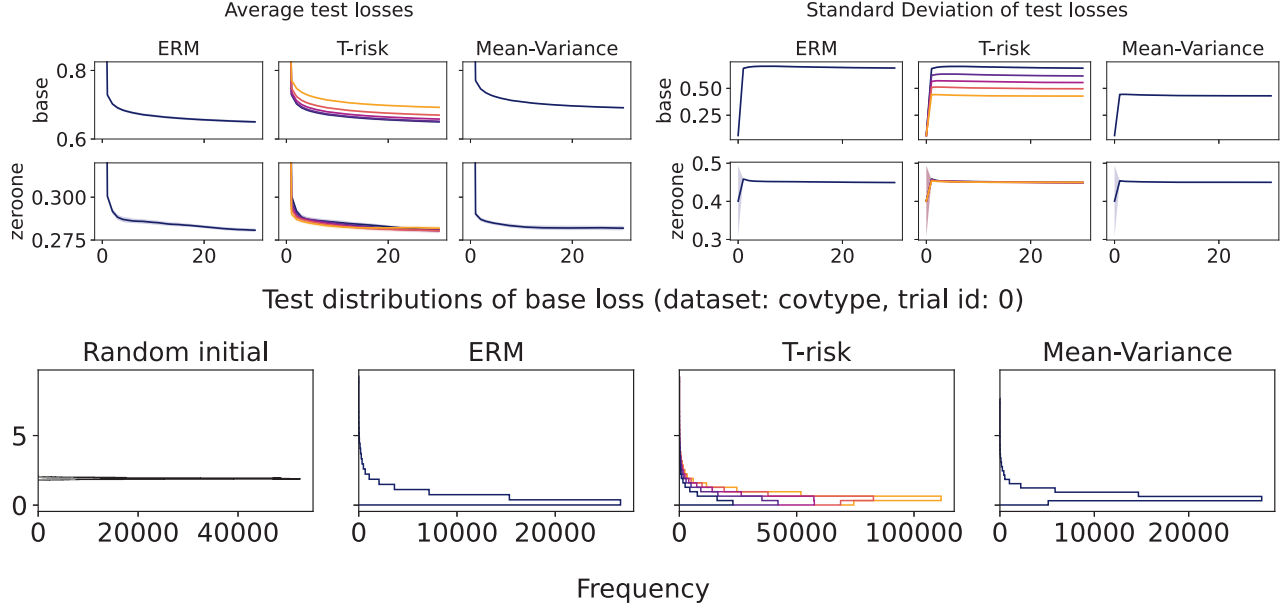


Figure 7: Analogous to Figure 6 (dataset: covtype).

where the limit is taken as $\alpha \rightarrow -\infty$, and follows from the classical limit characterization of the exponential function. For the limit as $\alpha \rightarrow 2_-$, first note that

$$|\alpha - 2| \left(1 + \frac{x^2}{|\alpha - 2|} \right)^{\alpha/2} = \left(|\alpha - 2|^{2/\alpha} + |\alpha - 2|^{(2/\alpha)-1} x^2 \right)^{\alpha/2}$$

and that as long as $\alpha < 2$, we can write

$$|\alpha - 2|^{(2/\alpha)-1} = (2 - \alpha)^{(2-\alpha)/\alpha} = (u^u)^{1/(2-u)}$$

where we have introduced $u := 2 - \alpha$. Taking $\alpha \rightarrow 2_-$ amounts to $u \rightarrow 0_+$, and thus using the fact that $u^u \rightarrow 1$ as $u \rightarrow 0_+$, the desired result follows from straightforward analysis.

D.3 Additional lemmas for Barron class and T-risk

Lemma 8 (Dispersion function convexity and smoothness). *Consider $\rho_\sigma(x) := \rho(x/\sigma; \alpha)$ with $\rho(\cdot; \alpha)$ from the Barron class (15), with parameter $-\infty \leq \alpha \leq 2$. The following properties hold for any choice of $\sigma > 0$.*

- Case of $\alpha = 2$:
 ρ_σ is convex and $1/\sigma^2$ -smooth on \mathbb{R} .
- Case of $\alpha = 0$:
 ρ_σ is convex on $[-\sqrt{2}\sigma, \sqrt{2}\sigma]$, and is $1/(\sqrt{2}\sigma)$ -Lipschitz and $1/\sigma^2$ -smooth on \mathbb{R} .
- Case of $\alpha = -\infty$:
 ρ_σ is convex on $[-\sigma, \sigma]$, and is $(1/\sigma) \exp(-1/2)$ -Lipschitz and $1/\sigma^2$ -smooth on \mathbb{R} .
- Otherwise:
 ρ_σ is $1/\sigma^2$ -smooth on \mathbb{R} . When $\alpha \geq 1$, ρ_σ is convex on \mathbb{R} . When $\alpha = 1$, ρ_σ is also $1/\sigma$ -Lipschitz on \mathbb{R} . Else, when $\alpha < 1$, we have that ρ_σ is convex between $\pm \sigma \sqrt{|\alpha - 2|/(1 - \alpha)}$, and is $(1/\sigma)(\sqrt{(1 - \alpha)/|\alpha - 2|})^{1-\alpha}$ -Lipschitz on \mathbb{R} .

Furthermore, all these coefficients are tight (see also Figure 9).




Figure 8: Here we consider the function $x \mapsto \rho(x/\sigma; \alpha)$, with ρ from the Barron class (15), for a variety of shape parameter values α , with scale parameter $\sigma = 0.5$ fixed. We plot graphs for this function (left), its first derivative (center), and its second derivative (right), each computed over $\pm 5\sigma$.

The dispersion $D_\rho(h; \theta) := \mathbf{E}_\mu \rho(L(h) - \theta)$ plays a prominent role in the risk definitions considered in this paper, and one is naturally interested in the properties of the map $\theta \mapsto D_\rho(h; \theta)$. The following lemma shows that using $\rho = \rho_\sigma$ from the Barron class, we can differentiate under the integral without needing any additional conditions beyond those required for finiteness.

Lemma 9. *Let ρ_σ be as in Lemma 8. Assume that the random loss $L(h)$ is \mathcal{F} -measurable in general, and that $\mathbf{E}_\mu |L(h)| < \infty$ holds whenever $1 < \alpha \leq 2$. It follows that the first two derivatives are μ -integrable, namely that*

$$|\mathbf{E}_\mu \rho'_\sigma(L(h) - \theta)| < \infty, \quad |\mathbf{E}_\mu \rho''_\sigma(L(h) - \theta)| < \infty \quad (26)$$

for any $\theta \in \mathbb{R}$. Furthermore, the function $\theta \mapsto D_\rho(h; \theta)$ is twice-differentiable on \mathbb{R} , and satisfies the Leibniz integration property for both derivatives, that is

$$\left. \frac{d}{d\theta} D_\rho(h; \theta) \right|_{\theta=u} = -\mathbf{E}_\mu \rho'_\sigma(L(h) - u), \quad \left. \frac{d^2}{d\theta^2} D_\rho(h; \theta) \right|_{\theta=u} = \mathbf{E}_\mu \rho''_\sigma(L(h) - u) \quad (27)$$

for any $u \in \mathbb{R}$.⁷

With first-order information about the expected dispersion in hand, one can readily obtain conditions under which the special case of T-risk $\underline{R}_\rho(h; \eta) := \inf_{\theta \in \mathbb{R}} R_\rho(h; \theta, \eta)$ is finite and determined by a meaningful “optimal threshold.”

Lemma 10. *Following the setup of Lemma 1, let $D_\rho(h; \theta) < \infty$ for all $\theta \in \mathbb{R}$. If $\alpha > 1$, then for any choice of $\sigma > 0$ and $\eta \in \mathbb{R}$, there exists a finite optimal threshold $\theta_\rho(h; \eta) \in \mathbb{R}$ such that*

$$\underline{R}_\rho(h; \eta) = \eta \theta_\rho(h; \eta) + D_\rho(h; \theta_\rho(h; \eta)). \quad (28)$$

In the case of $\alpha = 1$, we have that (28) holds if and only if $|\eta| < 1/\sigma$. In both cases, the optimal threshold is unique. In the case of $\alpha < 1$, then there is no minimum and thus $\underline{R}_\rho(h; \eta) = -\infty$.

Remark 11. We have overloaded our notation $\theta_\rho(h; \eta)$ in Lemma 10, recalling that we have used the same notation to denote the set of optimal thresholds for the T-risk in (11). This saves us from having to introduce additional symbols, and should not lead to any confusion since we only do this overloading when the solution set contains a single unique solution.

D.4 Smoothness of the T-risk

When the objective $R_\rho(h; \theta, \eta)$ in (11) is sufficiently smooth, we can apply well-established analytical techniques to control the gradient norms of stochastic gradient-based learning algorithms. Assuming we have unbiased first-order stochastic feedback as in (17)–(18), we will always have to deal with terms of the form $\mathbf{E}_\mu [\rho'_\sigma(L(h) - \theta)L'(h)]$. Defining

⁷Let us emphasize that ρ'_σ and ρ''_σ denote the first and second derivatives of $x \mapsto \rho(x/\sigma; \alpha)$, which differ from $\rho'(x/\sigma; \alpha)$ and $\rho''(x/\sigma; \alpha)$ by a σ -dependent factor; see §F.3 for details.




Figure 9: Here we show zoomed-in versions of the two right-most plots in Figure 8, now with dashed horizontal lines denoting the Lipschitz (left) and smoothness (right) coefficients stated in Lemma 8 (both positive and negative values). The Lipschitz coefficients depend on α and σ , and the colors in the left plot reflect the α value. The smoothness coefficient $1/\sigma^2$ is independent of α , and is drawn in light gray. Numerically, we can see that the coefficients are tight, though tighter *lower* bounds for the second derivatives are possible.

$f(h, \theta) := \rho'_\sigma(\mathbf{L}(h) - \theta)\mathbf{L}'(h)$ for readability, and considering the function difference at two arbitrary points (h_1, θ_1) and (h_2, θ_2) , first note that

$$\begin{aligned} f(h_1, \theta_1) - f(h_2, \theta_2) &= \underbrace{\rho'_\sigma(\mathbf{L}(h_1) - \theta_1) [\mathbf{L}'(h_1) - \mathbf{L}'(h_2)]}_A + \underbrace{[\rho'_\sigma(\mathbf{L}(h_1) - \theta_1) - \rho'_\sigma(\mathbf{L}(h_2) - \theta_2)] \mathbf{L}'(h_2)}_B. \end{aligned} \quad (29)$$

In the case of $\rho(\cdot) = \rho(\cdot; \alpha)$ from the Barron class (15), when $-\infty \leq \alpha \leq 1$, we have that ρ' is both bounded ($\|\rho'\|_\infty < \infty$) and Lipschitz continuous on \mathbb{R} (see Lemma 8). This means that all we need in order to control $\mathbf{E}_\mu A + \mathbf{E}_\mu B$ is for \mathbf{L} to be smooth (for control of A) and for $\mathbf{L}'(\cdot)$ to have a norm bounded over \mathcal{H} (for control of B); see §E.2 for more details. Things are more difficult in the case of $1 < \alpha \leq 2$, since the dispersion function ρ is not (globally) Lipschitz, meaning that $\|\rho'\|_\infty = \infty$. Even if \mathbf{L} is smooth, when the threshold parameter is left unconstrained, it will always be possible for $\|\mathbf{E}_\mu A\| \rightarrow \infty$ as $|\theta_1| \rightarrow \infty$, impeding smoothness guarantees for R_ρ in this setting.

Let us proceed by distilling the preceding discussion into a set of concrete conditions that are sufficient to make $(h, \theta) \mapsto R_\rho(h; \theta, \eta)$ amenable to standard analysis techniques for stochastic gradient-based algorithms. For readability, we write $\|\mathbf{L}'\|_{\mathcal{H}} := \sup\{\|\mathbf{L}'(h)\| : h \in \mathcal{H}\}$.

A1. Moment bound for loss gradient. For any choice of $h_1, h_2 \in \mathcal{H}$, $0 < c < 1$, and $k \in \{1, 2\}$, the loss \mathbf{L} is differentiable at $ch_1 + (1 - c)h_2$, and satisfies

$$\mathbf{E}_\mu \left(\sup_{0 < c < 1} \|\mathbf{L}'(ch_1 + (1 - c)h_2)\| \right)^k \leq \mathbf{E}_\mu \|\mathbf{L}'\|_{\mathcal{H}}^k < \infty. \quad (30)$$

A2. Loss is smooth in expectation. There exists $0 < \lambda_1 < \infty$ such that for any choice of $h_1, h_2 \in \mathcal{H}$, we have $\mathbf{E}_\mu \|\mathbf{L}'(h_1) - \mathbf{L}'(h_2)\| \leq \lambda_1 \|h_1 - h_2\|$.

