
Supervised Learning with General Risk Functionals

Liu Leqi¹ Audrey Huang² Zachary C. Lipton¹ Kamyar Azizzadenesheli³

Abstract

Standard uniform convergence results bound the generalization gap of the *expected* loss over a hypothesis class. The emergence of risk-sensitive learning requires generalization guarantees for functionals of the loss distribution beyond the expectation. While prior works specialize in uniform convergence of particular functionals, our work provides uniform convergence for a general class of Hölder risk functionals for which the closeness in the Cumulative Distribution Function (CDF) entails closeness in risk. We establish the first uniform convergence results for estimating the CDF of the loss distribution, which yield uniform convergence guarantees that hold simultaneously both over a class of Hölder risk functionals and over a hypothesis class. Thus licensed to perform empirical risk minimization, we develop practical gradient-based methods for minimizing distortion risks (widely studied subset of Hölder risks that subsumes the spectral risks, including the mean, conditional value at risk, cumulative prospect theory risks, and others) and provide convergence guarantees. In experiments, we demonstrate the efficacy of our learning procedure, both in settings where uniform convergence results hold and in high-dimensional settings with deep networks.

1. Introduction

To date, the vast majority of supervised, unsupervised, and reinforcement learning research has focused on objectives expressible as expectations (over some dataset or distribution) of an underlying loss (or reward) function. This focus is understandable. The expected loss is mathematically convenient and a reasonable default, and a special case of nearly every proposed family of risks. To be sure, this focus has

paid off: we now possess a rich body of theory and methods for evaluating, optimizing, and providing theoretical guarantees on the expected loss.

However, real-world concerns such as risk aversion, equitable allocations of benefits and harms, or alignment with human preferences, often demand that we address other functionals of the loss distribution. For example, in finance, the expectation of returns must be weighed against their variance to determine an ideal portfolio allocation, as codified, e.g., in the mean-variance objective (Björk et al., 2014). Focusing on supervised learning, consider the common scenario in which a population contains a minority (constituting fraction α of the population) but where group membership was not recorded in the available data. If the pattern relating the features to the label were different for different demographics, a naively trained model might adversely harm members of a minority group. Absent further information, one sensible strategy could be to optimize the worst case performance over all subsets (of size up to α). This would translate to the familiar Conditional Value at Risk (CVaR) objective (Rockafellar et al., 2000). In addition, even in settings where a model is evaluated in terms of the expected loss at test time, the training objective may be chosen as other functionals due to reasons including distribution shifts (Duchi & Namkoong, 2018), noisy labels (Lee et al., 2020), imbalanced dataset (Li et al., 2020), etc.

Risk-sensitive learning research addresses the problem of learning models under many families of (risk) functionals, including (among others) distortion risks (Wirch & Hardy, 2001), coherent risks (Artzner et al., 1999), spectral risks (Acerbi, 2002), and cumulative prospect theory risks (Prashanth et al., 2016). Subsuming these risks under a common framework addressing bounded losses/rewards, Huang et al. (2021) recently introduced Lipschitz risk functionals, for which differences in the risk are bounded by (sup norm) differences in the Cumulative Distribution Function (CDF) of losses. Thus, because a single CDF estimate can be used to estimate all Lipschitz risks, sup norm concentration of the CDF estimate entails corresponding (simultaneous) concentration of all Lipschitz risks calculated on that CDF estimate. However, this concentration result applies only to a single hypothesis. In contrast, most uniform convergence results in learning theory have concentrated largely on the expected loss (Vapnik, 1999; 2013; Bartlett & Mendelson, 2002).

¹Machine Learning Department, Carnegie Mellon University

²Department of Computer Science, University of Illinois Urbana-Champaign ³Department of Computer Science, Purdue University. Correspondence to: Liu Leqi <leqil@cs.cmu.edu>, Audrey Huang <audrehy5@illinois.edu>.

While uniform convergence results are known for several specific risk classes, including the spectral/rank-weighted risks (Khim et al., 2020) and optimized certainty equivalent risks (Lee et al., 2020), no results to date provide uniform convergence guarantees that hold simultaneously over both a hypothesis class and a broad class of risks.

Tackling this problem, we present, to our knowledge, the first uniform convergence guarantee on estimation of the loss CDF. Our bounds rely on appropriate complexity measures of the hypothesis class. In addition to the traditional Rademacher complexity and VC dimension, we propose a new notion of permutation complexity that is especially suited to CDF estimation.

For general risk estimation, we adopt the broader class (subsuming the Lipschitz risks) of Hölder risk functionals, for which closeness in distribution entails closeness in risk. Combined with our uniform convergence guarantees for CDF estimation, this property allows us to establish uniform convergence guarantees for risk estimation of supervised learning models, which hold simultaneously over all hypotheses in the model class and over all Hölder risks.

These results license us to optimize general risks, assuring that for appropriate model classes and given sufficient data, the empirical risk minimizer will indeed generalize and that whichever objective is optimized, all Hölder risk estimates will be close to their true values. Generalization aside, optimizing complex risks is non-trivial. To tackle this problem, we propose a new algorithm for optimizing distortion risks, a subset of Hölder risks that subsumes the spectral risks, including the expectation, CVaR, cumulative prospect theory risks, and others. Our approach extends traditional gradient-based empirical risk minimization methods to handle distortion risks (Denneberg, 1990; Wang, 1996). In particular, we calculate the empirical distortion risk by re-weighting losses based on CDF values and establish convergence guarantees for the proposed optimization method. Finally, we experimentally validate our algorithm, both in settings where uniform convergence results hold and in high-dimensional settings with deep networks.

In summary, we contribute the following:

1. The first uniform convergence result for CDF estimation together with corresponding new complexity measures suited to the task (Section 4).
2. Uniform convergence for risk estimation that holds simultaneously for all Hölder risks, yielding learning guarantees for empirical risk minimization for the broad class of distortion risks (Section 5).
3. A gradient-based method for minimizing distortion risks that re-weights examples dynamically based on the empirical CDF of losses, and corresponding con-

vergence guarantees (Section 6).

4. Experiments confirming the practical usefulness of our learning algorithm (Section 7).

2. Related Literature

Risk functionals have long been studied in diverse contexts (Sharpe, 1966; Artzner et al., 1999; Rockafellar et al., 2000; Krokmal, 2007; Shapiro et al., 2014; Acerbi, 2002; Prashanth et al., 2016; Jie et al., 2018). CVaR, value-at-risk, and mean-variance (Cassel et al.; Sani et al., 2013; Vakili & Zhao, 2015; Zimin et al., 2014) rank among the most widely studied risks. Prashanth et al. (2016) introduces the cumulative prospect theory risks, which have been studied in bandit (Gopalan et al., 2017) and supervised learning settings (Leqi et al., 2019). Many previous works have tackled the evaluation (Huang et al., 2021; Chandak et al., 2021) and optimization (Torossian et al., 2019; Munos, 2014) of risk functionals.

Recent work on risk-sensitive supervised learning has established the uniform convergence of a single risk functional when losses incurred by the models are bounded (Khim et al., 2020; Lee et al., 2020), or the excess risk of a particular learning procedure in cases where the loss could be unbounded (Holland & Haress, 2021). Collectively, these works have addressed the class of spectral risks (L-risks or rank-weighted risks) that includes the expected value, CVaR and cumulative prospect theory risks (Khim et al., 2020; Holland & Haress, 2021), as well as the class of optimized certainty equivalent risks that includes the expected value, CVaR and entropic risks (Ben-Tal & Teboulle, 1986; Lee et al., 2020).

To our knowledge, the aforementioned risk-sensitive learning results are considerably narrower: the analyses apply only to smaller *families* of risks and the guarantees hold only for a *single risk functional* (not simultaneously over the family). By contrast, we establish uniform convergence results that hold simultaneously over both a broader class of risks and over an entire model class (constrained by an appropriate complexity measure). The key to our approach is to estimate the CDF of losses and control its sup norm error uniformly over a hypothesis class. CDF estimation is a central topic in learning theory (Devroye et al., 2013). Strong approximation results provide concentration bounds on the Kolmogorov–Smirnov distance (sup norm) between the true and estimated CDF (Massart, 1990). As our uniform convergence results are over a hypothesis class of possibly infinite number of hypotheses, we control the complexity of the hypothesis class using data-dependent complexity notions (e.g., Rademacher complexity) and data-independent complexity notions (e.g., VC dimension) (Alexander, 1984; Vapnik, 2006; Gänsler & Stute, 1979).

3. Preliminaries

We use \mathcal{X} to denote the space of covariates, \mathcal{Y} the space of labels, and $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. Let $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ denote a loss function and \mathbb{F} a hypothesis class, where for any $f \in \mathbb{F}$, $f : \mathcal{X} \rightarrow \mathcal{Y}$. The set \mathbb{F}_ℓ , with elements ℓ_f , denotes the class of functions that are compositions of the loss function ℓ and a hypothesis $f \in \mathbb{F}$, i.e., $\forall z \in \mathcal{Z}, \ell_f(z) = \ell(f(x), y)$. Furthermore, we use $\ell_f(Z)$ to denote the random variable of the loss incurred by $f \in \mathbb{F}$ under data $Z = (X, Y)$. For any $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$.

We use \mathcal{U} to denote the space of real-valued random variables that admit CDFs. For any $U \in \mathcal{U}$, its CDF is denoted by F_U . A risk functional $\rho : \mathcal{U} \rightarrow \mathbb{R}$ is a mapping from a space of real-valued random variables to reals. A risk functional is called law-invariant (or version-independent) if for any pair of random variables $U, U' \in \mathcal{U}$ with the same law ($F_U = F_{U'}$), we have $\rho(U) = \rho(U')$ (Kusuoka, 2001). We work with law-invariant risk functionals in this paper, and with some abuse of notation, we refer to $\rho(F_U)$ and $\rho(U)$ interchangeably.