A3. Dispersion is Lipschitz and smooth. The function ρ is such that $\|\rho'\|_\infty < \infty$, and there exists $0 < \lambda_2 < \infty$ such that $|\rho'(x_1) - \rho'(x_2)| \leq \lambda_2 |x_1 - x_2|$ for all $x_1, x_2 \in \mathbb{R}$.

If \mathcal{H} is a convex set, then the first inequality in A1 holds trivially. Note that under A2, the right-hand side of (30) will be finite for $k = 1$ whenever \mathcal{H} has bounded diameter and $\mathbf{E}_\mu \|\mathbf{L}'(h)\| < \infty$ for some $h \in \mathcal{H}$. As for A3, all the requirements

are clearly met by the Barron class with $-\infty \leq \alpha \leq 1$. These conditions imply a Lipschitz property for the gradients, as summarized in the following lemma.

Lemma 12. *Let the conditions A1, A2, and A3 hold. Then, the T-risk map $(h, \theta) \mapsto R_\rho(h; \theta, \eta)$ defined in (11) is smooth on $\mathcal{H} \times \mathbb{R}$ in the sense that*

$$\|R'_\rho(h_1; \theta_1, \eta) - R'_\rho(h_2; \theta_2, \eta)\| \leq \left(\frac{\lambda_5}{\sigma} + \frac{\lambda_2}{\sigma^2} \mathbf{E}_\mu \|L'\|_{\mathcal{H}} \right) (\|h_1 - h_2\| + |\theta_1 - \theta_2|)$$

for any choice of $h_1, h_2 \in \mathcal{H}$ and $\theta_1, \theta_2 \in \mathbb{R}$. Here the factor λ_5 is defined $\lambda_5 := \lambda_3 + \lambda_4$, where

$$\lambda_3 := \left(\frac{\lambda_2}{\sigma} \right) \left[\mathbf{E}_\mu \|L'\|_{\mathcal{H}}^2 + \sup_{h \in \mathcal{H}} \mathbf{E}_\mu \|L'(h)\| \right], \quad \lambda_4 := \lambda_1 \|\rho'\|_\infty.$$

Remark 13 (Norm on product spaces). In Lemma 12 we have to deal with norms on product spaces, and in all cases we just use the traditional choice of summing the norms of the constituent elements, i.e., on $\mathcal{H} \times \mathbb{R}$, we have $\|(h, \theta)\| := \|h\| + |\theta|$. Similarly, we have that the gradient $R'_\rho(h; \theta, \eta) = (\partial_h R_\rho(h; \theta, \eta), \partial_\theta R_\rho(h; \theta, \eta))$, a pair of linear functionals. As such, the norm of $R'_\rho(h; \theta, \eta)$ defined as the sum of the norms of these two constituent functionals.

Proving Lemma 12 is straightforward but somewhat tedious. Detailed computations as well as a direct proof are organized in §E.2 for easy reference.

D.5 Experimental details

Here we provide some additional details for the empirical analysis carried out in §3.3 and §C. Detailed hyperparameter settings and seeds for exact re-creation of all the results in this paper are available at the GitHub repository cited in §1, and thus here we will not explicitly write all the step sizes, shape settings, etc., but rather focus on concise exposition of points on which we expect readers to desire clarification.

Static risk analysis For our experiments in §3.3, we gave just one plot using a log-Normal distribution, but analogous tests were run for a wide variety of parametric distributions. In total, we have run tests for Bernoulli, Beta, χ^2 , Exponential, Gamma, log-Normal, Normal, Pareto, Uniform, Wald, and Weibull distributions. The settings of each distribution to be sampled from has not been tweaked at all; we set the parameters rather arbitrarily before running any tests. As for the fixed value of $\sigma = 0.5$ in the T-risk across all tests, we tested several values of σ ranging from 0.1 to 10, and based on the rough position of $\theta_\rho(L; \eta)$ in the Normal case, we determined 0.5 as a reasonable representative value; indeed, this settings performs quite well across a very wide range of distributions. Regarding the optimization involved in solving for the optimal threshold θ (for T-risk, OCE risks, and DRO risk), we use `minimize_scalar` from SciPy, with bounded solver type, and valid brackets set manually.

Noisy linear classification In the tests described in §5.1, we only give error/norm trajectories for “representative” settings of each risk class (Figure 4). For T-risk we consider different choices of $1 \leq \alpha \leq 2$, for CVaR we consider $0.025 \leq \beta \leq 0.75$, for tilted risk we consider γ between ± 0.05 , and for χ^2 -DRO we consider $0.025 \leq \tilde{\alpha} \leq 0.35$. For each of these ranges, we evaluate five evenly-spaced candidates (via `np.linspace`), and representative settings were selected as those which achieved the best classification error (average zero-one loss) after the final iteration. In the event of ties, the smaller hyperparameter value was always selected (via `np.argmax`).

Regression under outliers For the tests introduced in §C.1, we have given results for learning algorithms started at a point which is quite accurate for the majority of the data points, but incurs extremely large errors on the outlying minority. This choice of initial value naturally has a strong impact on the behavior of learning algorithms under each risk. For some perspective, in Figure 13 we provide results using a different initial value (again, in gray), which complement Figure 5. Since the naive strategy fixing θ at the initial median sets the scale extremely large and close to the loss incurred at *most* points, even a large number of gradient-based updates result in minimal change. The basic reason for this is quite straightforward. Since the T-risk gradient is modulated by $\rho'_\sigma(L(h) - \theta)$, and all points are such that either $L_i(h) \ll \theta$ (the minority) or $L_i(h) \approx \theta$ (the majority), both cases shrink the norm of the update direction vector $\alpha \leq 1$. When implementing T-risk in the more traditional way (jointly in (h, θ)) and choosing $\alpha \in [1, 2]$, we see results that are very similar to CVaR. Finally, we remark that here we have set the hyperparameter ranges with upper bounds low enough that the learning procedure described here does not run into overflow errors.

Classification under larger benchmark datasets In §C.2 we make use of several well-known benchmark datasets, and in our figures we identify them respectively by the following keywords: `adult`,⁸ `australian`,⁹ `cifar10`,¹⁰ `covtype`,¹¹ `emnist_balanced`,¹² `fashion_mnist`,¹³ and `protein`.¹⁴ For further background on all of these datasets, please access the URLs provided in the footnotes. As mentioned in the §C, for a k -class problem with d features, the predictor is characterized by k weighting vectors (w_1, \dots, w_k) , each of which is $w_j \in \mathbb{R}^d$ and computes scores as $w_j^\top x$ for $x \in \mathbb{R}^d$. These weight vectors are penalized using the usual multi-class logistic loss, namely the negative log-likelihood of the k -Categorical distribution that arises after passing these scores through the (logistic) softmax function. Regarding step sizes, we consider choices of factor $c \in \{0.1, 0.5, 1.0, 1.5, 2.0\}$, and set step size to c/\sqrt{kd} , where d and k are as just stated. In Figures 14–18 of §D.6, we provide additional results for the datasets not covered in Figures 6–7.

D.6 Additional figures

Here we include a number of figures that complement those provided in §C. A brief summary is given below.

- Figures 10–11 are additional results from the experiments described in §5.1, here using logistic and hinge losses instead of the unhinged loss.
- Figures 12–13 are related to the regression under outliers task described in §C.1. The first figure shows how different regression lines incur very different loss distributions under different convex base loss functions. The second figure illustrates how a different initial value impacts the learned regression lines under each risk class.
- Figures 14–18 are all completely analogous to Figure 6 given in §C.2, but for additional benchmark datasets.

⁸<https://archive.ics.uci.edu/ml/datasets/Adult>

⁹[https://archive.ics.uci.edu/ml/datasets/statlog+\(australian+credit+approval\)](https://archive.ics.uci.edu/ml/datasets/statlog+(australian+credit+approval))

¹⁰<https://www.cs.toronto.edu/~kriz/cifar.html>

¹¹<https://archive.ics.uci.edu/ml/datasets/coverttype>

¹²<https://www.nist.gov/itl/products-and-services/emnist-dataset>

¹³<https://github.com/zalandoresearch/fashion-mnist>

¹⁴<https://www.kdd.org/kdd-cup/view/kdd-cup-2004/Data>




Figure 10: Analogous to Figure 4, this time using logistic loss.




Figure 11: Analogous to Figure 4, this time using the hinge loss.




Figure 12: Here the left-most plot is the Belgian phone call dataset for use in a one-dimensional regression task, with two linear regression candidates. The remaining three plots show histograms of the loss distributions incurred by each of these candidates using three loss functions commonly used in regression.




Figure 13: Same procedure as in Figure 5, but this time started from a different initial value.

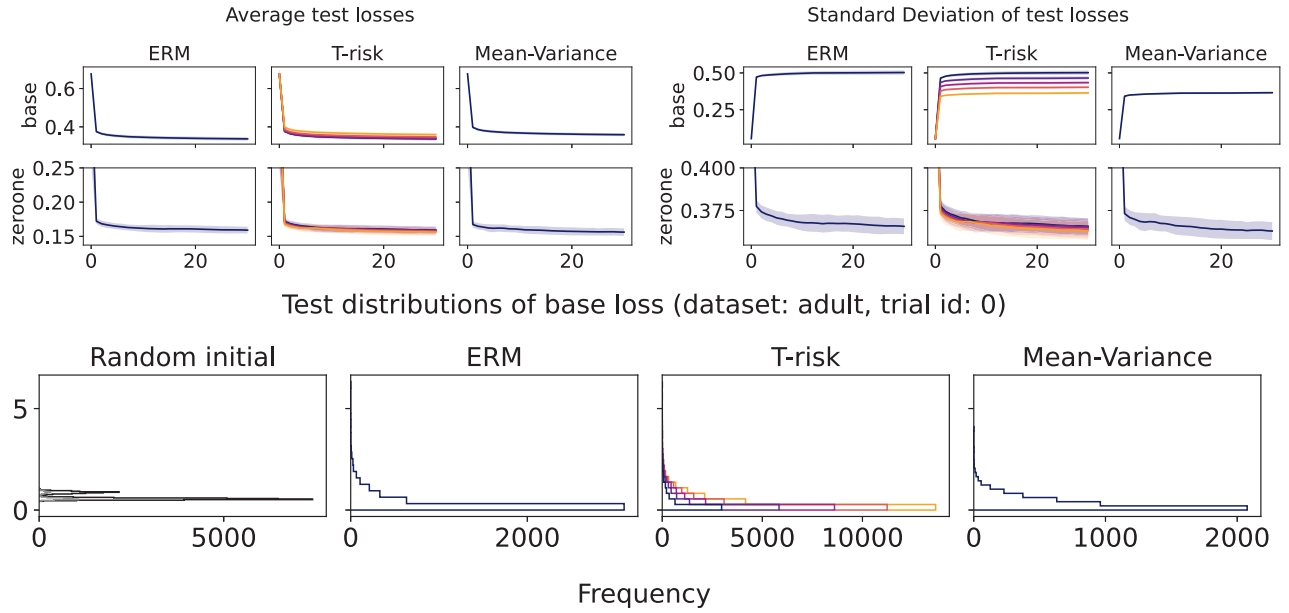


Figure 14: Using T-risk to interpolate between test loss distributions (dataset: adult).

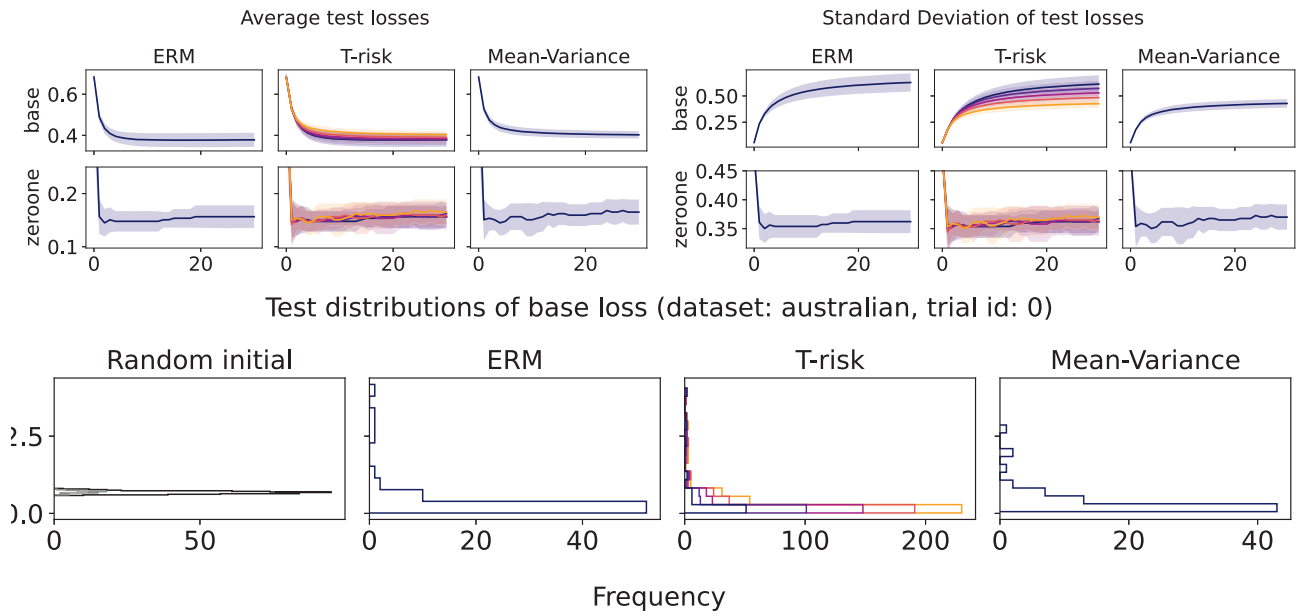


Figure 15: Using T-risk to interpolate between test loss distributions (dataset: australian).