4. Uniform Convergence for CDF

Estimation

We begin with an important building block for risk estimation—CDF estimation with uniform convergence guarantees. Given a loss function ℓ and a data set of n labeled data points $\{Z_i\}_{i=1}^n$ where $Z_i = (X_i, Y_i)$, we are interested in estimating the CDF of $\ell_f(Z)$ for all $f \in \mathbb{F}$. We use the unbiased empirical CDF estimator:

$$\widehat{F}(r; f) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell_f(Z_i) \leq r\}}, \quad (1)$$

where $\mathbb{E}[\widehat{F}(r; f)] = \mathbb{P}(\ell_f(Z) \leq r) = F(r; f)$. To establish the uniform convergence of the estimator, our central goal is to analyze the following quantity:

$$e_n(\mathbb{F}, \ell) = \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}(r; f) - F(r; f) \right|. \quad (2)$$

In Section 4.1, we exploit the special structure of CDF estimation and propose a new notion of permutation complexity that captures the complexity of the hypothesis class used for CDF estimation. In Section 4.2, we apply the more classical approach for analyzing uniform convergence that does not exploit any special structure of CDF estimation. Each approach offers a unique perspective and contributes to our understanding of CDF estimation. We highlight that the uniform convergence we provide hold for *any* loss distribution regardless of whether the loss is binary or bounded.

We first introduce notation key to our analysis. The Rademacher complexity in our setting (for a given loss

function ℓ) is given as follows:

$$\begin{aligned} \mathcal{R}(n, \mathbb{F}) &= \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \frac{1}{n} \left| \sum_{i=1}^n \xi_i \mathbb{1}_{\{\ell_f(Z_i) \leq r\}} \right| \right] \\ &= \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \frac{1}{n} \left| \sum_{i=1}^n \xi_i g(\ell_f(Z_i)) \right| \right], \quad (3) \end{aligned}$$

with \mathfrak{R} being a Rademacher measure on a set of Rademacher random variables $\{\xi_i\}_{i=1}^n$ and $\mathbb{G}(1) := \{\mathbb{1}_{\{\cdot \leq r\}} : \forall r \in \mathbb{R}\}$ is the set of indicator functions parameterized by a real-valued r . Using McDiarmid’s inequality and symmetrization, we obtain the following classical result that bounds $e_n(\mathbb{F}, \ell)$ in terms of the Rademacher complexity. All proofs in this section can be found in Appendix B.

Theorem 4.1. *Given a hypothesis class \mathbb{F} , any loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, and n samples $\{Z_i\}_{i=1}^n$, we have that with probability at least $1 - \delta$,*

$$e_n(\mathbb{F}, \ell) \leq 2\mathcal{R}(n, \mathbb{F}) + \sqrt{\frac{\log(\frac{1}{\delta})}{2n}}.$$

In general, the Rademacher complexity $\mathcal{R}(n, \mathbb{F})$ is hard to obtain. Researchers have come up with different ways to control it for various hypothesis classes, e.g., hypothesis classes with finite VC dimension (Wainwright, 2019). In the following, we discuss how we work with $\mathcal{R}(n, \mathbb{F})$.

4.1. Permutation Complexity

We first notice that $\mathcal{R}(n, \mathbb{F})$ depends *jointly* on both the hypothesis class \mathbb{F} and $\mathbb{G}(1)$. A direct approach that follows from the classical statistical learning theory is to work with the function class that combines \mathbb{F} and $\mathbb{G}(1)$, which we provide more details in Section 4.2. In this section, we propose a new way of thinking about $\mathcal{R}(n, \mathbb{F})$. By exploiting the special structure of CDF estimation (the structure of $\mathbb{G}(1)$), we uncover that $\mathcal{R}(n, \mathbb{F})$ can be controlled by *only* the complexity of the hypothesis class \mathbb{F} (or \mathbb{F}_ℓ with elements ℓ_f). In order to do so, we first introduce the notion of *permutation complexity*. This complexity measure is data-dependent and enables us to work with $\mathcal{R}(n, \mathbb{F})$ by disentangling the complexity of \mathbb{F} (or \mathbb{F}_ℓ) from that of $\mathbb{G}(1)$.

For a measurable space \mathcal{V} , let $\zeta : \mathcal{V} \rightarrow \mathbb{R}$ denote a measurable function and $\{v_i\}_{i=1}^n$ denote a set of n points in \mathcal{V} . Satisfying the conditions of selection and maximum theorems (Guide, 2006, Chapter 17), a permutation function $\pi : [n] \rightarrow [n]$ in the space $\Pi(n)$ of all permutation of size n exists, and permutes the indices of $\{v_i\}_{i=1}^n$ such that $\zeta(v_{\pi(1)}) \leq \zeta(v_{\pi(2)}) \leq \dots \leq \zeta(v_{\pi(n)})$. We note that the permutation function can depend on the specific data points $\{v_i\}_{i=1}^n$ and the function ζ of interests. In the following definitions, we consider a function class \mathbb{J} of functions

$\zeta : \mathcal{V} \rightarrow \mathbb{R}$ and a probability measure μ on $(\mathcal{V}, \sigma(\mathcal{V}))$, where $\sigma(\mathcal{V})$ denotes the σ -algebra generated by \mathcal{V} .

Definition 4.1 (Permutation Complexity). *The instance-dependent permutation complexity of \mathbb{J} at n data points $\{v_i\}_{i=1}^n$, denoted as $\mathcal{N}_\Pi(\mathbb{J}, \{v_i\}_{i=1}^n)$, is the minimum number of permutation functions $\pi \in \Pi(n)$ needed to sort elements of $\{\zeta(v_i)\}_{i=1}^n$, $\forall \zeta \in \mathbb{J}$. The permutation complexity of \mathbb{J} at n (random) data points $\{V_i\}_{i=1}^n$ with measure μ is*

$$\mathcal{N}_\Pi(n, \mathbb{J}, \mu) := \mathbb{E}_\mu [\mathcal{N}_\Pi(\mathbb{J}, \{V_i\}_{i=1}^n)].$$

To obtain a better understanding of permutation complexity, we provide the permutation complexity of monotone real-valued functions.

Lemma 4.1. *When \mathcal{V} is the space of reals, \mathbb{J} is a set of real-valued non-decreasing functions on \mathcal{V} , the permutation complexity $\mathcal{N}_\Pi(n, \mathbb{J}, \mu) = 1$.*

Proof. For a set of real-valued points $\{v_i\}_{i=1}^n$, we can construct a permutation function π such that $v_{\pi(1)} \leq v_{\pi(2)} \leq \dots \leq v_{\pi(n)}$. Since all functions in \mathbb{J} are non-decreasing, for any $\zeta \in \mathbb{J}$, we have $\zeta(v_{\pi(1)}) \leq \zeta(v_{\pi(2)}) \leq \dots \leq \zeta(v_{\pi(n)})$. Thus, $\mathcal{N}_\Pi(\mathbb{J}, \{v_i\}_{i=1}^n) = 1$ and $\mathcal{N}_\Pi(n, \mathbb{J}, \mu) = 1$ for any measure μ . \square

Remark 4.1. *This result implies that the permutation complexity of threshold functions $\mathbb{G}(1)$ is one.*

Mapping the definition to our setting, the function class \mathbb{J} of interest is \mathbb{F}_ℓ with elements ℓ_f . The data points $V_i = Z_i = (X_i, Y_i)$ and the measure $\mu = \mathbb{P}$ (the probability measure for Z).

An immediate observation is that when $|\mathbb{F}|$ is finite, we only need at most $|\mathbb{F}|$ permutation functions to sort $\{\ell_f(z_i)\}_{i=1}^n$ (one permutation function for each $f \in \mathbb{F}$), i.e., $\mathcal{N}_\Pi(\mathbb{F}_\ell, \{z_i\}_{i=1}^n) \leq |\mathbb{F}|$. For the special case of binary classification, where $\mathcal{V} = \{0, 1\}$ and the loss function ℓ is the 0/1 loss, a coarse upper bound on the permutation complexity when \mathbb{F} has finite VC dimension $\nu(\mathbb{F})$ is $\mathcal{N}_\Pi(n, \mathbb{F}_\ell, \mathbb{P}) \leq (n+1)^{\nu(\mathbb{F})}$. This is due to Sauer's Lemma: there are at most $(n+1)^{\nu(\mathbb{F})}$ ways of labeling the data, which suggests that we need at most $(n+1)^{\nu(\mathbb{F})}$ permutation functions. However, as one may have noticed, the number of permutation functions needed may be (much) smaller than this number. For example, consider a binary classification setting where we have 3 data points and \mathbb{F} is large enough such that all 2^3 possible losses (000, 001, ...) can be incurred. In such a case, we only need 4 permutation functions, since loss sequences that are non-decreasing (or non-increasing) can share the same permutation function, e.g., for loss sequences 111, 011, 001, 000, we can use the same permutation function $\pi(i) = i, \forall i \in [3]$. We note that the permutation complexity is defined for not just

binary-valued function classes. Precisely characterizing the permutation complexity for different combinations of function classes, data distributions and loss functions is of future interest.

Theorem 4.2. *For any hypothesis class \mathbb{F} and loss function $\ell : \mathcal{Y} \times \mathcal{X} \rightarrow \mathbb{R}$, we have that*

$$\mathcal{R}(n, \mathbb{F}) \leq \sqrt{\frac{\log(4\mathcal{N}_\Pi(n, \mathbb{F}_\ell, \mathbb{P}))}{2n}}.$$

Theorem 4.2 indicates that, despite $\mathcal{R}(n, \mathbb{F})$ depending on the supremum over both the function class \mathbb{F} and $\mathbb{G}(1)$ where $\mathbb{G}(1)$ is an infinite set, $\mathcal{R}(n, \mathbb{F})$ can be controlled by just the complexity of \mathbb{F}_ℓ . When the permutation complexity $\mathcal{N}_\Pi(n, \mathbb{F}_\ell, \mathbb{P})$ is polynomial in the number of samples n , we obtain that $\mathcal{R}(n, \mathbb{F}) = O(\sqrt{\log(n)/n})$. The immediate consequence of Theorem 4.2 when the hypothesis class \mathbb{F} is finite is provided below.

Corollary 4.1. *For a finite hypothesis class \mathbb{F} ,*

$$\mathcal{R}(n, \mathbb{F}) \leq \sqrt{\frac{\log(4|\mathbb{F}|)}{2n}}.$$

Corollary 4.1 suggests that the generalization bound of CDF estimation follows the same rate as classical generalization bound of the expected loss for finite function classes, i.e., $O(\sqrt{\log(|\mathbb{F}|)/n})$.