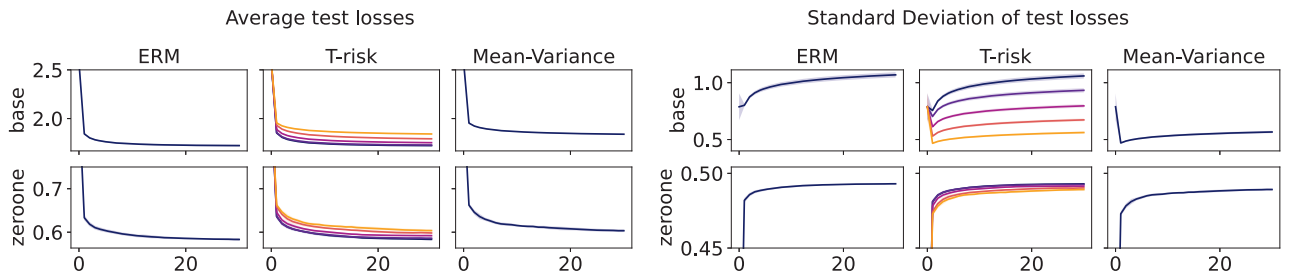


Figure 16: Using T-risk to interpolate between test loss distributions (dataset: cifar10).

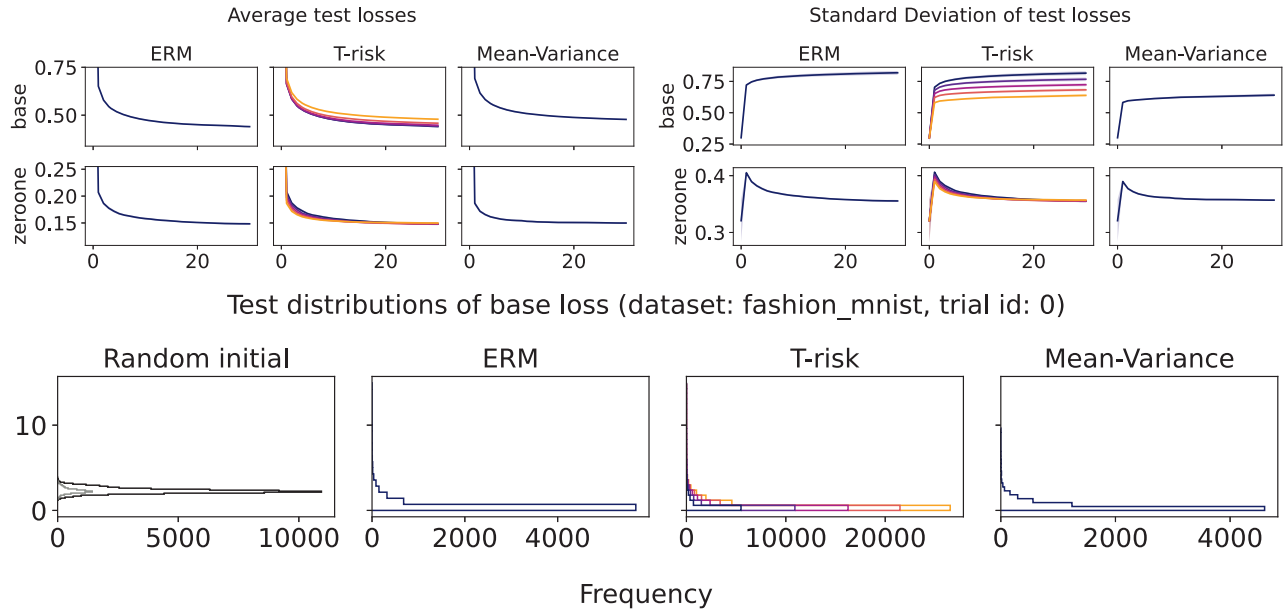


Figure 17: Using T-risk to interpolate between test loss distributions (dataset: `fashion_mnist`).

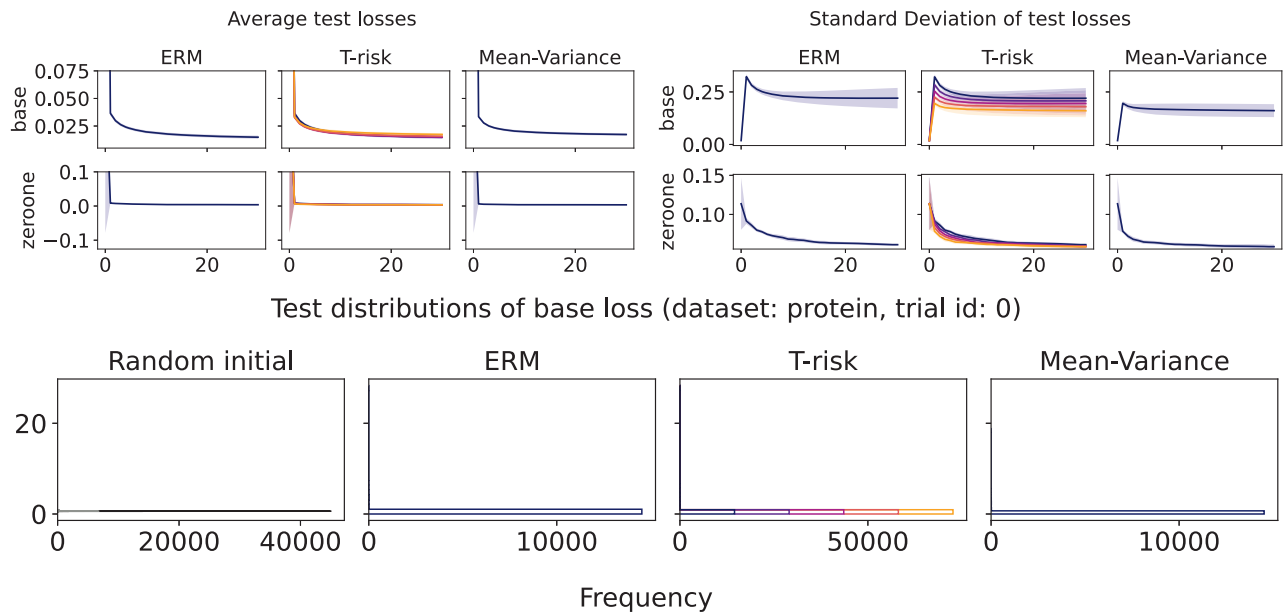


Figure 18: Using T-risk to interpolate between test loss distributions (dataset: `protein`).

E Detailed proofs

E.1 Proofs of results in the main text

Proof of Lemma 8. In this proof, without further mention, we will make regular use of the following two helper results: Lemma 14 (bounded gradient implies Lipschitz continuity) and Lemma 15 (positive definite Hessian implies convexity). For reference, the first and second derivatives of ρ_σ are given in §F.3. We take up each α setting one at a time.

First, the case of $\alpha = 2$. For this case, clearly ρ'_σ is unbounded, and thus ρ_σ is not (globally) Lipschitz on \mathbb{R} . On the other hand, since $\rho''_\sigma(x) = 1/\sigma^2$, we have that ρ'_σ is λ -Lipschitz with $\lambda = (1/\sigma^2)$.

Next, the case of $\alpha = 0$. For any fixed $\sigma > 0$, in both the limits $x \rightarrow 0$ and $|x| \rightarrow \infty$, we have $\rho'_\sigma(x) \rightarrow 0$. Maximum and minimum values are achieved when $\rho''_\sigma(x) = 0$, and this occurs if and only if $x^2 = 2\sigma^2$. It follows from direct computation that $\rho'_\sigma(\pm\sqrt{2}\sigma) = \pm 1/(\sqrt{2}\sigma)$, and thus ρ_σ is λ -Lipschitz with $\lambda = 1/(\sqrt{2}\sigma)$. Next, recalling that ρ''_σ takes the form

$$\rho''_\sigma(x) = \frac{2}{x^2 + 2\sigma^2} \left(1 - \frac{2x^2}{x^2 + 2\sigma^2} \right),$$

we see that this is a product of two factors, one taking values in $(0, 1/\sigma^2)$, and one taking values in $(-1, 1)$. The absolute value of both of these factors is maximized when $x = 0$, and so $|\rho''_\sigma(x)| \leq \rho''_\sigma(0) = 1/\sigma^2$, meaning that ρ'_σ is λ -Lipschitz with $\lambda = 1/\sigma^2$. Finally, regarding convexity, we have that $\rho''_\sigma(x) \geq 0$ if and only if $|x| \leq \sqrt{2}\sigma$.

Next, the case of $\alpha = -\infty$. For any fixed $\sigma > 0$, we have $\rho'_\sigma(x) \rightarrow 0$ in both the limits $x \rightarrow 0$ and $|x| \rightarrow \infty$. Furthermore, it is immediate that $\rho''_\sigma(x) = 0$ at the points $x = \pm\sigma$. Evaluating ρ'_σ at these stationary points we have $\rho'_\sigma(\pm\sigma) = \pm(1/\sigma) \exp(-1/2)$, and thus ρ_σ is λ -Lipschitz with $\lambda = (1/\sigma) \exp(-1/2)$. Regarding bounds on ρ''_σ , first note that $\rho''_\sigma(x) \rightarrow 0$ as $|x| \rightarrow \infty$, and $\rho''_\sigma(0) = 1/\sigma^2$. Then to identify stationary points, note that

$$\rho'''_\sigma(x) = \frac{1}{\sigma^2} \exp\left(-\frac{1}{2} \left(\frac{x}{\sigma}\right)^2\right) \left[\frac{x}{\sigma^2} \left(\frac{x^2}{\sigma^2} - 1\right) - \frac{2x}{\sigma^2} \right]$$

and thus $\rho'''_\sigma(x) = 0$ if and only if $(x/\sigma)^2 - 1 = 2$, i.e., the stationary points are $x = \pm\sqrt{3}\sigma$, both of which yield the same value, namely $\rho''_\sigma(\pm\sqrt{3}\sigma) = -(2/\sigma^2) \exp(-3/2)$. Since $2 \exp(-3/2) \approx 0.45 < 1$, we conclude that ρ'_σ is λ -Lipschitz with $\lambda = 1/\sigma^2$. Finally, since $\rho''_\sigma(x) \geq 0$ if and only if $|x| \leq \sigma$, this specifies the region on which ρ_σ is convex.

Finally, all that remains is the general case of $-\infty < \alpha < 2$ where $\alpha \neq 0$. Note that in order for $\rho''_\sigma(x) = 0$ to hold, we require

$$\frac{2(x/\sigma)^2}{1 + (x/\sigma)^2/|\alpha - 2|} = \frac{|\alpha - 2|}{1 - (\alpha/2)},$$

which via some basic algebra is equivalent to

$$\left(\frac{x}{\sigma}\right)^2 = \frac{|\alpha - 2|}{1 - \alpha}.$$

Clearly, this is only possible when $\alpha < 1$, so we consider this sub-case first. This implies stationary points $\pm x^* := \pm\sigma\sqrt{|\alpha - 2|/(1 - \alpha)}$, for which we have

$$\rho'_\sigma(\pm x^*) = \pm \frac{1}{\sigma} \sqrt{\frac{|\alpha - 2|}{1 - \alpha}} \left(\frac{2 - \alpha}{1 - \alpha}\right)^{(\alpha/2) - 1} = \pm \frac{1}{\sigma} \left(\sqrt{\frac{|\alpha - 2|}{1 - \alpha}}\right)^{\alpha - 1} = \pm \frac{1}{\sigma} \left(\sqrt{\frac{1 - \alpha}{|\alpha - 2|}}\right)^{1 - \alpha}.$$

Since $\rho'_\sigma(x) \rightarrow 0$ in both the limits $x \rightarrow 0$ and $|x| \rightarrow \infty$, we have obtained a maximum value for ρ'_σ at x^* , thus implying for the case of $\alpha < 1$ that ρ_σ is λ -Lipschitz, with a coefficient of $\lambda = (1/\sigma)(\sqrt{(1 - \alpha)/|\alpha - 2|})^{1 - \alpha}$. For the case of $\alpha = 1$, direct inspection shows

$$|\rho'_\sigma(x)| = \frac{|x/\sigma^2|}{\sqrt{1 + (x/\sigma)^2}} = \frac{1}{\sigma^2 \sqrt{(1/x^2) + (1/\sigma^2)}},$$

a value which is maximized in the limit $|x| \rightarrow \infty$. As such, for $\alpha = 1$, we have that ρ_σ is λ -Lipschitz with $\lambda = 1/\sigma$. For the case of $1 < \alpha < 2$, ρ'_σ is unbounded. To see this, note that for $x > 0$ we have

$$\rho'_\sigma(x) = \frac{1}{\sigma^2} \frac{(1 + (x/\sigma)^2/|\alpha - 2|)}{((1/x) + x/(\sigma^2|\alpha - 2|))} \geq \left(\frac{1}{\sigma^{2+\alpha} \sqrt{|\alpha - 2|^\alpha}}\right) \frac{x^\alpha}{((1/x) + x/(\sigma^2|\alpha - 2|))},$$

and since $\alpha > 1$, sending $x \rightarrow \infty$ clearly implies $\rho'_\sigma(x) \rightarrow \infty$, and this means that ρ_σ cannot be Lipschitz on \mathbb{R} when $\alpha > 1$. As for bounds on ρ''_σ , recall that

$$\rho''_\sigma(x) = \frac{1}{\sigma^2} \underbrace{\left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{(\alpha/2)-1}}_{A(x)} \left(1 - \underbrace{\frac{1 - (\alpha/2)}{|\alpha - 2|} \frac{2(x/\sigma)^2}{1 + (x/\sigma)^2/|\alpha - 2|}}_{B(x)}\right)$$

where we have introduced the labels $A(x)$ and $B(x)$ just as convenient notation. Fixing any $\sigma > 0$, first note that since $\alpha < 2$, we have $(\alpha/2) - 1 < 0$ and thus $0 \leq A(x) \leq 1$. Next, direct inspection shows $0 \leq B(x) \leq 2(1 - (\alpha/2))$. These two facts immediately imply an upper bound $\rho''_\sigma(x) \leq 1/\sigma^2$ and a lower bound $\rho''_\sigma(x) \geq -(1 - \alpha)/\sigma^2$, both of which hold for any $\alpha < 2$. Furthermore, for the case of $1 \leq \alpha < 2$, we thus have $0 \leq \rho''_\sigma(x) \leq 1/\sigma^2$. When $\alpha < 1$ however, ρ''_σ can be negative. To get matching lower bounds requires $A(x)(1 - B(x)) \geq -1$, or $A(x)(B(x) - 1) \leq 1$. To study conditions under which this holds, first note that $B(x)$ can be re-written as

$$B(x) = \left(\frac{2 - \alpha}{|\alpha - 2|}\right) \frac{(x/\sigma)^2}{1 + \frac{(x/\sigma)^2}{|\alpha - 2|}} = \frac{(x/\sigma)^2}{1 + \frac{(x/\sigma)^2}{|\alpha - 2|}},$$

and thus we have

$$A(x)(B(x) - 1) = \frac{(x/\sigma)^2}{\left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{2 - (\alpha/2)}} - \frac{1}{\left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{1 - (\alpha/2)}}. \quad (31)$$

To get a more convenient upper bound on this, observe that $(1+x)^{1-(\alpha/2)} \leq (1+x)^{2-(\alpha/2)}$ for any $x \geq 0$ and $-\infty \leq \alpha \leq 2$. It follows immediately that

$$A(x)(B(x) - 1) \leq \frac{(x/\sigma)^2 - 1}{\left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{2 - (\alpha/2)}}. \quad (32)$$

To get the right-hand side of (32) to be no greater than 1 is equivalent to

$$(x/\sigma)^2 - 1 \leq \left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{2 - (\alpha/2)}. \quad (33)$$

For the case of $0 \leq \alpha < 1$, note that $1 \leq |\alpha - 2| = 2 - \alpha < 2 - (\alpha/2)$, and using the helper inequality (69), we have

$$\left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{2 - (\alpha/2)} \geq \left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{|\alpha - 2|} \geq 1 + (x/\sigma)^2 > (x/\sigma)^2 - 1,$$

which implies (33) for $0 \leq \alpha < 1$. All that remains is the case of $-\infty < \alpha < 0$, which requires a bit more care. Returning to the exact form of $A(x)(B(x) - 1)$ given in (31), note that the inequality

$$(x/\sigma)^2 - \left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right) \leq \left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{2 - (\alpha/2)} \quad (34)$$

is equivalent to the desired property, i.e., $A(x)(B(x) - 1) \leq 1 \iff$ (34). Using Bernoulli's inequality (71), we can bound the right-hand side of (34) as

$$\left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{2 - (\alpha/2)} \geq 1 + \left(\frac{2 - (\alpha/2)}{|\alpha - 2|}\right) (x/\sigma)^2.$$

Subtracting the left-hand side of (34) from the right-hand side of the preceding inequality, we obtain

$$\begin{aligned} \left(1 + \frac{(x/\sigma)^2}{|\alpha - 2|}\right)^{2 - (\alpha/2)} - \left[(x/\sigma)^2 - 1 - \frac{(x/\sigma)^2}{|\alpha - 2|}\right] &\geq 2 + \left(\frac{2 - (\alpha/2)}{|\alpha - 2|} - 1 + \frac{1}{|\alpha - 2|}\right) (x/\sigma)^2 \\ &= 2 + \left(\frac{1 - (|\alpha|/2)}{2 + |\alpha|}\right) (x/\sigma)^2, \end{aligned} \quad (35)$$

where the second step uses the fact that for $\alpha < 0$, we can write $|\alpha - 2| = 2 + |\alpha|$ and $2 - (\alpha/2) = 2 + (|\alpha|/2)$. Note that the right-hand side of (35) is non-negative for all $x \in \mathbb{R}$ whenever $-2 \leq \alpha < 0$, which via (34) tells us that $A(x)(B(x) - 1) \leq 1$ indeed holds in this case as well. For the case of $-\infty < \alpha < -2$, note that showing (34) holds is equivalent to showing $f_\alpha(x) \geq 0$ for all $x \geq 0$, where for convenience we define the polynomial

$$f_\alpha(x) := 1 + \left(\frac{1}{2 + |\alpha|} - 1 \right) x + \left(1 + \frac{x}{2 + |\alpha|} \right)^{2 + (|\alpha|/2)}.$$

The first derivative is

$$f'_\alpha(x) = \frac{2 + (|\alpha|/2)}{2 + |\alpha|} \left(1 + \frac{x}{2 + |\alpha|} \right)^{1 + (|\alpha|/2)} + \frac{1}{2 + |\alpha|} - 1,$$

and with this form in hand, solving for $f_\alpha(x) = 0$, it is straightforward to confirm that x_α^* given below is a stationary point:

$$x_\alpha^* := (2 + |\alpha|) \left[\left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{1}{1 + (|\alpha|/2)}} - 1 \right].$$

Furthermore, it is clear that $f''_\alpha \geq 0$, implying that f_α is convex, and that x_α^* is a minimum. As such, the minimum value taken by f_α on \mathbb{R}_+ is

$$\begin{aligned} f_\alpha(x_\alpha^*) &= \left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{1}{1 + (|\alpha|/2)}} - (2 + |\alpha|) \left[\left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{1}{1 + (|\alpha|/2)}} - 1 \right] + \left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{2 + (|\alpha|/2)}{1 + (|\alpha|/2)}} \\ &= (2 + |\alpha|) + \left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{2 + (|\alpha|/2)}{1 + (|\alpha|/2)}} - (1 + |\alpha|) \left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{1}{1 + (|\alpha|/2)}} \\ &= (2 + |\alpha|) - (1 + |\alpha|) \left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{1}{1 + (|\alpha|/2)}} \left[1 - \frac{1}{2 + (|\alpha|/2)} \right] \\ &= 1 + (1 + |\alpha|) \left[1 - \left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{1}{1 + (|\alpha|/2)}} \left[1 - \frac{1}{2 + (|\alpha|/2)} \right] \right]. \end{aligned}$$

We require $f_\alpha(x_\alpha^*) \geq 0$ for all $-\infty < \alpha < -2$. From the preceding equalities, note that a simple sufficient condition for $f_\alpha(x_\alpha^*) \geq 1$ is

$$\left(\frac{1 + |\alpha|}{2 + (|\alpha|/2)} \right)^{\frac{1}{1 + (|\alpha|/2)}} \left[1 - \frac{1}{2 + (|\alpha|/2)} \right] \leq 1$$

or equivalently

$$\left(1 - \frac{1}{2 + (|\alpha|/2)} \right)^{1 + (|\alpha|/2)} \leq \left(\frac{2 + (|\alpha|/2)}{1 + |\alpha|} \right). \quad (36)$$

Applying the helper inequality (70) to the left-hand side of (36), we have

$$\begin{aligned} \left(1 - \frac{1}{2 + (|\alpha|/2)} \right)^{1 + (|\alpha|/2)} &\leq 1 + \frac{\left(\frac{-(1 + (|\alpha|/2))}{2 + (|\alpha|/2)} \right)}{\left(1 - \frac{-|\alpha|/2}{2 + (|\alpha|/2)} \right)} = 1 - \left(\frac{1 + (|\alpha|/2)}{2 + |\alpha|} \right) = \frac{1 + (|\alpha|/2)}{2 + |\alpha|} \\ &\leq \frac{2 + (|\alpha|/2)}{1 + |\alpha|}. \end{aligned}$$

This is precisely the desired inequality (36), implying $f_\alpha(x_\alpha^*) \geq 1 > 0$ for all $-\infty < \alpha < -2$, and in fact all real $\alpha < 0$. To summarize, we have $A(x)(B(x) - 1) \leq 1$ for all $x \in \mathbb{R}$, and thus the desired $1/\sigma^2$ -smoothness result follows, concluding the proof. \square

Proof of Lemma 1. Let X denote any \mathcal{F} -measurable random variable. The continuity of ρ implies that the integral $\mathbf{E}_\mu \rho_\sigma(X - \theta)$ exists for any $\sigma > 0$; we just need to prove it is finite.¹⁵ Since we are taking ρ from the Barron class (15), we consider each α setting separately. Starting with $\alpha = 2$, note that

$$\rho_\sigma(X - \theta; 2) = \frac{1}{2} \left(\frac{X - \theta}{\sigma} \right)^2$$

and thus $\mathbf{E}_\mu X^2 < \infty$ is sufficient and necessary. For $\alpha = 0$, first note that we have

$$\rho_\sigma(X - \theta; 0) = \log \left(1 + \frac{1}{2} \left(\frac{X - \theta}{\sigma} \right)^2 \right).$$

Let $f_1(x) := \log(1 + x)$ and $f_2(x) := x^c/c$, where $0 < c < 1$. Note that $f_1(0) = f_2(0) = 0$, and furthermore that for any $x > 0$,

$$f_1'(x) = \frac{1}{1+x} < \left(\frac{1}{1+x} \right)^{1-c} < \left(\frac{1}{x} \right)^{1-c} = f_2'(x).$$

We may thus conclude that $f_1(x) \leq f_2(x)$ for all $x \geq 0$, and thus for any $0 < c < 1$ we have

$$\log \left(1 + \frac{1}{2} \left(\frac{X - \theta}{\sigma} \right)^2 \right) \leq \frac{1}{c} \left(\frac{X - \theta}{\sqrt{2}\sigma} \right)^{2c}.$$

It follows that to ensure $\mathbf{E}_\mu \rho_\sigma(X - \theta; 0) < \infty$, it is sufficient if we assume $\mathbf{E}_\mu |X|^c < \infty$ for some $c > 0$. Proceeding to the case of $\alpha = -\infty$, we have

$$\rho_\sigma(X - \theta; -\infty) = 1 - \exp \left(-\frac{1}{2} \left(\frac{X - \theta}{\sigma} \right)^2 \right).$$

Any composition of measurable functions is measurable, and since the right-hand side is bounded above by 1 and below by 0, we have that $\rho_\sigma(X - \theta; -\infty)$ is μ -integrable without requiring any extra assumptions on X besides measurability. All that remains for the Barron class is the case of non-zero $-\infty < \alpha < 2$, where we have

$$\rho_\sigma(X - \theta; \alpha) = \frac{|\alpha - 2|}{\alpha} \left(\left(1 + \frac{1}{|\alpha - 2|} \left(\frac{X - \theta}{\sigma} \right)^2 \right)^{\alpha/2} - 1 \right).$$

Let us break this into two cases: $-\infty < \alpha < 0$ and $0 < \alpha < 2$. Starting with the former case, this is easy since

$$(1 + x^2)^{\alpha/2} = \frac{1}{(\sqrt{1 + x^2})^{-\alpha}}$$

which is bounded above by 1 and below by 0 for any $\alpha < 0$ and $x \in \mathbb{R}$, which means the random variable $\rho_\sigma(X - \theta; \alpha)$ is μ -integrable without any extra assumptions on X . As for the latter case of $0 < \alpha < 2$, first note that the monotonicity of $(\cdot)^{\alpha/2}$ on \mathbb{R}_+ implies

$$(1 + x^2)^{\alpha/2} \geq |x|^\alpha$$

which means $\mathbf{E}_\mu |X|^\alpha < \infty$ is necessary. That this condition is also sufficient is immediate from the form of $\rho_\sigma(X - \theta; \alpha)$ just given. This concludes the proof; the desired result stated in the lemma follows from setting $X = L(h)$ and the observation that the choice of $\theta \in \mathbb{R}$ in the preceding discussion was arbitrary. \square

Proof of Lemma 9. Referring to the derivatives (67)–(68) in §F.3, we know that ρ'_σ is measurable, and by the proof of Lemma 8, we know that $\|\rho'\|_\infty < \infty$ for all $\alpha \leq 1$. Thus, as long as $L(h)$ is \mathcal{F} -measurable, we have that $\rho'_\sigma(L(h) - \theta)$ is μ -integrable. For the case of $1 < \alpha \leq 2$, note that $|\rho'_\sigma(x)| \leq |x|/\sigma^2$ holds, meaning that $\mathbf{E}_\mu |L(h)| < \infty$ implies integrability.