4.2. Classical Approach

In this section, we present a more classical approach for analyzing the uniform convergence without exploiting the specific structure of CDF estimation. As we have noted before, $\mathcal{R}(n, \mathbb{F})$ depends on both \mathbb{F} and $\mathbb{G}(1)$. In this approach, we directly work with the function class that combines \mathbb{F} and $\mathbb{G}(1)$: for a given loss function ℓ , we define

$$\begin{aligned} \mathbb{H} &:= \{h : \mathcal{Z} \rightarrow \{0, 1\} : \\ &h(z; r) = \mathbb{1}_{\{\ell_f(z) \leq r\}}, f \in \mathbb{F}, r \in \mathbb{R}\}. \end{aligned}$$

We note that even when \mathbb{F} is finite, \mathbb{H} is an infinite set. The Rademacher complexity (3) can be re-written as

$$\mathcal{R}(n, \mathbb{F}) = \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\sup_{h \in \mathbb{H}} \frac{1}{n} \left| \sum_{i=1}^n \xi_i h(Z_i) \right| \right].$$

Lemma 4.2. (Wainwright, 2019, Lemma 4.14) *Let $|\mathbb{H}(\{z_i\}_{i=1}^n)|$ denote the maximum cardinality of the set $\mathbb{H}(\{z_i\}_{i=1}^n) = \{(h(z_1), \dots, h(z_n)) : h \in \mathbb{H}\}$, where $n \in \mathbb{N}$ is fixed and $\{z_i\}_{i=1}^n$ can be any data collection for $z_i \in \mathcal{Z}$. Then, we have*

$$\mathcal{R}(n, \mathbb{F}) \leq 2\sqrt{\frac{\log(|\mathbb{H}(\{z_i\}_{i=1}^n)|)}{n}}.$$

Since \mathbb{H} consists of binary functions, for any data collection $\{z_i\}_{i=1}^n$, the set $\mathbb{H}(\{z_i\}_{i=1}^n)$ is finite. When \mathbb{F} is finite, we obtain that $|\mathbb{H}(\{z_i\}_{i=1}^n)| \leq (n+1)|\mathbb{F}|$. This is true since for a given data collection $\{z_i\}_{i=1}^n$ and hypothesis $f \in \mathbb{F}$, after sorting $\{\ell_f(z_i)\}_{i=1}^n$, we have at most $(n+1)$ ways of labeling them using the indicator functions $\mathbb{1}_{\{\cdot \leq r\}}$ for $r \in \mathbb{R}$. Thus, if we directly apply Lemma 4.2, we obtain that the Rademacher complexity $\mathcal{R}(n, \mathbb{F})$ is on the order of $O(\sqrt{(\log |\mathbb{F}| + \log n)/n})$, which has an extra $\log(n)$ term in the numerator compared to our bound in Corollary 4.1. In general, when \mathbb{H} has finite VC dimension $\nu(\mathbb{H})$, we obtain that $\mathcal{R}(n, \mathbb{F}) \leq O(\sqrt{\nu(\mathbb{H}) \log(n+1)/n})$. However, this result is far from being sharp. Using more advanced techniques, e.g., chaining and Dudley’s entropy integral (Wainwright, 2019), one can remove the extra $\log(n)$ factor on the numerator, and obtain that $\mathcal{R}(n, \mathbb{F}) \leq O(\sqrt{\nu(\mathbb{H})/n})$ (for more details, see Wainwright (2019, Example 5.24)). As a consequence, when $|\mathbb{F}| < \infty$, we have $\nu(\mathbb{H}) \leq \log(|\mathbb{F}|)$ and thus obtain that $\mathcal{R}(n, \mathbb{F}) \leq O(\sqrt{\log(|\mathbb{F}|)/n})$ which is at the same rate as our bound in Corollary 4.1.

5. Uniform Convergence for Hölder Risk Estimation

Using our results in Section 4, we show uniform convergence for a broad class of risk functionals—Hölder risks. As illustrated in Section 6, uniform convergence for a single risk functional provides grounding for learning models through minimizing the empirical risk. In addition to uniform convergence for a single risk, we also provide uniform convergence results hold for a collection of risks simultaneously. The second result is important due to the following reasons: Although models are trained to optimize a single risk objective, evaluating their performance under multiple risks can give a holistic assessment of their behavior—a task we call *risk assessment*. For example, models minimizing CVaR at different α levels may have different tradeoffs with their expected loss, and monitoring the progress of both objectives throughout the training process can inform choice of the best model. In addition, when given a set of models obtained under different learning mechanisms, one may want to compare them in terms of different risks. To this end, using our results on uniform convergence for CDF estimation (Section 4), we demonstrate how a collection of models may be assessed under many risks simultaneously, with estimation errors of the same order as the CDF estimation error.

5.1. Hölder Risk Functionals

We begin by introducing a new class of risks—the Hölder risk functionals—that includes many popularly studied risks and generalizes the notion of Lipschitz risk functionals (Huang et al., 2021) and Hölder continuous functionals in Wasserstein distance (Bhat & LA, 2019).

Definition 5.1. Let d denote a quasi-metric¹ on the space of CDFs. A risk functional ρ is $L(\rho, p, d)$ Hölder on a space of real-valued random variables \mathcal{U} if there exist constants $p > 0$ and $L(\rho, p, d) > 0$ such that for all $U, U' \in \mathcal{U}$ with CDF F_U and $F_{U'}$ respectively, the following holds:

$$|\rho(F_U) - \rho(F_{U'})| \leq L(\rho, p, d)d(F_U, F_{U'})^p.$$

The class of Hölder risk functionals subsumes many other risk functional classes. In particular, Lipschitz risk functionals (Huang et al., 2021) are $L(\cdot, 1, \mathcal{L}_\infty)$ Hölder on bounded random variables. As a direct result, distortion risk functionals with Lipschitz distortion functions (Denneberg, 1990; Wang, 1996; Wang et al., 1997; Balbás et al., 2009; Wirth & Hardy, 1999; 2001), cumulative prospect theory risks (Prashanth et al., 2016; Leqi et al., 2019), variance, and linear combinations of aforementioned functionals are all Hölder on bounded random variables. In addition, the optimized certainty equivalent risks (Lee et al., 2020) and spectral risks (Khim et al., 2020; Holland & Haress, 2021) recently studied in risk-sensitive supervised learning literatures, are Lipschitz (hence Hölder) on the space of bounded random variables. We provide proofs and further details in Appendix A.

To be more specific, we present a subset of Hölder risk functionals—distortion risks with Lipschitz distortion functions—that consists of many well-studied risks including the expected value, CVaR and cumulative prospect theory risks. When the loss is non-negative, the distortion risk of $\ell_f(Z)$ is defined to be:

$$\rho(F(\cdot; f)) = \int_0^\infty g(1 - F(r; f))dr, \quad (4)$$

where the distortion function $g : [0, 1] \rightarrow [0, 1]$ is non-decreasing with $g(0) = 0$ and $g(1) = 1$. In the case of expected value, the distortion function g is 1-Lipschitz. For CVaR at level α , the distortion function g is $\frac{1}{\alpha}$ -Lipschitz. For more details, we refer the readers to Huang et al. (2021). In the following, we show uniform convergence for estimating Hölder risks using our proposed estimator (Section 5.2) and develop optimization procedures to minimize distortion risks (Section 6).

5.2. Uniform Convergence for Risk Estimation

For a given hypothesis $f \in \mathbb{F}$ and loss function ℓ , we estimate the risk $\rho(F(\cdot; f))$ using the CDF estimator $\hat{F}(\cdot; f)$ by plugging it in the functional of interest: $\rho(\hat{F}(\cdot; f))$. Many existing risk estimators in the supervised learning literatures, including the traditional empirical risk and estimators used for estimating spectral risks (Khim et al., 2020) and optimized certainty equivalent risks (Lee et al., 2020) can be viewed as examples of the above estimator.

¹Quasi-metrics are defined in Appendix A.

Leveraging the uniform convergence results of the CDF estimator, we present uniform convergence result of the proposed risk estimator. The uniform convergence holds both over the hypothesis class \mathbb{F} and over a set of Hölder risk functionals. As the Hölder class contains a large set of popularly studied risks, our result demonstrates that models can be assessed under many risks without loss of statistical power. This is formalized in Theorem 5.1 below.

Theorem 5.1. *For a hypothesis class \mathbb{F} , a bounded loss function ℓ , and $\delta \in (0, 1]$, if $\mathbb{P}(e_n(\mathbb{F}, \ell) \leq \epsilon) \geq 1 - \delta$, then with probability $1 - \delta$, for all $\rho \in \mathbb{T}$, we have*

$$\sup_{f \in \mathbb{F}} |\rho(F(\cdot; f)) - \rho(\widehat{F}(\cdot; f))| \leq L(\rho, 1, \mathcal{L}_\infty)\epsilon,$$

where \mathbb{T} is the set of $L(\cdot, 1, \mathcal{L}_\infty)$ Hölder risk functionals on the space of bounded random variables.

In cases where the CDF estimation error ϵ is of order $O(1/\sqrt{n})$, we can estimate the set of Hölder risks \mathbb{T} for all hypotheses in \mathbb{F} at rate $O(1/\sqrt{n})$. Because our result is uniform over both the hypothesis class \mathbb{F} and the risk functional class \mathbb{T} , it is a generalization of existing uniform convergence results that are uniform over \mathbb{F} , but for a single risk functional (Khim et al., 2020; Lee et al., 2020).

In Appendix C, we provide similar uniform convergence results where the set of risk functionals are Hölder smooth in Wasserstein distance. As a direct consequence to Theorem 5.1, uniform convergence of a single Hölder risk, e.g., a distortion risk (4), is given below.

Corollary 5.1. *For a hypothesis class \mathbb{F} , a bounded loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, D]$, and $\delta \in (0, 1]$, if $\mathbb{P}(e_n(\mathbb{F}, \ell) \leq \epsilon) \geq 1 - \delta$, then with probability $1 - \delta$, we have*

$$\sup_{f \in \mathbb{F}} |\rho(F(\cdot; f)) - \rho(\widehat{F}(\cdot; f))| \leq L\epsilon,$$

where ρ is a distortion risk with $\frac{L}{D}$ -Lipschitz distortion function.