¹⁵This uses the fact that any composition of (Borel) measurable functions is itself measurable (Ash and Doléans-Dade, 2000, Lem. 1.5.7).

Similarly for the second derivatives, from the proof of Lemma 8, we see that $\|\rho''\|_\infty < \infty$ for all $-\infty \leq \alpha \leq 1$, implying the μ -integrability of $\rho''_\sigma(\mathbf{L}(h) - \theta)$.

The Leibniz integration property follows using a straightforward dominated convergence argument, which we give here for completeness. Letting (a_k) be any non-zero real sequence such that $a_k \rightarrow 0$, we can write

$$\begin{aligned} \frac{d}{d\theta} D_\rho(h; \theta) &= \lim_{k \rightarrow \infty} \frac{D_\rho(h; \theta + a_k) - D_\rho(h; \theta)}{a_k} \\ &= \lim_{k \rightarrow \infty} \mathbf{E}_\mu \left[\frac{\rho_\sigma(\mathbf{L}(h) - (\theta + a_k)) - \rho_\sigma(\mathbf{L}(h) - \theta)}{a_k} \right]. \end{aligned}$$

For notational convenience, let us denote the key sequence of functions by

$$f_k := \frac{\rho_\sigma(\mathbf{L}(h) - (\theta + a_k)) - \rho_\sigma(\mathbf{L}(h) - \theta)}{a_k}$$

and note that $f_k \rightarrow f := -\rho'_\sigma(\mathbf{L}(h) - \theta)$ pointwise as $k \rightarrow \infty$. We can then say the following: for all k , we have that

$$|f_k| \leq \sup_{0 < c < 1} |\rho'_\sigma(\mathbf{L}(h) - (\theta + ca_k))| \leq g := |\rho'_\sigma(\mathbf{L}(h) - \theta')|$$

for an appropriate choice of $\theta' \in \mathbb{R}$. The first inequality follows from the helper Lemma 14. We can always find an appropriate θ' because the sequence (a_k) is bounded and ρ' is eventually monotone, regardless of the choice of α . With the fact $|f_k| \leq g$ in hand, recall that we have already proved that $\mathbf{E}_\mu g < \infty$ under the assumptions we have made, and thus $\mathbf{E}_\mu f_k \rightarrow \mathbf{E}_\mu f$ by dominated convergence.¹⁶ As such, we have

$$\frac{d}{d\theta} D_\rho(h; \theta) = \lim_{k \rightarrow \infty} \mathbf{E}_\mu f_k = \mathbf{E}_\mu f = -\mathbf{E}_\mu \rho'_\sigma(\mathbf{L}(h) - \theta),$$

which is the desired Leibniz property for the first derivative. A completely analogous argument holds for the second derivative, yielding the desired result. \square

Proof of Lemma 10. From Lemma 9, we know that the map $\theta \mapsto D_\rho(h; \theta)$ is differentiable and thus continuous. Using continuity, taking any $a < b$ and constructing a closed interval $[a, b]$, the Weierstrass extreme value theorem tells us that $D_\rho(h; \theta)$ achieves its maximum and minimum on $[a, b]$. Furthermore, note that ρ taken from the Barron class (15) satisfies all the requirements of our helper Lemma 16, and thus implies $D_\rho(h; \theta) \rightarrow \sup\{\rho(x) : x \in \mathbb{R}\}$ as $|\theta| \rightarrow \infty$. We can thus always take the interval $[a, b]$ wide enough that

$$\theta \notin [a, b] \implies D_\rho(h; \theta) \geq \max_{a \leq x \leq b} D_\rho(h; x) \geq \min_{a \leq x \leq b} D_\rho(h; x).$$

This proves the existence of a minimizer of $\theta \mapsto D_\rho(h; \theta)$ on \mathbb{R} .

Next, considering the T-risk $R_\rho(h; \theta, \eta)$ and minimization with respect to θ , since we are doing unconstrained optimization, any solution $\theta_\rho(h; \eta)$ must satisfy $\eta + d_\theta D_\rho(h; \theta_\rho(h; \eta)) = 0$, where $d_\theta := d/d\theta$. Using Lemma 9 again, this can be equivalently re-written as

$$\mathbf{E}_\mu \rho'_\sigma(\mathbf{L}(h) - \theta_\rho(h; \eta)) = \eta. \tag{37}$$

When $\alpha > 1$, the derivative of the dispersion function has unbounded range, i.e., $\rho'_\sigma(\mathbb{R}) = \mathbb{R}$. As such, an argument identical to that used in the proof of Lemma 16 implies that for any $\eta \in \mathbb{R}$, we can always find a $\theta_\eta(h) \in \mathbb{R}$ such that (37) holds, recalling that continuity follows via Lemma 9. Combining this with convexity gives us a valid solution. The special case of $\alpha = 1$ requires additional conditions, since from Lemma 8, we know that in this case $|\rho'_\sigma| \leq 1/\sigma$, and thus by an analogous argument, whenever $|\theta_\rho(h; \eta)| < 1/\sigma$ we can find a finite solution.

To prove uniqueness under $1 \leq \alpha \leq 2$, direct inspection of the second derivative in (68) shows us that $\rho''_\sigma(x) > 0$ on \mathbb{R} whenever we have

$$\frac{2(x/\sigma)^2}{\left(1 + \frac{(x/\sigma)^2}{|\alpha-2|}\right)} < \frac{|\alpha-2|}{1 - (\alpha/2)}.$$

¹⁶See for example Ash and Doléans-Dade (2000, Thm. 1.6.9).

Re-arranging the above inequality yields an equivalent condition of $(x/\sigma)^2(1-\alpha) < |\alpha-2|$, a condition which holds on \mathbb{R} if and only if $1 \leq \alpha \leq 2$. Since ρ''_σ is positive on \mathbb{R} , this implies that $\mathbf{E}_\mu \rho''_\sigma(\mathbf{L}(h) - \theta) > 0$ for all $\theta \in \mathbb{R}$. Using Lemma 9, we have that $\mathbf{E}_\mu \rho''_\sigma(\mathbf{L}(h) - \theta)$ is equal to the second derivative of $D_\rho(h; \theta)$ with respect to θ , which implies that $\theta \mapsto D_\rho(h; \theta)$ and $\theta \mapsto D_\rho(h; \theta) + \eta\theta$ are strictly convex on \mathbb{R} , and thus their minimum must be unique.¹⁷ \square

Proof of Lemma 5. For random loss \mathbf{L} , using Lemma 9, first-order optimality conditions require

$$\mathbf{E}_\mu \rho'_\sigma(\mathbf{L} - M_\rho(\mathbf{L})) = 0, \quad \mathbf{E}_\mu \rho'_\sigma(\mathbf{L} - \theta_\rho(\mathbf{L}; \eta)) = \eta. \quad (38)$$

If these conditions hold, then from direct inspection, the same conditions will clearly hold if we replace \mathbf{L} by $\mathbf{L} + c$, $M_\rho(\mathbf{L})$ by $M_\rho(\mathbf{L}) + c$, and $\theta_\rho(\mathbf{L}; \eta)$ by $\theta_\rho(\mathbf{L}; \eta) + c$. This implies both translation-invariance of the dispersions and the translation-equivariance of the optimal thresholds. Non-negativity follows trivially from the fact that $\rho(\cdot) \geq 0$. Noting that $\rho(x) > 0$ for all $x \neq 0$, we have that $D_\rho(\mathbf{L}; \theta) = 0$ if and only if $\mathbf{L} = \theta$ almost surely.¹⁸ Since $D_\rho(\mathbf{L}; M_\rho(\mathbf{L})) \leq D_\rho(\mathbf{L}; \theta_\rho(\mathbf{L}; \eta))$ by the optimality of $M_\rho(\mathbf{L})$, it follows that for any non-constant \mathbf{L} , we must have $D_\rho(\mathbf{L}; \theta_\rho(\mathbf{L}; \eta)) > 0$. Furthermore, from the optimality condition (38) for $\theta_\rho(\mathbf{L}; \eta)$, even when \mathbf{L} is constant, we must have $D_\rho(\mathbf{L}; \theta_\rho(\mathbf{L}; \eta)) > 0$ whenever $\eta \neq 0$, since $\rho'_\sigma(x) = 0$ if and only if $x = 0$.

In the special case where $1 \leq \alpha \leq 2$, we have that ρ''_σ is positive on \mathbb{R} (see §F.3 and Fig. 8). This implies that ρ_σ is strictly convex, and ρ'_σ is monotonically increasing. Let $\mathbf{L}_1 \leq \mathbf{L}_2$ almost surely, but say $M_\rho(\mathbf{L}_1) > M_\rho(\mathbf{L}_2)$. Using the optimality condition (38), uniqueness of the solution via Lemma 10, and the aforementioned monotonicity of ρ'_σ , we have

$$0 = \mathbf{E}_\mu \rho'_\sigma(\mathbf{L}_1 - M_\rho(\mathbf{L}_1)) < \mathbf{E}_\mu \rho'_\sigma(\mathbf{L}_1 - M_\rho(\mathbf{L}_2)) \leq \mathbf{E}_\mu \rho'_\sigma(\mathbf{L}_2 - M_\rho(\mathbf{L}_2)) = 0.$$

This is a contradiction, and thus we must have $M_\rho(\mathbf{L}_1) \leq M_\rho(\mathbf{L}_2)$. An identical argument using the exact same properties proves that $\theta_\rho(\mathbf{L}_1; \eta) \leq \theta_\rho(\mathbf{L}_2; \eta)$ also holds. Finally, to prove convexity, take any $\mathbf{L}_1, \mathbf{L}_2 \in \mathcal{L}$, $\theta_1, \theta_2 \in \mathbb{R}$, and $a \in (0, 1)$, and note that

$$\begin{aligned} \underline{\mathbf{R}}_\rho(a\mathbf{L}_1 + (1-a)\mathbf{L}_2; \eta) &\leq D_\rho(a\mathbf{L}_1 + (1-a)\mathbf{L}_2; a\theta_1 + (1-a)\theta_2) + \eta(a\theta_1 + (1-a)\theta_2) \\ &= \mathbf{E}_\mu \rho_\sigma(a(\mathbf{L}_1 - \theta_1) + (1-a)(\mathbf{L}_2 - \theta_2)) + \eta(a\theta_1 + (1-a)\theta_2) \\ &\leq a(D_\rho(\mathbf{L}_1; \theta_1) + \eta\theta_1) + (1-a)(D_\rho(\mathbf{L}_2; \theta_2) + \eta\theta_2). \end{aligned}$$

The first inequality uses optimality of the threshold in the definition of $\underline{\mathbf{R}}_\rho$, whereas the second inequality uses the convexity of ρ_σ . Since the choice of θ_1 and θ_2 here were arbitrary, we can set $\theta_1 = \theta_\rho(\mathbf{L}_1; \eta)$ and $\theta_2 = \theta_\rho(\mathbf{L}_2; \eta)$ to obtain the desired inequality

$$\underline{\mathbf{R}}_\rho(a\mathbf{L}_1 + (1-a)\mathbf{L}_2; \eta) \leq a\underline{\mathbf{R}}_\rho(\mathbf{L}_1; \eta) + (1-a)\underline{\mathbf{R}}_\rho(\mathbf{L}_2; \eta)$$

giving us convexity of the threshold risk. As a direct corollary, setting $\eta = 0$ yields the convexity result for $\mathbf{L} \mapsto D_\rho(\mathbf{L}; M_\rho(\mathbf{L}))$. \square

Proof of Lemma 2. The crux of this result is an analogue to Lemma 9 regarding the differentials of $D_\rho(h; \theta)$, this time taken with respect to h , rather than θ . Fixing arbitrary $g, h \in \mathcal{H}$, let us start by considering the following sequence of random variables:

$$f_k := \frac{\rho_\sigma(\mathbf{L}(h + a_k g) - \theta) - \rho_\sigma(\mathbf{L}(h) - \theta)}{a_k} \quad (39)$$

where (a_k) is any sequence of real values such that $a_k \rightarrow 0_+$ as $k \rightarrow \infty$. Before getting into the details, let us unpack the differentiability assumption made on the base loss. Before random sampling, the map $h \mapsto \mathbf{L}(h)$ is of course a map from \mathcal{H} to the set of measurable functions $\{\mathbf{L}(h) : h \in \mathcal{H}\}$, but after sampling, there is no randomness and it is simply a map from \mathcal{H} to \mathbb{R} . Having sampled the random loss, the property we desire is that for each $h \in \mathcal{H}$, there exists a continuous linear functional $\mathbf{L}'(h) : \mathcal{U} \rightarrow \mathbb{R}$ such that

$$\lim_{\|g\| \rightarrow 0} \frac{|\mathbf{L}(h + g) - \mathbf{L}(h) - \mathbf{L}'(h)(g)|}{\|g\|} = 0. \quad (40)$$

¹⁷See for example Boyd and Vandenberghe (2004, Sec. 3.1.4).

¹⁸This fact follows from basic Lebesgue integration theory (Ash and Doléans-Dade, 2000, Thm. 1.6.6).

The differentiability condition in the lemma statement is simply that

$$\mu\{\text{equality (40) holds}\} = 1. \quad (41)$$

On this “good” event, since the map $x \mapsto \rho_\sigma(x)$ is differentiable by definition, we have that the composition $h \mapsto \rho_\sigma(\mathbf{L}(h) - \theta)$ is also differentiable for any choice of $\theta \in \mathbb{R}$, and a general chain rule can be applied to compute the differentials.¹⁹ In particular, we have a pointwise limit of

$$f := \lim_{k \rightarrow \infty} f_k = \rho'_\sigma(\mathbf{L}(h) - \theta)\mathbf{L}'(h)(g) \quad (42)$$

which also uses the fact that the Fréchet and Gateaux differentials are equal here.²⁰ Technically, it just remains to obtain conditions which imply $\mathbf{E}_\mu f_k \rightarrow \mathbf{E}_\mu f$. In pursuit of a μ -integrable upper bound on the sequence (f_k) , note that for large enough k , we have

$$\begin{aligned} |f_k| &\leq \frac{1}{a_k} \|a_k g\| \sup_{0 < a < a_k} \|\rho'_\sigma(\mathbf{L}(h + ag) - \theta)\mathbf{L}'(h + ag)\| \\ &\leq \|g\| \sup \{ \|\rho'_\sigma(\mathbf{L}(h_0) - \theta)\mathbf{L}'(h_0)\| : h_0 \in \mathcal{H} \} \\ &\leq \frac{\|g\|}{\sigma^2} \sup \{ \|\mathbf{L}(h_0) - \theta\| \|\mathbf{L}'(h_0)\| : h_0 \in \mathcal{H} \}. \end{aligned} \quad (43)$$

The key to the first of the preceding inequalities is a generalized mean value theorem.²¹ Both the first and second inequalities also use the fact that $h + a_k g \in \mathcal{H}$ eventually. The final inequality uses the fact that $\rho'_\sigma(x) = \rho'(x/\sigma)/\sigma \leq |x|/\sigma^2$ for any choice of $-\infty \leq \alpha \leq 2$. This inequality suggests a natural condition of

$$\mathbf{E}_\mu \left[\sup_{h_0 \in \mathcal{H}} \|\mathbf{L}(h_0)\mathbf{L}'(h_0)\| \right] < \infty \quad (44)$$

under which we can apply a standard dominated convergence argument.²² In particular, the key implication is that

$$(44) \implies \lim_{k \rightarrow \infty} \mathbf{E}_\mu f_k = \mathbf{E}_\mu f. \quad (45)$$

Since we have

$$\lim_{k \rightarrow \infty} \mathbf{E}_\mu f_k = \lim_{a \rightarrow 0_+} \mathbf{E}_\mu \left[\frac{\rho_\sigma(\mathbf{L}(h + ag) - \theta) - \rho_\sigma(\mathbf{L}(h) - \theta)}{a} \right] = D'_\rho(h; \theta)(g),$$

where $D'_\rho(h; \theta) : \mathcal{U} \rightarrow \mathbb{R}$ denotes the gradient of $h \mapsto D_\rho(h; \theta)$, we see that by applying the preceding argument (culminating in (45)) to the modified losses (16), we readily obtain the desired result. \square

Proof of Theorem 3. To begin, let us consider the smoothness of the objective $h \mapsto R_\rho(h; \theta, \eta)$ under the present assumptions. From Lemma 12 and the basic properties of the Barron class of dispersion functions (Lemma 8), it follows that this function is λ -smooth with coefficient

$$\lambda := \frac{\lambda_5}{\sigma} + \frac{\lambda_2}{\sigma^2} \mathbf{E}_\mu \|\mathbf{L}'\|_{\mathcal{H}}, \quad (46)$$

where $\lambda_5 := \lambda_3 + \lambda_4$, and

$$\lambda_3 := \left(\frac{\lambda_2}{\sigma} \right) \left[\mathbf{E}_\mu \|\mathbf{L}'\|_{\mathcal{H}}^2 + \sup_{h \in \mathcal{H}} \mathbf{E}_\mu \|\mathbf{L}'(h)\| \right], \quad \lambda_4 := \lambda_1 \|\rho'\|_\infty.$$

From here, we can leverage the main argument of [Cutkosky and Mehta \(2021, Thm. 2\)](#), utilizing the smoothness property given by (46) above, and the Γ -bound of (19). For completeness and transparency we include the key details here. First,

¹⁹See [Penot \(2012, Thm. 2.47\)](#) for this key fact, where “ X ” is \mathcal{U} here, and both “ Y ” and “ Z ” are \mathbb{R} here.

²⁰[Luenberger \(1969, §7.2, Prop. 2\)](#)

²¹Considering the proof of Lemma 14 due to [Luenberger \(1969, §7.3, Prop. 2\)](#), just generalize the one-dimensional part of the argument from the original interval $[0, 1]$ to the interval $[0, a_k]$ here.

²²See for example [Ash and Doléans-Dade \(2000, Thm. 1.6.9\)](#). If (43) holds for say all $k \geq k_0$, then we can just bound $|f_k|$ by the greater of $\max_{j \leq k_0} |f_j|$ (clearly μ -integrable) and the right-hand side of (43).

note that if we define $\epsilon_t := G_t - \partial_h R_\rho(h_t; \theta, \eta)$, $\widehat{\epsilon}_t := M_t - \partial_h R_\rho(h_t; \theta, \eta)$, and $S(h_t, h_{t+1}) := \partial_h R_\rho(h_t; \theta, \eta) - \partial_h R_\rho(h_{t+1}; \theta, \eta)$, our definitions imply that for each $t \geq 1$, we have

$$M_{t+1} = \partial_h R_\rho(h; \theta, \eta) + b(\widehat{\epsilon}_t + S(h_t, h_{t+1})) + (1-b)\epsilon_{t+1} \quad (47)$$

Using the form (47), it follows immediately that

$$\widehat{\epsilon}_{t+1} = b(\widehat{\epsilon}_t + S(h_t, h_{t+1})) + (1-b)\epsilon_{t+1} \quad (48)$$

again for each $t \geq 1$. By setting $M_0 := 0$ and $h_0 := h_1$, we trivially have

$$\widehat{\epsilon}_0 = -\partial_h R_\rho(h_0; \theta, \eta) = -\partial_h R_\rho(h_1; \theta, \eta)$$

and one can then easily check that (48) holds for all $t \geq 0$. Expanding the recursion of (48), we have

$$\widehat{\epsilon}_{t+1} = (1-b) \sum_{k=0}^t b^k \epsilon_{t-k+1} + \sum_{k=1}^{t+1} b^k S(h_{t-k+1}, h_{t-k+2}) + b^{t+1} \widehat{\epsilon}_0. \quad (49)$$

We take the summands of (49) one at a time. Since the stochastic gradients are Γ -bounded, we have that $b^k \|\epsilon_{t-k+1}\| \leq 2\Gamma$ for all $0 \leq k \leq t$. Furthermore, we have $\mathbf{E}_\mu \epsilon_{t-k+1} = 0$ (Lemma 2) and $\mathbf{E}_\mu (b^k \|\epsilon_{t-k+1}\|)^2 \leq (2b^k \Gamma)^2$ for each k . These bounds can be passed to standard concentration inequalities for martingales on Banach spaces (Cutkosky and Mehta, 2021, Lem. 14), which using the smoothness property of \mathcal{H} that we assumed tell us that with probability no less than $1 - \delta$, we have

$$\left\| \sum_{k=0}^t b^k \epsilon_{t-k+1} \right\| \leq 10\Gamma \max\{1, \log(3\delta^{-1})\} + 8\Gamma \sqrt{\max\{1, \log(3\delta^{-1})\} \sum_{k=0}^t b^{2k}}. \quad (50)$$

Moving on to the second term of (49), note that using λ -smoothness of the risk function with coefficient λ given by (46), along with the definition of the update procedure (20)–(21), we have

$$\|S(h_t, h_{t+1})\| \leq \lambda \|h_t - h_{t+1}\| = \lambda a_t \|\widetilde{M}_t\| = \lambda a_t.$$

This implies that using a constant step size $a_t = a$, we can control the sum as

$$\left\| \sum_{k=1}^{t+1} b^k S(h_{t-k+1}, h_{t-k+2}) \right\| \leq \lambda \sum_{k=1}^{t+1} a_{t-k+1} b^k \leq \frac{a\lambda}{1-b}. \quad (51)$$

Finally, the third term of (49) is easily controlled as $\|\widehat{\epsilon}_0\| = \|\partial_h R_\rho(h_1; \theta, \eta)\| \leq \Gamma$. Taking this bound along with (50) and (51), we see that $\|\widehat{\epsilon}_{t+1}\|$ can be bounded above by

$$(1-b) \left(10\Gamma \max\{1, \log(3\delta^{-1})\} + 8\Gamma \sqrt{\max\{1, \log(3\delta^{-1})\} \sum_{k=0}^t b^{2k}} \right) + \frac{a\lambda}{1-b} + b^{t+1} \Gamma \quad (52)$$

on the high-probability event mentioned earlier. To make use of the bound (52), note that using the λ -smoothness of the losses, the update procedure used here can be shown (Cutkosky and Mehta, 2021, Lem. 1) to satisfy

$$\sum_{k=1}^t \|\partial_h R_\rho(h_k; \theta, \eta)\| \leq \frac{R_\rho(h_1; \theta, \eta) - R_\rho(h_{t+1}; \theta, \eta)}{a} + 2 \sum_{k=1}^t \|\widehat{\epsilon}_k\| + \left(\frac{\lambda a}{2}\right) t. \quad (53)$$

Using a union bound, we have that the bound of (52) holds for all $\widehat{\epsilon}_1, \dots, \widehat{\epsilon}_t$ with probability no less than $1 - t\delta$. To get some clean bounds out of (52) and (53), first we loosen

$$(1-b) \sqrt{\sum_{k=0}^t b^{2k}} \leq \frac{1-b}{\sqrt{1-b^2}} \leq \frac{1-b}{\sqrt{1-b}} = \sqrt{1-b}$$

and note that $0 < \delta < 1$ and $t \geq 1 \geq e/3$ implies $\log(3t\delta^{-1}) \geq 1$. This means that with probability no less than $1 - \delta$, we have

$$2 \sum_{k=1}^t \|\widehat{\epsilon}_k\| \leq (1-b)20\Gamma t \log(3t\delta^{-1}) + 16\Gamma t \sqrt{(1-b) \log(3t\delta^{-1})} + \frac{2a\lambda t}{1-b} + \frac{2\Gamma}{1-b}.$$

Dividing both sides by t and plugging in the prescribed settings of a and b , we have

$$\frac{2}{t} \sum_{k=1}^t \|\widehat{\epsilon}_k\| \leq \frac{20\Gamma \log(3t\delta^{-1})}{\sqrt{t}} + \frac{16\Gamma \sqrt{\log(3t\delta^{-1})}}{t^{1/4}} + \frac{2\lambda}{t^{1/4}} + \frac{2\Gamma}{\sqrt{t}}. \quad (54)$$