Remark 5.1. *As an example, consider a binary classification setting where the loss function ℓ is the 0/1 loss and the hypothesis class \mathbb{F} is finite, using our results in Corollary 5.1, we obtain that the generation error for the expected value is $O(\sqrt{\log(|\mathbb{F}|)/n})$ and the generation error for the CVaR is $O(\sqrt{\log(|\mathbb{F}|)/\alpha^2 n})$, which are of the same rates (with better constants) as the ones in Lee et al. (2020).*

6. Empirical Risk Minimization

For a single risk functional, one may want to learn models that optimize it. Our uniform convergence results for risk estimation license us to learn models that minimize the population risk through Empirical Risk Minimization (ERM). We denote the population and empirical risk minimizers as

follows:

$$f_\star \in \arg \min_{f \in \mathbb{F}} \rho(F(\cdot; f)), \quad \widehat{f}_\star \in \arg \min_{f \in \mathbb{F}} \rho(\widehat{F}(\cdot; f)). \quad (5)$$

The excess risk of the empirical risk minimizer \widehat{f}_\star can be bounded by

$$\begin{aligned} \rho(F(\cdot; \widehat{f}_\star)) - \rho(F(\cdot; f_\star)) &= \rho(F(\cdot; \widehat{f}_\star)) - \rho(\widehat{F}(\cdot; \widehat{f}_\star)) \\ &+ \rho(\widehat{F}(\cdot; \widehat{f}_\star)) - \rho(\widehat{F}(\cdot; f_\star)) + \rho(\widehat{F}(\cdot; f_\star)) - \rho(F(\cdot; f_\star)) \\ &\leq 2 \sup_{f \in \mathbb{F}} |\rho(\widehat{F}(\cdot; f)) - \rho(F(\cdot; f))|. \end{aligned}$$

We study ERM when the loss function is non-negative and the risk functional of interest is a distortion risk with a Lipschitz distortion function (4). Such distortion risk functionals consist many well-studied risks, including the expected value, CVaR, cumulative prospect theory risks, and other spectral risks (Bäuerle & Glauner, 2021). Using Corollary 5.1, we obtain that when $e_n(\mathbb{F}, \ell) = O(1/\sqrt{n})$ and the loss ℓ is bounded, the excess risk of \widehat{f}_\star is $O(1/\sqrt{n})$.

We consider settings where the hypothesis class \mathbb{F} is a class of parameterized functions, e.g., linear models and neural networks and use $\Theta \subseteq \mathbb{R}^d$ to denote the set of parameters. For a hypothesis $f \in \mathbb{F}$ parameterized by $\theta \in \Theta$, we denote $F(\cdot; f)$ and $\ell_f(z_i)$ by F_θ and $\ell_\theta(i)$, respectively. Similarly, we use θ_\star and $\widehat{\theta}_\star$ for referring to f_\star and \widehat{f}_\star . As in Section 4.1, we use $\pi_\theta : [n] \rightarrow [n]$ to denote the permutation function such that $\ell_\theta(\pi_\theta(i))$ is the i -th smallest loss under the current model θ and the fixed dataset $\{z_i\}_{i=1}^n$. Using the CDF estimator \widehat{F}_θ (1), the empirical distortion risk $\rho(\widehat{F}_\theta)$ can be re-written as

$$\sum_{i=1}^n g \left(1 - \frac{i-1}{n} \right) \cdot (\ell_\theta(\pi_\theta(i)) - \ell_\theta(\pi_\theta(i-1))),$$

where for all $\theta \in \Theta$, we set $\ell_\theta(\pi_\theta(0)) := 0$ since the losses are non-negative.

To employ first-order methods for minimizing the empirical distortion risk, it is natural to first identify when $\rho(\widehat{F}_\theta)$ is differentiable.

Lemma 6.1. *If $\{\ell_\theta(z_i)\}_{i=1}^n$ are Lipschitz continuous in $\theta \in \Theta$, then for all $i \in [n]$, $\ell_\theta(\pi_\theta(i))$, i.e., the i -th smallest loss, is Lipschitz continuous in θ and $\rho(\widehat{F}_\theta)$ is differentiable in θ almost everywhere.*

When $\rho(\widehat{F}_\theta)$ (and $\ell_\theta(\pi_\theta(i))$) is differentiable, the gradient $\nabla_\theta \rho(\widehat{F}_\theta)$ can be written as

$$\sum_{i=1}^n \left(g \left(1 - \frac{i-1}{n} \right) - g \left(1 - \frac{i}{n} \right) \right) \cdot \nabla_\theta \ell_\theta(\pi_\theta(i)). \quad (6)$$

	VGG-11	GoogLeNet	ShuffleNet	Inception	ResNet-18
Accuracy	69.022%	69.772%	69.356%	69.542%	69.756%
$\mathbb{E}[\ell_f]$	1.261	1.283	1.360	1.829	1.247
$\text{CVaR}_{.05}(\ell_f)$	1.327	1.350	1.431	1.925	1.313
$\mathbb{E}[\ell_f] + 0.5\text{Var}(\ell_f)$	5.215	4.376	6.718	14.416	5.353
$\text{HRM}_{.3,.4}(\ell_f)$	1.374	1.336	1.542	2.214	1.382
$\text{HRM}_{.2,.8}(\ell_f)$	1.233	1.239	1.344	1.845	1.225

Table 1. Risks for different ImageNet classification models evaluated on the validation set. $\ell_f(Z)$ is the cross-entropy loss for each model f . For simplicity, we omitted the arguments Z in the table. CVaR_α is the expected value of the top 100 α percent losses. $\text{HRM}_{a,b}$ is the cumulative prospect theory risk defined in Leqi et al. (2019). All results are rounded to 3 digits.

When $g(x) = x$, the distortion risk is the same as the expected loss and we recover the gradient for the traditional empirical risk.

To avoid the non-differentiable points, we add a small noise to the gradient descent steps. By doing so, we ensure that the descent steps will end up in differentiable points almost surely. Choose initial point $\theta_1 \in \Theta$. At iteration t , the parameter is updated as follows

$$\theta_{t+1} \leftarrow \theta_t - \eta \left(\nabla_{\theta} \rho(\widehat{F}_{\theta}) + w_t \right), \quad (7)$$

where η is the learning rate, $\nabla_{\theta} \rho(\widehat{F}_{\theta})$ is given in (6) and w_t is sampled from a d -dimensional Gaussian with mean 0 and variance $\frac{1}{d}$.

In general, even when the loss function ℓ is convex in the parameter θ , the empirical distortion risk $\rho(\widehat{F}_{\theta})$ may not be convex in θ . In Corollary 6.1, we show local convergence of θ_t obtained through following (7).

Corollary 6.1. *If $\{\ell_{\theta}(z_i)\}_{i=1}^n$ are Lipschitz continuous and $\rho(\widehat{F}_{\theta})$ is β -smooth in θ , then the following holds almost surely when the learning rate in (7) is $\eta = \frac{1}{\beta\sqrt{T}}$:*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\|\nabla_{\theta} \rho(\widehat{F}_{\theta_t})\|^2 \right] \leq \frac{2\beta}{\sqrt{T}} \left(\rho(\widehat{F}_{\theta_1}) - \rho(\widehat{F}_{\theta_*}) + \frac{1}{2\beta} \right).$$

Corollary 6.1 demonstrates that the average gradient magnitude over T iterations shrinks as T goes to infinity, suggesting that performing gradient descent by following (7) will converge to an approximate stationary point.

7. Experiment

In our experiments, we demonstrate the efficacy of our proposed estimator for risk estimation and proposed learning procedure for obtaining risk-sensitive models. In Section 7.1, we work with a risk assessment setting where we simultaneously inspect a finite set of models in terms of multiple risks. In Section 7.2, we show the performance of empirical risk minimization under various distortion risks.

After showing that the classifier learned under different risk objectives behave differently in a toy example, we learn risk-sensitive models for CIFAR-10.

7.1. Risk Assessment on ImageNet Models

We perform risk assessments on pretrained Pytorch models for ImageNet classification. In particular, we choose models with similar accuracy (both reported on the official Pytorch website and confirmed by us) on the validation set (50,000 images) for the ImageNet classification challenge (Rusakovsky et al., 2015). The models are VGG-11 (Simonyan & Zisserman, 2014), GoogLeNet (Szegedy et al., 2015), ShuffleNet (Ma et al., 2018), Inception (Szegedy et al., 2016) and ResNet-18 (He et al., 2016) and the accuracy of these models evaluated on the validation set are around 69% (Table 1). By assessing the risks of models with similar accuracy, we highlight how models with similar performance under traditional metrics (e.g., accuracy) could have different risk performances. For example, though Inception has similar accuracy compared to other models, its loss variance is much higher compared to others, which may be detrimental in settings where high-varying performance is not preferred. We also evaluated the CVaR of these models under different α 's (Figure 3 in Appendix E). Our theoretical results suggest that all these evaluations hold simultaneously across the risk functionals and models of interest with the error being $O(\sqrt{\log |\mathbb{F}|/n})$ ($|\mathbb{F}| = 5$ in this experiment). In addition to showcasing the power of our theoretical results, this example demonstrates how model assessments under multiple risk notions provide a better understanding of model behaviors.

7.2. Empirical Distortion Risk Minimization

To illustrate the difference among models learned under different risk objectives, we first present a toy example for comparing models learned under the expected loss objective and the CVaR objective respectively. We then show the efficacy of our proposed optimization procedure through training deep neural networks on CIFAR-10. In both cases, the models are learned by following (7). For more details on these




Figure 1. Prediction (predicted probability of a covariate being labeled as 1) contours and loss histograms of two models learned under the expected loss and the $\text{CVaR}_{0.05}$ objective, respectively. The blue pluses and orange dots represent two classes. The loss distribution for the expected loss model has extremely high values for a small subset of the covariates.

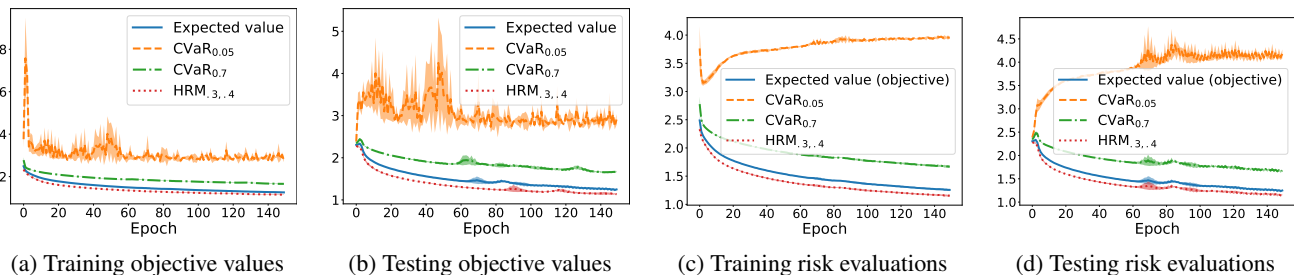


Figure 2. Performance of VGG-16 models trained under expected loss, $\text{CVaR}_{0.05}$, $\text{CVaR}_{0.7}$ and $\text{HRM}_{3,4}$. In Figures 2a and 2b, each model is trained and evaluated on the same objective. In Figures 2c and 2d, we only train one model under the expected loss but report all four objectives of that model. All results are averaged over 5 runs.

experiments, we refer the readers to Appendix E.

Toy Example In the toy example, we work with a binary classification task where the covariates are 2-dimensional. In Figure 1, the blue pluses and orange dots represent two classes, respectively. We have learned logistic regression models to minimize the expected loss and the $\text{CVaR}_{0.05}$ (expected value of the top 5% losses) through minimizing their empirical risks. The loss distribution along with the prediction contours of the two classifiers showcase the difference between the two models. In particular, the model learned under expected loss suffers high loss for a small subset of the covariates while the model learned under $\text{CVaR}_{0.05}$ have all losses concentrated around a small value. Indicated by the (uniform) grey color in the contour plot, the predictions (predicted probability of a covariate being labeled as 1) for the $\text{CVaR}_{0.05}$ model are around 0.5. In contrast, the predictions for the expected loss model spread across a wide range between 0 and 1.

CIFAR-10 We have trained VGG-16 models on CIFAR-10 through minimizing the empirical risks for expected loss, $\text{CVaR}_{0.05}$, $\text{CVaR}_{0.7}$ and $\text{HRM}_{3,4}$ (Leqi et al., 2019) using the gradient descent step presented in (7). The models are trained over 150 epochs and the learning rate is chosen to be 0.005. As shown in Figure 2a and 2b, in general, the

objective values are decreasing over the epochs during training and testing. In addition, we observe that minimizing the empirical risk for expected loss does not necessarily imply minimizing other risks, e.g., $\text{CVaR}_{0.05}$ (Figure 2c). These results suggest the efficacy of our proposed optimization procedure for minimizing distortion risks.

8. Discussion

We have presented a principled framework, including analytic tools and algorithms for risk-sensitive learning and assessment that: (1) obtains the empirical CDF; (2) estimates the risks of interest through plugging in the empirical CDF; and (3) minimizes the empirical risk (for risk-sensitive learning). Our theoretical results on the uniform convergence of the proposed risk estimators hold simultaneously over a hypothesis class (constrained by an appropriate complexity measure) and over Hölder risks. The key building block for these results is the uniform convergence of the CDF estimator.

There are multiple future directions of our work. First, we hope to more precisely characterize the permutation complexity (under various hypothesis classes). Second, our gradient descent procedure (7) requires sorting all losses. An important next step would be to allow minibatches (sorting only a small subset of losses) when minimizing empirical distortion risks. Third, as shown in Figure 1, models learned

under different risk objectives behave distinctly. Characterizing these model behaviors theoretically and empirically, and understanding the trade-offs among these objectives is crucial for building future models.

Acknowledgements

LL is generously supported by an Open Philanthropy AI Fellowship.

References

- Acerbi, C. Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking & Finance*, 26(7):1505–1518, 2002.
- Alexander, K. S. Probability inequalities for empirical processes and a law of the iterated logarithm. *The Annals of Probability*, pp. 1041–1067, 1984.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.
- Balbás, A., Garrido, J., and Mayoral, S. Properties of distortion risk measures. *Methodology and Computing in Applied Probability*, 11(3):385–399, 2009.
- Bartlett, P. L. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Bäuerle, N. and Glauner, A. Minimizing spectral risk measures applied to markov decision processes. *Mathematical Methods of Operations Research*, pp. 1–35, 2021.
- Ben-Tal, A. and Teboulle, M. Expected utility, penalty functions, and duality in stochastic nonlinear programming. *Management Science*, 32(11):1445–1466, 1986.
- Bhat, S. P. and LA, P. Concentration of risk measures: A wasserstein distance approach. *Advances in Neural Information Processing Systems*, 32:11762–11771, 2019.
- Björk, T., Murgoci, A., and Zhou, X. Y. Mean–variance portfolio optimization with state-dependent risk aversion. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 24(1):1–24, 2014.
- Bottou, L., Curtis, F. E., and Nocedal, J. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- Boucheron, S., Lugosi, G., and Massart, P. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Cassel, A., Mannor, S., and Zeevi, A. A general framework for bandit problems beyond cumulative objectives.
- Chandak, Y., Shankar, S., and Thomas, P. S. High-confidence off-policy (or counterfactual) variance estimation, 2021.
- Denneberg, D. Distorted probabilities and insurance premiums. *Methods of Operations Research*, 63(3):3–5, 1990.
- Devroye, L., Györfi, L., and Lugosi, G. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- Duchi, J. and Namkoong, H. Learning models with uniform performance via distributionally robust optimization. *Annals of Statistics*, 2018.
- Gänssler, P. and Stute, W. Empirical processes: a survey of results for independent and identically distributed random variables. *The Annals of Probability*, pp. 193–243, 1979.
- Gopalan, A., Prashanth, L., Fu, M., and Marcus, S. Weighted bandits or: How bandits learn distorted values that are not expected. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- Guide, A. H. *Infinite dimensional analysis*. Springer, 2006.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Holland, M. J. and Haress, E. M. Spectral risk-based learning using unbounded losses. *arXiv preprint arXiv:2105.04816*, 2021.
- Huang, A., Leqi, L., Lipton, Z. C., and Azizzadenesheli, K. Off-policy risk assessment in contextual bandits. *arXiv preprint arXiv:2104.08977*, 2021.
- Jie, C., Prashanth, L., Fu, M., Marcus, S., and Szepesvári, C. Stochastic optimization in a cumulative prospect theory framework. *IEEE Transactions on Automatic Control*, 63(9):2867–2882, 2018.
- Khim, J., Leqi, L., Prasad, A., and Ravikumar, P. Uniform convergence of rank-weighted learning. In *International Conference on Machine Learning*, pp. 5254–5263. PMLR, 2020.
- Krokhmal, P. A. Higher moment coherent risk measures. 2007.
- Kusuoka, S. On law invariant coherent risk measures. In *Advances in mathematical economics*, pp. 83–95. Springer, 2001.

- Lee, J., Park, S., and Shin, J. Learning bounds for risk-sensitive learning. In *34th Conference on Neural Information Processing Systems (NeurIPS) 2020*. Neural Information Processing Systems, 2020.
- Leqi, L., Prasad, A., and Ravikumar, P. On human-aligned risk minimization. In *Advances in Neural Information Processing Systems*, 2019.
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. Tilted empirical risk minimization. In *International Conference on Learning Representations*, 2020.
- Ma, N., Zhang, X., Zheng, H.-T., and Sun, J. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 116–131, 2018.
- Massart, P. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, pp. 1269–1283, 1990.
- Massart, P. Some applications of concentration inequalities to statistics. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 9, pp. 245–303, 2000.
- Munos, R. From bandits to monte-carlo tree search: The optimistic principle applied to optimization and planning. 2014.
- Oliveira, L. P. R. and Yang, S. S. Topics in statistical mathematics: Lecture 2. https://w3.impa.br/~rimfo/estatistica_a16/notas_alunos/shangjie.pdf. Accessed: 2021-9-15.
- Prashanth, L., Jie, C., Fu, M., Marcus, S., and Szepesvári, C. Cumulative prospect theory meets reinforcement learning: Prediction and control. In *International Conference on Machine Learning*, pp. 1406–1415. PMLR, 2016.
- Rockafellar, R. T. and Wets, R. J.-B. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- Rockafellar, R. T., Uryasev, S., et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi:10.1007/s11263-015-0816-y.
- Sani, A., Lazaric, A., and Munos, R. Risk-aversion in multi-armed bandits. *arXiv preprint arXiv:1301.1936*, 2013.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. *Lectures on stochastic programming: modeling and theory*. SIAM, 2014.
- Sharpe, W. F. Mutual fund performance. *The Journal of business*, 39(1):119–138, 1966.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Torossian, L., Garivier, A., and Picheny, V. \mathcal{X} -armed bandits: Optimizing quantiles, cvar and other risks. In *Asian Conference on Machine Learning*, pp. 252–267. PMLR, 2019.
- Vakili, S. and Zhao, Q. Mean-variance and value at risk in multi-armed bandit problems. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1330–1335. IEEE, 2015.
- Vallender, S. Calculation of the wasserstein distance between probability distributions on the line. *Theory of Probability & Its Applications*, 18(4):784–786, 1974.
- Vapnik, V. *Estimation of dependences based on empirical data*. Springer Science & Business Media, 2006.
- Vapnik, V. *The nature of statistical learning theory*. Springer science & business media, 2013.
- Vapnik, V. N. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999.
- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Wang, S. Premium calculation by transforming the layer premium density. *ASTIN Bulletin: The Journal of the IAA*, 26(1):71–92, 1996.
- Wang, S. S., Young, V. R., and Panjer, H. H. Axiomatic characterization of insurance prices. *Insurance: Mathematics and economics*, 21(2):173–183, 1997.

Wirch, J. L. and Hardy, M. R. A synthesis of risk measures for capital adequacy. *Insurance: mathematics and Economics*, 25(3):337–347, 1999.

Wirch, J. L. and Hardy, M. R. Distortion risk measures: Coherence and stochastic dominance. In *International congress on insurance: Mathematics and economics*, pp. 15–17, 2001.