Taking this bound and our a setting and applying it to (53), we obtain

$$\frac{1}{t} \sum_{k=1}^t \|\partial_h R_\rho(h_t; \theta, \eta)\| \leq \frac{R_\rho(h_1; \theta, \eta) - R_\rho(h_{t+1}; \theta, \eta)}{t^{1/4}} + (\text{RHS of (54)}) + \frac{\lambda}{2t^{3/4}}. \quad (55)$$

To clean this all up, we have

$$\begin{aligned} \frac{1}{t} \sum_{k=1}^t \|\partial_h R_\rho(h_t; \theta, \eta)\| &\leq \frac{1}{t^{1/4}} \left(R_\rho(h_1; \theta, \eta) - R_\rho(h_{t+1}; \theta, \eta) + 16\Gamma \sqrt{\log(3t\delta^{-1})} + 2\lambda \right) \\ &\quad + \frac{1}{\sqrt{t}} \left(20\Gamma \log(3t\delta^{-1}) + 2\Gamma \right) + \frac{\lambda}{2t^{3/4}}. \end{aligned}$$

For readability, the proof uses a slightly looser choice of λ_3 , and instead of t iterations, it is stated for T iterations. \square

Proof of Corollary 4. Let $\alpha = 0$, and note from §F.3 that we have

$$\rho_\sigma(x) = \frac{2x}{x^2 + 2\sigma^2}.$$

It thus follows from (17) that

$$\partial_h L_\rho(h; \theta, \eta) = \frac{2(L(h) - \theta)}{(L(h) - \theta)^2 + 2\sigma^2} L'(h).$$

In the case of the quadratic loss with a linear model as assumed here, this becomes

$$\begin{aligned} \partial_h L_\rho(h; \theta, \eta) &= \frac{2(L(h) - \theta)}{(L(h) - \theta)^2 + 2\sigma^2} (h(\mathbf{X}) - \mathbf{Y}) \mathbf{X} \\ &= \frac{2(L(h) - \theta) \sqrt{2L(h)}}{(L(h) - \theta)^2 + 2\sigma^2} \text{sign}(h(\mathbf{X}) - \mathbf{Y}) \mathbf{X}. \end{aligned}$$

Regarding growth in \mathbf{X} , note that since $L(h) = \mathcal{O}(\|\mathbf{X}\|^2)$, both the numerator and denominator are $\mathcal{O}(\|\mathbf{X}\|^4)$. As for growth in $L(h)$ which accounts for the random noise ε as well, the numerator is $\mathcal{O}(|L|^{3/2})$ whereas the denominator is $\mathcal{O}(|L|^2)$, and thus accounting for the randomness of \mathbf{X} and ε which can both be potentially unbounded and heavy-tailed, we see that the norm of $\partial_h L_\rho(h; \theta, \eta)$ must be bounded as the norm of the inputs, noise, and/or loss grow large. Trivially, since $\sigma > 0$, the limits where $L(h) \rightarrow \theta$ also result in $\partial_h L_\rho(h; \theta, \eta)$ with a norm that is almost surely bounded.

Proceeding to the case of the logistic loss, an analogous argument yields the desired result. First note that under the linear model assumed here, the gradient with respect to any h_j takes the form

$$\partial_{h_j} L_\rho(h; \theta, \eta) = \frac{2(L(h) - \theta)}{(L(h) - \theta)^2 + 2\sigma^2} (p_j(h) - \tilde{Y}_j) \mathbf{X}$$

where $p_j(h) := \exp(h_j(\mathbf{X})) / \sum_{i=1}^k \exp(h_i(\mathbf{X}))$, i.e., the softmax transformation of the score assigned by h_j . By definition, the coefficients $(p_j(h) - \tilde{Y}_j)$ are bounded. Furthermore, using our linear model assumption, we have that $L(h) = \mathcal{O}(\|\mathbf{X}\|)$, and as such the numerator and denominator are both $\mathcal{O}(\|\mathbf{X}\|^2)$, implying the desired boundedness.

Taking the previous two paragraphs together, we have that the bound (19) is satisfied for a finite Γ under $\alpha = 0$, even when the data is unbounded and potentially heavy-tailed. For $\alpha < 0$, the derivative of ρ shrinks even faster, so the same result follows *a fortiori* from the $\alpha = 0$ case. Finally, the λ_1 -smoothness assumption in Theorem 3 follows from direct inspection of the forms of $L'(h)$ given here for each loss, using our assumption of $\|\mathbf{X}\|^2$ having bounded second moments. \square

Proof of Proposition 7. To begin, recall from §3 the notation $M_\rho(h) := \arg \min_{\theta \in \mathbb{R}} D_\rho(h; \theta)$ for the M-locations and $\theta_\rho(h; \eta) := \arg \min_{\theta \in \mathbb{R}} R_\rho(h; \theta, \eta)$ for the T-risk optimal thresholds. By Lemma 10, both $M_\rho(h)$ and $\theta_\rho(h; \eta)$ have unique solutions, and thus we overload this notation to represent the unique solution. By definition, we have

$$\begin{aligned} \underline{R}_\rho(h; \eta) &= R_\rho(h; \theta_\rho(h; \eta)) \leq \eta M_\rho(h) + D_\rho(h; M_\rho(h)) \\ &= \eta R(h) + \eta(M_\rho(h) - R(h)) + D_\rho(h; M_\rho(h)). \end{aligned}$$

Similarly, we can obtain a lower bound using the optimality of $M_\rho(h)$ and $\theta_\rho(h; \eta)$ as

$$\begin{aligned} \underline{R}_\rho(h; \eta) &= \eta \theta_\rho(h; \eta) + D_\rho(h; \theta_\rho(h; \eta)) \\ &\geq \eta \theta_\rho(h; \eta) + D_\rho(h; M_\rho(h)) \\ &= \eta R(h) + \eta(\theta_\rho(h; \eta) - R(h)) + D_\rho(h; M_\rho(h)). \end{aligned}$$

Taking these two bounds together, we have

$$\eta(\theta_\rho(h; \eta) - R(h)) + D_\rho(h; M_\rho(h)) \leq \underline{R}_\rho(h; \eta) - \eta R(h) \leq \eta(M_\rho(h) - R(h)) + D_\rho(h; M_\rho(h)) \quad (56)$$

for any choice of $\eta \in \mathbb{R}$. The bounds in (56) are stated for ideal risk quantities for the true distribution under μ , but an identical argument holds if we replace μ by the empirical measure induced by an iid sample L_1, \dots, L_n . Writing this out explicitly, let \widehat{D}_ρ , \widehat{R}_ρ , and \widehat{R} denote the empirical analogues of D_ρ , \underline{R}_ρ , and R , and similarly let \widehat{M}_ρ and $\widehat{\theta}_\rho$ be the empirical analogues of M_ρ and θ_ρ . From the argument leading to (56), it follows that

$$\eta(\widehat{\theta}_\rho(h; \eta) - \widehat{R}(h)) + \widehat{D}_\rho(h; \widehat{M}_\rho(h)) \leq \widehat{R}_\rho(h; \eta) - \eta \widehat{R}(h) \leq \eta(\widehat{M}_\rho(h) - \widehat{R}(h)) + \widehat{D}_\rho(h; \widehat{M}_\rho(h)). \quad (57)$$

Next, using the lower bound in (57), for any $\eta \geq 0$ we have that

$$\begin{aligned} \eta R(h) &= \eta \left(\widehat{R}(h) + (R(h) - \widehat{R}(h)) \right) \\ &\leq \eta \left(\widehat{R}(h) + \|R - \widehat{R}\|_{\mathcal{H}} \right) \\ &\leq \widehat{R}_\rho(h; \eta) - \eta(\widehat{\theta}_\rho(h; \eta) - \widehat{R}(h)) - \widehat{D}_\rho(h; \widehat{M}_\rho(h)) + \eta \|R - \widehat{R}\|_{\mathcal{H}} \\ &\leq \widehat{R}_\rho(h; \eta) + \eta(\widehat{R}(h) - \widehat{\theta}_\rho(h; \eta)) + \eta \|R - \widehat{R}\|_{\mathcal{H}}. \end{aligned} \quad (58)$$

Letting \widehat{h} be a minimizer of \widehat{R}_ρ , using the upper bound in (57) and any choice of h^* , we have

$$\begin{aligned} \widehat{R}_\rho(\widehat{h}; \eta) &\leq \widehat{R}_\rho(h^*; \eta) \leq \eta \widehat{R}(h^*) + \eta(\widehat{M}_\rho(h^*) - \widehat{R}(h^*)) + \widehat{D}_\rho(h^*; \widehat{M}_\rho(h^*)) \\ &\leq \eta R(h^*) + \eta(\widehat{M}_\rho(h^*) - \widehat{R}(h^*)) + \widehat{D}_\rho(h^*; \widehat{M}_\rho(h^*)) + \eta \|R - \widehat{R}\|_{\mathcal{H}}. \end{aligned} \quad (59)$$

Combining (58) and (59), we have that $\eta R(\widehat{h})$ is bounded above by

$$\eta R(h^*) + \eta(\widehat{M}_\rho(h^*) - \theta_\rho(\widehat{h}; \eta)) + \eta(\widehat{R}(\widehat{h}) - \widehat{R}(h^*)) + \widehat{D}_\rho(h^*; \widehat{M}_\rho(h^*)) + 2\eta \|R - \widehat{R}\|_{\mathcal{H}}. \quad (60)$$

Some elementary manipulations let us bound the key differences in (60) as

$$\widehat{M}_\rho(h^*) - \theta_\rho(\widehat{h}; \eta) + \widehat{R}(\widehat{h}) - \widehat{R}(h^*) \leq 2\|\widehat{M}_\rho - R\|_{\mathcal{H}} + 2\|R - \widehat{R}\|_{\mathcal{H}} + \|\widehat{M}_\rho - \theta_\rho\|_{\mathcal{H}},$$

and thus dividing by $\eta > 0$, we end up with a final bound taking the form

$$R(\widehat{h}) \leq R(h^*) + \|\widehat{M}_\rho - \theta_\rho\|_{\mathcal{H}} + 2\|\widehat{M}_\rho - R\|_{\mathcal{H}} + \frac{1}{\eta} \widehat{D}_\rho(h^*; \widehat{M}_\rho(h^*)) + 4\|R - \widehat{R}\|_{\mathcal{H}}.$$

The desired result is just the special case where we set h^* is the expected loss minimizer. \square

E.2 Smoothness computations (proof of Lemma 12)

Here we provide detailed computations for the smoothness coefficients used in Lemma 12. We assume here that the assumptions A1, A2, and A3 are satisfied. Starting with the difference of expected gradients, using Jensen's inequality and the smoothness assumption A2, we have

$$\begin{aligned} \|\mathbf{E}_\mu [\mathbf{L}'(h_1) - \mathbf{L}'(h_2)]\| &\leq \mathbf{E}_\mu \|\mathbf{L}'(h_1) - \mathbf{L}'(h_2)\| \\ &\leq \lambda_1 \|h_1 - h_2\|. \end{aligned} \quad (61)$$

As discussed in §D.4, differences of gradients modulated by ρ' are slightly more complicated. In particular, recalling the equality (29), the norm of the difference

$$\mathbf{E}_\mu \rho'_\sigma(\mathbf{L}(h_1) - \theta_1)\mathbf{L}'(h_1) - \mathbf{E}_\mu \rho'_\sigma(\mathbf{L}(h_2) - \theta_2)\mathbf{L}'(h_2) \quad (62)$$

can be bounded above by the sum of

$$\mathbf{E}_\mu \|\mathbf{L}'(h_1)\| |\rho'_\sigma(\mathbf{L}(h_1) - \theta_1) - \rho'_\sigma(\mathbf{L}(h_2) - \theta_2)| \quad (63)$$

and

$$\mathbf{E}_\mu |\rho'_\sigma(\mathbf{L}(h_2) - \theta_2)| \|\mathbf{L}'(h_1) - \mathbf{L}'(h_2)\|. \quad (64)$$

We take up (63) and (64) one at a time. Starting with (63), from A3 we know that the dispersion derivative ρ' is λ_2 -Lipschitz, and thus we have

$$\begin{aligned} (63) &\leq \left(\frac{\lambda_2}{\sigma}\right) \mathbf{E}_\mu \|\mathbf{L}'(h_1)\| (|\mathbf{L}(h_1) - \mathbf{L}(h_2)| + |\theta_1 - \theta_2|) \\ &\leq \left(\frac{\lambda_2}{\sigma}\right) \mathbf{E}_\mu \|\mathbf{L}'(h_1)\| \left(\|h_1 - h_2\| \sup_{0 < c < 1} \|\mathbf{L}'((1-c)h_1 - ch_2)\| + |\theta_1 - \theta_2| \right) \\ &\leq \left(\frac{\lambda_2}{\sigma}\right) \left(\|h_1 - h_2\| \mathbf{E}_\mu \|\mathbf{L}'\|_{\mathcal{H}}^2 + |\theta_1 - \theta_2| \sup_{h \in \mathcal{H}} \mathbf{E}_\mu \|\mathbf{L}'(h)\| \right) \\ &\leq \lambda_3 (\|h_1 - h_2\| + |\theta_1 - \theta_2|). \end{aligned} \quad (65)$$

Here, the second inequality uses the helper Lemma 14 and our assumption of differentiability, while the third inequality uses our assumption on the expected squared norm of the gradient. We have set the Lipschitz coefficient λ_3 in (65) to be

$$\lambda_3 := \left(\frac{\lambda_2}{\sigma}\right) \left[\mathbf{E}_\mu \|\mathbf{L}'\|_{\mathcal{H}}^2 + \sup_{h \in \mathcal{H}} \mathbf{E}_\mu \|\mathbf{L}'(h)\| \right].$$

This gives us a bound on (63). Moving on to (64), if ρ' is bounded on \mathbb{R} , then we have

$$\begin{aligned} (64) &\leq \|\rho'\|_\infty \mathbf{E}_\mu \|\mathbf{L}'(h_1) - \mathbf{L}'(h_2)\| \\ &\leq \lambda_4 \|h_1 - h_2\| \end{aligned} \quad (66)$$

with $\lambda_4 := \lambda_1 \|\rho'\|_\infty$, recalling the bound (61). To summarize, we can use (65) and (66) to control (62) as follows:

$$\begin{aligned} (62) &\leq (63) + (64) \\ &\leq \lambda_3 (\|h_1 - h_2\| + |\theta_1 - \theta_2|) + \lambda_4 \|h_1 - h_2\| \\ &\leq \lambda_5 (\|h_1 - h_2\| + |\theta_1 - \theta_2|) \end{aligned}$$

where $\lambda_5 := \lambda_3 + \lambda_4$. With these preparatory details organized, it is straightforward to obtain a Lipschitz property on the gradient of $R_\rho(\cdot)$, as summarized in Lemma 12, and detailed in the proof below.

Proof of Lemma 12. Using our upper bounds on (61) and (62), we have

$$\begin{aligned} &\|\partial_h R_\rho(h_1; \theta_1, \eta) - \partial_h R_\rho(h_2; \theta_2, \eta)\| \\ &\leq \left(\frac{1}{\sigma}\right) \left\| \mathbf{E}_\mu \left[\rho' \left(\frac{\mathbf{L}(h_1) - \theta_1}{\sigma} \right) \mathbf{L}'(h_1) - \rho' \left(\frac{\mathbf{L}(h_2) - \theta_2}{\sigma} \right) \mathbf{L}'(h_2) \right] \right\| \\ &\leq \left(\frac{\lambda_5}{\sigma}\right) (\|h_1 - h_2\| + |\theta_1 - \theta_2|). \end{aligned}$$

Next, let us look at the partial derivative taken with respect to the threshold parameter θ . To bound the absolute value of these differences, using the generalized mean value theorem (Lemma 14), we have

$$\begin{aligned} |\partial_\theta R_\rho(h_1; \theta_1, \eta) - \partial_\theta R_\rho(h_2; \theta_2, \eta)| &\leq \left(\frac{1}{\sigma} \right) \left| \mathbf{E}_\mu \left[\rho' \left(\frac{L(h_1) - \theta_1}{\sigma} \right) - \rho' \left(\frac{L(h_2) - \theta_2}{\sigma} \right) \right] \right| \\ &\leq \left(\frac{\lambda_2}{\sigma^2} \right) (\mathbf{E}_\mu |L(h_1) - L(h_2)| + |\theta_1 - \theta_2|) \\ &\leq \left(\frac{\lambda_2}{\sigma^2} \mathbf{E}_\mu \|L'\|_{\mathcal{H}} \right) (\|h_1 - h_2\| + |\theta_1 - \theta_2|). \end{aligned}$$

Taking the preceding upper bounds together, the gradient difference for R_ρ can be bounded as

$$\begin{aligned} &\|R'_\rho(h_1; \theta_1, \eta) - R'_\rho(h_2; \theta_2, \eta)\| \\ &= \|\partial_h R_\rho(h_1; \theta_1, \eta) - \partial_h R_\rho(h_2; \theta_2, \eta)\| + |\partial_\theta R_\rho(h_1; \theta_1, \eta) - \partial_\theta R_\rho(h_2; \theta_2, \eta)| \\ &\leq \left(\left(\frac{\lambda_5}{\sigma} \right) + \frac{\eta \lambda_2}{\sigma^2} \mathbf{E}_\mu \|L'\|_{\mathcal{H}} \right) (\|h_1 - h_2\| + |\theta_1 - \theta_2|), \end{aligned}$$

noting that the initial equality follows from the fact that we are using the sum of norms for our product space norm here (see also Remark 13). These bounds on the gradient differences are precisely the desired result. \square

F Additional technical facts

F.1 Lipschitz properties

Here we give a fundamental property of differentiable functions that generalizes the mean value theorem.

Lemma 14. *Let \mathcal{U} and \mathcal{V} be normed linear spaces, and let $f : \mathcal{U} \rightarrow \mathcal{V}$ be Fréchet differentiable on an open set $S \subseteq \mathcal{U}$. Taking any $u \in S$, we have*

$$\|f(u + u') - f(u)\| \leq \|u'\| \sup_{0 < c < 1} \|f'(u + cu')\|$$

for any $u' \in \mathcal{U}$ such that $u + cu' \in S$ for all $0 \leq c \leq 1$.

Proof. See Luenberger (1969, §7.3, Prop. 2). \square

Note that Lemma 14 has the following important corollary: *bounded gradients imply Lipschitz continuity*. In particular, if $\|f'(u)\| \leq \lambda < \infty$ for all $u \in S$, then it follows immediately that f is λ -Lipschitz on S .

A closely related result goes in the other direction. Let $f : \mathcal{U} \rightarrow \overline{\mathbb{R}}$ be convex and λ -Lipschitz. If f is sub-differentiable at a point $u \in \mathcal{U}$, then we have

$$|\langle \partial f(u), u' - u \rangle| \leq |f(u') - f(u)| \leq \lambda \|u' - u\|.$$

As such, for convex, sub-differentiable functions, λ -Lipschitz continuity implies that all sub-gradients are bounded as $\|\partial f(x)\| \leq \lambda$.

F.2 Convexity

Lemma 15. *Let function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be twice continuously differentiable. Then f is convex if and only if its Hessian is positive semi-definite, namely when*

$$\langle f''(v)u, u \rangle \geq 0$$

for all $u, v \in \mathbb{R}^d$.

Proof. See Nesterov (2004, Thm. 2.1.4). \square

F.3 Derivatives for the Barron class

Let $\rho(\cdot; \alpha)$ be defined according to (15). Here we compute derivatives of the map $x \mapsto \rho_\sigma(x; \alpha)$, using the shorthand notation $\rho_\sigma(x; \alpha) := \rho(x/\sigma; \alpha)$. We denote the first derivative of $\rho_\sigma(\cdot; \alpha)$ evaluated at $x \in \mathbb{R}$ by $\rho'_\sigma(x; \alpha)$, which is computed as

$$\rho'_\sigma(x; \alpha) = \begin{cases} x/\sigma^2, & \text{if } \alpha = 2 \\ 2x/(x^2 + 2\sigma^2), & \text{if } \alpha = 0 \\ (x/\sigma^2) \exp(-(x/\sigma)^2/2), & \text{if } \alpha = -\infty \\ \frac{x}{\sigma^2} \left(1 + \frac{(x/\sigma)^2}{|\alpha-2|}\right)^{(\alpha/2)-1}, & \text{otherwise.} \end{cases} \quad (67)$$

In the same way, letting $\rho''_\sigma(x; \alpha)$ denote the second derivative of $\rho_\sigma(\cdot; \alpha)$ evaluated at $x \in \mathbb{R}$, this is computed as

$$\rho''_\sigma(x; \alpha) = \begin{cases} 1/\sigma^2, & \text{if } \alpha = 2 \\ \frac{2}{x^2+2\sigma^2} \left(1 - \frac{2x^2}{x^2+2\sigma^2}\right), & \text{if } \alpha = 0 \\ (1/\sigma^2) \exp(-(x/\sigma)^2/2) \left(1 - \left(\frac{x}{\sigma}\right)^2\right), & \text{if } \alpha = -\infty \\ \frac{1}{\sigma^2} \left(1 + \frac{(x/\sigma)^2}{|\alpha-2|}\right)^{(\alpha/2)-1} \left(1 - \frac{1-(\alpha/2)}{|\alpha-2|} \frac{2(x/\sigma)^2}{1+(x/\sigma)^2/|\alpha-2|}\right), & \text{otherwise.} \end{cases} \quad (68)$$

We emphasize that $\rho'_\sigma(x; \alpha)$ and $\rho''_\sigma(x; \alpha)$ are not equal to $\rho'(x/\sigma; \alpha)$ and $\rho''(x/\sigma; \alpha)$, but by a simple application of the chain rule are easily seen to satisfy the relations

$$\rho'_\sigma(x; \alpha) = \frac{1}{\sigma} \rho'(x/\sigma; \alpha), \quad \rho''_\sigma(x; \alpha) = \frac{1}{\sigma^2} \rho''(x/\sigma; \alpha)$$

for any $x \in \mathbb{R}$, $\sigma > 0$, and $\alpha \in [-\infty, 2]$.

F.4 Elementary inequalities

The following elementary inequalities will be of use.

$$\left(1 + \frac{x}{p}\right)^p \geq \left(1 + \frac{x}{q}\right)^q, \quad \forall x \geq 0, \quad p > q > 0 \quad (69)$$

$$(1+x)^c \leq 1 + \frac{cx}{1-(c-1)x}, \quad -1 \leq x < \frac{1}{c-1}, \quad c > 1 \quad (70)$$

The inequality below is sometimes referred to as Bernoulli's inequality.

$$(1+x)^a \geq 1+ax, \quad \forall x > -1, \quad a \geq 1. \quad (71)$$

F.5 Expected dispersion is coercive

Lemma 16 (Expected dispersion is coercive). *Let $f : \mathbb{R} \rightarrow \mathbb{R}_+$ be any non-negative function which is even (i.e., $f(x) = f(-x)$ for all $x \in \mathbb{R}$) and non-decreasing on \mathbb{R}_+ . Let X be any random variable such that $\mathbf{E}_\mu f(X - \theta) < \infty$ for all $\theta \in \mathbb{R}$. Then, we have*

$$\lim_{|\theta| \rightarrow \infty} \mathbf{E}_\mu f(X - \theta) = \lim_{x \rightarrow \infty} f(x)$$

and note that this includes the divergent case where $f(x) \rightarrow \infty$ as $|x| \rightarrow \infty$.

Proof of Lemma 16. By our assumptions, we have $f(x) \geq 0$ and $f(-x) = f(x)$ for all $u \in \mathbb{R}$, and $f(x_1) \leq f(x_2)$ whenever $0 \leq x_1 \leq x_2$. With these facts in place, note that for any choice of $a \geq 0$ and θ such that $|\theta| \geq a$, we have

$$\begin{aligned} \mathbf{E}_\mu f(X - \theta) &= \mathbf{E}_\mu f(|\theta - X|) \\ &\geq \mathbf{E}_\mu f(|\theta| - |X|) \\ &\geq \mathbf{E}_\mu \mathbf{I}_{\{|X| \leq a\}} f(|\theta| - |X|) \\ &\geq f(|\theta| - a) \mathbf{P}\{|X| \leq a\}. \end{aligned} \quad (72)$$

For readability, let us write $f(+\infty)$ for the limit of $f(x)$ as $|x| \rightarrow \infty$. Trivially, we know that $\mathbf{E}_\mu f(X - \theta) \leq f(+\infty)$. Using the preceding inequality (72), we have a lower bound of $\mathbf{E}_\mu f(X - \theta) \geq f(+\infty) \mathbf{P}\{|X| \leq a\}$ that holds for any $a \geq 0$. When $f(+\infty) = \infty$, the desired result is immediate. When $f(+\infty) < \infty$, simply note that $\{|X| \leq a\} \uparrow \Omega$ as $a \uparrow \infty$, and thus using the continuity of probability measures, we have $\mathbf{P}\{|X| \leq a\} \rightarrow 1$ as $a \rightarrow \infty$.²³ Thus, the lower bound (72) can be taken arbitrarily close to $f(+\infty)$, implying the desired result. \square

²³All countably additive set functions on σ -fields satisfy such continuity properties (Ash and Doléans-Dade, 2000, Thm. 1.2.7).