Zimin, A., Ibsen-Jensen, R., and Chatterjee, K. Generalized risk-aversion in stochastic multi-armed bandits. *arXiv preprint arXiv:1405.0833*, 2014.

A. Hölder Risk Functionals

In the definition of Hölder risk functionals, we require d to be a quasi-metric, which we provide the definition here.

Definition A.1. A function $d : \mathcal{L}_\infty(\mathbb{R}, \mathbb{B}(\mathbb{R})) \times \mathcal{L}_\infty(\mathbb{R}, \mathbb{B}(\mathbb{R})) \rightarrow [0, +\infty)$ is a quasi-metric if the following two conditions hold:

- For all $F_U, F_{U'} \in \mathcal{L}_\infty(\mathbb{R}, \mathbb{B}(\mathbb{R}))$, $d(F_U, F_{U'}) = 0$ if and only if $F_U = F_{U'}$;
- For all $F_U, F_{U'}, F_{Z''} \in \mathcal{L}_\infty(\mathbb{R}, \mathbb{B}(\mathbb{R}))$, $d(F_U, F_{Z''}) \leq d(F_U, F_{U'}) + d(F_{U'}, F_{Z''})$.

If a quasimetric is symmetric, i.e., for all $F_U, F_{U'} \in \mathcal{L}_\infty(\mathbb{R}, \mathbb{B}(\mathbb{R}))$, $d(F_U, F_{U'}) = d(F_{U'}, F_U)$, it is also a *metric*. The set of quasi-metrics contains symmetric quasi-metrics, e.g., sup norms \mathcal{L}_∞ , Wasserstein distance, along with non-symmetric quasi-metrics, e.g., Kullback-Leibler divergence.

We will now discuss why optimized certainty equivalent (OCE) risks (e.g., mean-variance, entropic risk, CVaR) and spectral risks with bounded spectrum (e.g., CVaR, certain CPT-inspired Risks) are Lipschitz on bounded random variables. OCE risks, first introduced by Ben-Tal & Teboulle (1986), are defined as

$$\rho_{\text{oce}}(F_U) := \inf_{\lambda \in \mathbb{R}} \{ \lambda + \mathbb{E}[\phi(U - \lambda)] \},$$

where $\phi : \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$ is a nondecreasing, closed and convex function with $\phi(0) = 0$ and $1 \in \partial\phi(0)$. To complement the risk-averse OCEs, a risk-seeking version (inverted OCE) is proposed:

$$\rho_{\text{oce}}^{\text{inv}}(F_U) := \sup_{\lambda \in \mathbb{R}} \{ \lambda - \mathbb{E}[\phi(\lambda - U)] \}.$$

Proposition A.1. If ϕ is continuously differentiable, then the OCE risks ρ_{oce} and inverted OCE risks $\rho_{\text{oce}}^{\text{inv}}$ are Lipschitz on the space of bounded random variables with support $[0, D]$:

$$\begin{aligned} |\rho_{\text{oce}}(F_U) - \rho_{\text{oce}}(F_{U'})| &\leq \max_{x \in [0, D]} (\phi(D - x) - \phi(-x)) \|F_U - F_{U'}\|_\infty, \\ |\rho_{\text{oce}}^{\text{inv}}(F_U) - \rho_{\text{oce}}^{\text{inv}}(F_{U'})| &\leq \max_{x \in [0, D]} (\phi(x - D) - \phi(x)) \|F_U - F_{U'}\|_\infty. \end{aligned}$$

Remark A.1. Similar to Huang et al. (2021, Lemma 4.1), Proposition A.1 shows that the expected value and CVaR_α are D - and $\frac{D}{\alpha}$ -Lipschitz on random variables with support $[0, D]$, respectively. In addition, this result also provides Lipschitzness of the entropic risks (since the corresponding ϕ is continuously differentiable (Lee et al., 2020, Table 1)) and other OCE and inverted OCE risks that do not belong to distortion risk functionals.

Proof. When U has support $[0, D]$, as shown in Lee et al. (2020, Lemma 9), we can re-write the OCE and inverted OCE risks as follows:

$$\begin{aligned} \rho_{\text{oce}}(F_U) &= \min_{\lambda \in [0, D]} \{ \lambda + \mathbb{E}[\phi(U - \lambda)] \} \\ \rho_{\text{oce}}^{\text{inv}}(F_U) &= \max_{\lambda \in [0, D]} \{ \lambda - \mathbb{E}[\phi(\lambda - U)] \}. \end{aligned}$$

For any $U, U' \in [0, D]$, denote $\lambda_U \in \arg \min_{\lambda \in [0, D]} \lambda + \mathbb{E}_U[\phi(U - \lambda)]$ and $\lambda_{U'} \in \arg \min_{\lambda \in [0, D]} \lambda + \mathbb{E}_{U'}[\phi(U' - \lambda)]$.

Consider the case where $\rho_{\text{oce}}(F_U) < \rho_{\text{oce}}(F_{U'})$.

$$\begin{aligned}
 |\rho_{\text{oce}}(F_{U'}) - \rho_{\text{oce}}(F_U)| &= \rho_{\text{oce}}(F_{U'}) - \rho_{\text{oce}}(F_U) \\
 &= \lambda_{U'} + \mathbb{E}_{U'}[\phi(U' - \lambda_{U'})] - \lambda_U - \mathbb{E}_U[\phi(U - \lambda_U)] \\
 &\stackrel{(i)}{\leq} \lambda_U + \mathbb{E}_{U'}[\phi(U' - \lambda_U)] - \lambda_U - \mathbb{E}_U[\phi(U - \lambda_U)] \\
 &= \int_0^D \phi(u - \lambda_U) d(F_{U'}(u) - F_U(u)) \\
 &\stackrel{(ii)}{=} \phi(u - \lambda_U) (F_{U'}(u) - F_U(u)) \Big|_{u=0}^{u=D} - \int_0^D \phi'(u - \lambda_U) (F_{U'}(u) - F_U(u)) du \\
 &\stackrel{(iii)}{\leq} \|F_U - F_{U'}\|_{\infty} \int_0^D \phi'(u - \lambda_U) du \\
 &= (\phi(D - \lambda_U) - \phi(-\lambda_U)) \|F_U - F_{U'}\|_{\infty},
 \end{aligned}$$

where (i) comes from the definition of $\lambda_{U'}$, (ii) uses integration by parts, and (iii) uses the fact that ϕ is non-decreasing, i.e., ϕ' is non-negative, and $F_U(0) = F_{U'}(0) = 0$, $F_U(D) = F_{U'}(D) = 1$. The case when $\rho_{\text{oce}}(F_{U'}) < \rho_{\text{oce}}(F_U)$ proceeds similarly:

$$\begin{aligned}
 |\rho_{\text{oce}}(F_U) - \rho_{\text{oce}}(F_{U'})| &= \rho_{\text{oce}}(F_U) - \rho_{\text{oce}}(F_{U'}) \\
 &\leq \int_0^D \phi(u - \lambda_{U'}) d(F_U(u) - dF_{U'}(u)) \\
 &= \phi(u - \lambda_{U'}) (F_U(u) - F_{U'}(u)) \Big|_{u=0}^{u=D} - \int_0^D \phi'(u - \lambda_{U'}) (F_U(u) - F_{U'}(u)) du \\
 &\leq (\phi(D - \lambda_{U'}) - \phi(-\lambda_{U'})) \|F_U - F_{U'}\|_{\infty}.
 \end{aligned}$$

Putting it together, we have that

$$|\rho_{\text{oce}}(F_U) - \rho_{\text{oce}}(F_{U'})| \leq \max_{\lambda \in [0, D]} (\phi(D - x) - \phi(-x)) \|F_U - F_{U'}\|_{\infty}.$$

For inverted OCE risks, denote $\lambda_U \in \arg \max_{\lambda \in [0, D]} \lambda + \mathbb{E}_U[\phi(\lambda - U)]$ and $\lambda_{U'} \in \arg \max_{\lambda \in [0, D]} \lambda + \mathbb{E}_{U'}[\phi(\lambda - U)]$. The proof proceeds similarly by using the fact that $\rho_{\text{occe}}(F_{U'}) - \rho_{\text{occe}}(F_U) \leq \mathbb{E}_U[\phi(\lambda_{U'} - U)] - \mathbb{E}_{U'}[\phi(\lambda_{U'} - U)]$ and $\rho_{\text{occe}}(F_U) - \rho_{\text{occe}}(F_{U'}) \leq \mathbb{E}_{U'}[\phi(\lambda_U - U')] - \mathbb{E}_U[\phi(\lambda_U - U)]$. Following similar steps, we obtain that

$$|\rho_{\text{occe}}(F_U) - \rho_{\text{occe}}(F_{U'})| \leq \max_{\lambda \in [0, D]} (\phi(x - D) - \phi(x)) \|F_U - F_{U'}\|_{\infty}.$$

□

Spectral risks (also known as L -risks or rank-weighted risks) are a subset of distortion risk functionals. As noted in Bäuerle & Glauner (2021), a spectral risk can be written as a distortion risk (Equation (4)) with the following distortion function: for $t \in [0, 1]$,

$$g(t) = \int_0^t h(s) ds,$$

where $h : [0, 1] \rightarrow \mathbb{R}_+$ is the non-decreasing spectrum function that integrates to 1. Since g is Lipschitz when h is bounded (i.e., $g'(t) = h(t)$ for $t \in (0, 1)$), spectral risks are Lipschitz on the space of bounded random variables when their spectrum is bounded.

Finally, examples of risk functionals that are Hölder but are not Lipschitz include distortion risks whose distortion functions are Hölder but not Lipschitz.

B. Proof of Results in Section 4

We note that in the following proofs, Z is used to denote a generic random variable and the loss function is denoted by $\ell : \mathcal{X} \times \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$. (The loss function presented in the main text is a special case of this.) For a given $(X, Y, f(X))$, the loss is denoted by $\ell(X, Y, f(X))$.

B.1. Auxillary Lemmas

The below two auxillary lemmas are mainly adaptations to the class note from Professor Roberto Imbuzeiro Oliveira (Oliveira & Yang).

Lemma B.1. *Let $\mathbb{G}(1) := \{g(\cdot; r) := \mathbb{1}_{\{\cdot \leq r\}} : \forall r \in \mathbb{R}\}$. For a fixed $f \in \mathbb{F}$, iid sample $\{X_i, Y_i\}_{i=1}^n$ with joint probability measure \mathbb{P} , and a loss function $\ell : \mathcal{X} \times \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, we have that*

$$\sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| = \max_{j \in [n]} \left| \sum_{i=1}^j \xi_i \right|, \quad (8)$$

and further

$$\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| \right) \right] \leq 2 \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_{j \in [n]} \left(\exp \left(\frac{\lambda}{n} \sum_{i=1}^j \xi_i \right) \mathbb{1}_{\{\sum_{i=1}^j \xi_i \geq 0\}} \right) \right]. \quad (9)$$

Remark B.1. *We note that an important property of Lemma B.1 is that the bound is independent of the samples $\{X_i, Y_i\}_{i=1}^n$.*

Proof. Let $\{R_i\}_{i=1}^n$ denote the sorted sequence of $\{\ell(X_i, Y_i, f(X_i))\}_{i=1}^n$, where $R_1 \leq R_2 \leq \dots \leq R_n$. Using $\{R_i\}_{i=1}^n$, we have

$$\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| \right) \right] = \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(R_i) \right| \right) \right]$$

Consider a function $g(t; r) = \mathbb{1}_{\{t \leq r\}}$. For such a function, $\sum_{i=1}^n \xi_i g(R_i; r)$ is equal to

- 0 if $r < \min_{i \in [n]} R_i$,
- $\sum_{i=1}^j \xi_i$ when $R_j \leq r < R_{j+1}$ for some $j \in \{1, \dots, n-1\}$,
- $\sum_{i=1}^n \xi_i$ otherwise.

Therefore, we have that

$$\sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| = \max_{j \in [n]} \left| \sum_{i=1}^j \xi_i \right|.$$

Finally, we notice that

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\sup_{g \in \mathbb{G}(1)} \exp \left(\frac{\lambda}{n} \left| \sum_{i=1}^n \xi_i g(R_i) \right| \right) \right] = \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_{j \in [n]} \exp \left(\frac{\lambda}{n} \left| \sum_{i=1}^j \xi_i \right| \right) \right] \\ &= \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_{j \in [n]} \left(\exp \left(\frac{\lambda}{n} \sum_{i=1}^j \xi_i \right) \mathbb{1}_{\{\sum_{i=1}^j \xi_i \geq 0\}} + \exp \left(-\frac{\lambda}{n} \sum_{i=1}^j \xi_i \right) \mathbb{1}_{\{\sum_{i=1}^j \xi_i < 0\}} \right) \right] \\ &= \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_{j \in [n]} \left(\exp \left(\frac{\lambda}{n} \sum_{i=1}^j \xi_i \right) \mathbb{1}_{\{\sum_{i=1}^j \xi_i \geq 0\}} \right) \right] + \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_{j \in [n]} \left(\exp \left(-\frac{\lambda}{n} \sum_{i=1}^j \xi_i \right) \mathbb{1}_{\{\sum_{i=1}^j \xi_i < 0\}} \right) \right] \\ &\leq 2 \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_{j \in [n]} \left(\exp \left(\frac{\lambda}{n} \sum_{i=1}^j \xi_i \right) \mathbb{1}_{\{\sum_{i=1}^j \xi_i \geq 0\}} \right) \right]. \end{aligned} \quad (10)$$

□

Lemma B.2. Let $\{Z_i\}_{i=1}^n$ taking values in \mathcal{Z} denote independent samples drawn from \mathbb{P} and $w : \mathcal{Z} \rightarrow \mathbb{R}_+$. For n independent Rademacher random variables $\{\xi_i\}_{i=1}^n$, we have that for all $\lambda \geq 0$,

$$\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\max_{j \in [n]} \frac{\lambda}{n} \sum_{i=1}^j w(Z_i) \xi_i \right) \right] \leq 2 \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sum_{i=1}^n w(Z_i) \xi_i \right) \right].$$

Further, if $\frac{1}{n} \sum_{i=1}^n w(Z_i) \xi_i$ is mean zero $\frac{\gamma^2}{n}$ -subGaussian, then

$$\mathbb{E}_{\mathfrak{R}} \left[\exp \left(\max_{j \in [n]} \frac{\lambda}{n} \sum_{i=1}^j \xi_i \right) \right] \leq 2 \exp \left(\frac{\lambda^2 \gamma^2}{2n} \right).$$

Remark B.2. We note that an immediate consequence of Lemma B.2 is

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_{j \in [n]} \left(\exp \left(\frac{\lambda}{n} \sum_{i=1}^j w(Z_i) \xi_i \right) \mathbb{1}_{\{\sum_{i=1}^j w(Z_i) \xi_i \geq 0\}} \right) \right] \\ & \leq \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\max_{j \in [n]} \frac{\lambda}{n} \sum_{i=1}^j w(Z_i) \xi_i \right) \right] \leq 2 \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sum_{i=1}^n w(Z_i) \xi_i \right) \right]. \end{aligned}$$

Proof. The proof contains two main steps:

Step 1: We will show that for all $t > 0$,

$$\mathbb{P} \left(\max_{j \in [n]} \sum_{i=1}^j w(Z_i) \xi_i \geq t \right) \leq 2 \mathbb{P} \left(\sum_{i=1}^n w(Z_i) \xi_i \geq t \right). \quad (11)$$

To show (11), for $t > 0$, consider events $E_0 := \emptyset$ and $E_j := \{\sum_{i=1}^j w(Z_i) \xi_i \geq t, \sum_{i=1}^l w(Z_i) \xi_i < t, \forall l < j\}$ for $j \in [n]$, which states that j is the first index such that the partial sum $\sum_{i=1}^j w(Z_i) \xi_i$ is at least t . We first notice that

$$\left\{ \max_j \sum_{i=1}^j w(Z_i) \xi_i \geq t \right\} \subset \bigcup_{j=0}^n E_j.$$

Since E_j and $\sum_{i=j+1}^n w(Z_i) \xi_i \geq 0$ implies that $\sum_{i=1}^n w(Z_i) \xi_i \geq t$, we obtain

$$\bigcup_{j=0}^n \left(E_j \cap \left\{ \sum_{i=j+1}^n w(Z_i) \xi_i \geq 0 \right\} \right) \subset \left\{ \sum_{i=1}^n w(Z_i) \xi_i \geq t \right\}. \quad (12)$$

Moreover, for any $j \in [n]$, we have

$$\mathbb{P} \left(\sum_{i=j+1}^n w(Z_i) \xi_i \geq 0 \right) \geq \frac{1}{2}, \quad (13)$$

since $\sum_{i=j+1}^n w(Z_i) \xi_i$ is symmetric around 0 (i.e., $\sum_{i=j+1}^n w(Z_i) \xi_i \stackrel{d}{=} -\sum_{i=j+1}^n w(Z_i) \xi_i$) for all $i \in \{0, \dots, n\}$. For any $j \in [n]$, since the event E_j (depending on $\{w(Z_i), \xi_i\}_{i \leq j}$) is independent of $\{\sum_{i=j+1}^n w(Z_i) \xi_i \geq 0\}$, we have

$$\mathbb{P} \left(E_j \cap \left\{ \sum_{i=j+1}^n w(Z_i) \xi_i \geq 0 \right\} \right) = \mathbb{P}(E_j) \mathbb{P} \left(\sum_{i=j+1}^n w(Z_i) \xi_i \geq 0 \right) \geq \frac{\mathbb{P}(E_j)}{2}.$$

As a result, we have

$$\begin{aligned}
 \mathbb{P}\left(\sum_{i=1}^n w(Z_i)\xi_i \geq t\right) &\geq \mathbb{P}\left(\bigcup_{j=0}^n \left(E_j \cap \left\{\sum_{i=j+1}^n w(Z_i)\xi_i \geq 0\right\}\right)\right) \\
 &= \sum_{j=0}^n \mathbb{P}\left(E_j \cap \left\{\sum_{i=j+1}^n w(Z_i)\xi_i \geq 0\right\}\right) \\
 &\geq \sum_{j=1}^n \frac{\mathbb{P}(E_j)}{2} \\
 &\geq \frac{1}{2} \mathbb{P}\left(\max_{j \in [n]} \sum_{i=1}^j w(Z_i)\xi_i \geq t\right),
 \end{aligned}$$

where the first equality holds because for $i \neq j$, $E_i \cap E_j = \emptyset$, the first inequality follows from (12) and the last inequality comes from the union bound.

Step 2: For any differentiable f and random variable X ,

$$\begin{aligned}
 \mathbb{E}[f(X)\mathbb{1}_{\{X \geq 0\}}] &= \mathbb{E}\left[\left(f(0) + \int_0^X f'(t)dt\right)\mathbb{1}_{\{X \geq 0\}}\right] \\
 &= f(0)\mathbb{E}[\mathbb{1}_{\{X \geq 0\}}] + \mathbb{E}\left[\int_0^\infty f'(t)\mathbb{1}_{\{X \geq t\}}dt\right] \\
 &= f(0)\mathbb{P}(X \geq 0) + \int_0^\infty f'(t)\mathbb{P}(X \geq t)dt,
 \end{aligned} \tag{14}$$

where the last equality follows from Fubini's theorem. Putting it altogether, we have

$$\begin{aligned}
 &\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp\left(\max_{j \in [n]} \frac{\lambda}{n} \sum_{i=1}^j w(Z_i)\xi_i\right)\right] \\
 &= 1 + \lambda \int_0^\infty \exp(\lambda t) \mathbb{P}\left(\max_{j \in [n]} \frac{1}{n} \sum_{i=1}^j w(Z_i)\xi_i \geq t\right) dt \\
 &\leq 1 + 2\lambda \int_0^\infty \exp(\lambda t) \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n w(Z_i)\xi_i \geq t\right) dt \\
 &= 1 + 2\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\left(\exp\left(\frac{\lambda}{n} \sum_{i=1}^n w(Z_i)\xi_i\right) - 1\right) \mathbb{1}_{\{\sum_{i=1}^n w(Z_i)\xi_i \geq 0\}} \right] \\
 &= 1 + 2\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp\left(\frac{\lambda}{n} \sum_{i=1}^n w(Z_i)\xi_i\right) \right] - 2\mathbb{P}\left(\sum_{i=1}^n w(Z_i)\xi_i \geq 0\right) \\
 &\leq 2\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp\left(\frac{\lambda}{n} \sum_{i=1}^n w(Z_i)\xi_i\right) \right],
 \end{aligned}$$

where the first inequality follows from (11), the second equality uses (14) with $f(t) = e^{\lambda t}$ and $X = \max_{j \in [n]} \frac{1}{n} \sum_{i=1}^j w(Z_i)\xi_i$, and the last inequality follows from (13). Finally, if $\frac{1}{n} \sum_{i=1}^n w(Z_i)\xi_i$ is mean zero $\frac{\gamma^2}{n}$ -subGaussian, then using the definition of a subGaussian random variable, we obtain

$$2\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp\left(\frac{\lambda}{n} \sum_{i=1}^n w(Z_i)\xi_i\right) \right] \leq 2 \exp\left(\frac{\lambda^2 \gamma^2}{2n}\right).$$

□

B.2. Proof of Theorem 4.1

Theorem 4.1. *Given a hypothesis class \mathbb{F} , any loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, and n samples $\{Z_i\}_{i=1}^n$, we have that with probability at least $1 - \delta$,*

$$e_n(\mathbb{F}, \ell) \leq 2\mathcal{R}(n, \mathbb{F}) + \sqrt{\frac{\log(\frac{1}{\delta})}{2n}}.$$

Proof. We first analyze the sensitivity of the sup-norm of the CDF estimator over \mathbb{F} and $\mathbb{G}(1)$. For a given two sets $\{x_i, y_i\}_{i=1}^n$ and $\{x'_i, y'_i\}_{i=1}^n$, which just differ in j 'th entry, let

$$\widehat{F}_1(r; f) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(y_i, f(x_i)) \leq r\}} \text{ and } \widehat{F}_2(r; f) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(y'_i, f(x'_i)) \leq r\}}$$

Then, if $\sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}_1(r; f) - F(r; f) \right| \geq \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}_2(r; f) - F(r; f) \right|$, we have

$$\begin{aligned} & \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}_1(r; f) - F(r; f) \right| - \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}_2(r; f) - F(r; f) \right| \\ &= \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(X_i, y_i, f(x_i)) \leq r\}} - F(r; f) \right| - \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(x_i, y'_i, f(x'_i)) \leq r\}} - F(r; f) \right| \\ &= \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(x'_i, y'_i, f(x'_i)) \leq r\}} + \frac{1}{n} \left(\mathbb{1}_{\{\ell(x_j, y_j, f(x_j)) \leq r\}} - \mathbb{1}_{\{\ell(x'_j, y'_j, f(x'_j)) \leq r\}} \right) - F(r; f) \right| \\ & \quad - \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(x'_i, y'_i, f(x'_i)) \leq r\}} - F(r; f) \right| \\ &= \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \frac{1}{n} \left| \mathbb{1}_{\{\ell(x_j, y_j, f(x_j)) \leq r\}} - \mathbb{1}_{\{\ell(x'_j, y'_j, f(x'_j)) \leq r\}} \right| \leq \frac{1}{n}. \end{aligned}$$

$$\begin{aligned} & \left| \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}_1(r; f) - F(r; f) \right| - \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}_2(r; f) - F(r; f) \right| \right| \\ & \leq \sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \frac{1}{n} \left| \mathbb{1}_{\{\ell(x_j, y_j, f(x_j)) \leq r\}} - \mathbb{1}_{\{\ell(x'_j, y'_j, f(x'_j)) \leq r\}} \right| = \frac{1}{n} \end{aligned}$$

This bound holds no matter what j and what data set we choose. Using bounded difference inequality, i.e., McDiarmid's inequality (Boucheron et al., 2013), we have,

$$\mathbb{P} \left(\sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}(r; f) - F(r; f) \right| - \mathbb{E}_{\mathbb{P}} \left[\sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}(r; f) - F(r; f) \right| \right] \leq \sqrt{\frac{\log(\frac{1}{\delta})}{2n}} \right) \quad (15)$$

with probability at least $1 - \delta$. Using a ghost sample set $\{X'_i, Y'_i\}_{i=1}^n$ we have

$$\begin{aligned} & \mathbb{E}_{\mathbb{P}} \left[\sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \widehat{F}(r; f) - F(r; f) \right| \right] \\ &= \mathbb{E}_{\mathbb{P}} \left[\sup_{f \in \mathbb{F}} \sup_{r \in \mathbb{R}} \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(X_i, Y_i, f(X_i)) \leq r\}} - \mathbb{E}_{\mathbb{P}} \left[\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\ell(X'_i, Y'_i, f(X'_i)) \leq r\}} \right] \right| \right] \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{\mathbb{P}} \left[\sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \frac{1}{n} \sum_{i=1}^n g(\ell(X_i, Y_i, f(X_i))) - \mathbb{E}_{\mathbb{P}} \left[\frac{1}{n} \sum_{i=1}^n g(\ell(X'_i, Y'_i, f(X'_i))) \right] \right| \right] \\
 &= \mathbb{E}_{\mathbb{P}} \left[\sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \frac{1}{n} \mathbb{E}_{\mathbb{P}} \left[\sum_{i=1}^n g(\ell(X_i, Y_i, f(X_i))) - \sum_{i=1}^n g(\ell(X'_i, Y'_i, f(X'_i))) \right] \middle| \sigma(\{X_i, Y_i\}_{i=1}^n) \right| \right] \\
 &\leq \mathbb{E}_{\mathbb{P}} \left[\sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \frac{1}{n} \sum_{i=1}^n (g(\ell(X_i, Y_i, f(X_i))) - g(\ell(X'_i, Y'_i, f(X'_i)))) \right| \right] \\
 &\leq \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \frac{1}{n} \sum_{i=1}^n \xi_i (g(\ell(X_i, Y_i, f(X_i))) - g(\ell(X'_i, Y'_i, f(X'_i)))) \right| \right] \\
 &\leq 2 \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \frac{1}{n} \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| \right] = 2\mathcal{R}(n, \mathbb{F}). \tag{16}
 \end{aligned}$$

Putting (15) and (16) together concludes the proof. \square

B.3. Permutation Complexity

We note that this proof, along with many other proofs for Section 4, is based on the machinery and techniques in Massart's finite class Lemma (Massart, 2000) and DKW inequality in (Devroye et al., 2013).

Theorem 4.2. *For any hypothesis class \mathbb{F} and loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, we have that*

$$\mathcal{R}(n, \mathbb{F}) \leq \sqrt{\frac{\log(4\mathcal{N}_{\Pi}(n, \mathbb{F}, \mathbb{P}))}{2n}}.$$

Proof. For a positive λ , we have,

$$\begin{aligned}
 \exp(\lambda \mathcal{R}(n, \mathbb{F})) &= \exp \left(\frac{\lambda}{n} \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| \right] \right) \\
 &\leq \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| \right) \right] \\
 &\leq \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sup_{f \in \mathbb{F}} \sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| \right) \middle| \sigma(\{X_i, Y_i\}_{i=1}^n) \right] \right] \tag{17}
 \end{aligned}$$

For any $f \in \mathbb{F}$, let π_f denote a permutation such that $\ell(X_{\pi_f(i)}, Y_{\pi_f(i)}, f(X_{\pi_f(i)})) \leq \ell(X_{\pi_f(j)}, Y_{\pi_f(j)}, f(X_{\pi_f(j)}))$ for any $i, j \in [n]$ and $i \leq j$. Therefore we have,

$$\sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_i g(\ell(X_i, Y_i, f(X_i))) \right| = \sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_{\pi_f(i)} g(\ell(X_{\pi_f(i)}, Y_{\pi_f(i)}, f(X_{\pi_f(i)}))) \right|$$

Consider a function $g(r') = \mathbb{1}_{\{r' \leq r\}}$. For such a function, $\sum_{i=1}^n \xi_{\pi_f(i)} g(\ell(X_{\pi_f(i)}, Y_{\pi_f(i)}, f(X_{\pi_f(i)})))$ is equal to,

- 0 if $r < \min_i \{\ell(X_{\pi_f(i)}, Y_{\pi_f(i)}, f(X_{\pi_f(i)}))\}^n$,
- $\sum_i^j \xi_{\pi_f(i)}$ when $\ell(X_{\pi_f(j)}, Y_{\pi_f(j)}, f(X_{\pi_f(j)})) \leq r < \ell(X_{\pi_f(j+1)}, Y_{\pi_f(j+1)}, f(X_{\pi_f(j+1)}))$ for a $j \in \{1, \dots, n-1\}$,
- $\sum_i^n \xi_{\pi_f(i)}$ otherwise.

Using this property, we have,

$$\sup_{g \in \mathbb{G}(1)} \left| \sum_{i=1}^n \xi_{\pi_f(i)} g(\ell(X_{\pi_f(i)}, Y_{\pi_f(i)}, f(X_{\pi_f(i)}))) \right| = \max_j \left| \sum_i^j \xi_{\pi_f(i)} \right| \tag{18}$$

Using this equality, we can further extend the Eq. 17,

$$\begin{aligned}
 \exp(\lambda \mathcal{R}(n, \mathbb{F})) &\leq \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \sup_{f \in \mathbb{F}} \max_j \left| \sum_i^j \xi_{\pi_f(i)} \right| \right) \middle| \sigma(\{X_i, Y_i\}_{i=1}^n) \right] \right] \\
 &\leq \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\mathcal{N}_{\Pi}(\mathbb{F}_{\ell}, \{X_i, Y_i\}_{i=1}^n) \exp \left(\frac{\lambda}{n} \max_j \left| \sum_i^j \xi_i \right| \right) \middle| \sigma(\{X_i, Y_i\}_{i=1}^n) \right] \right] \\
 &\leq \mathcal{N}_{\Pi}(n, \mathbb{F}_{\ell}, \mathbb{P}) \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\exp \left(\frac{\lambda}{n} \max_j \left| \sum_i^j \xi_i \right| \right) \right], \tag{19}
 \end{aligned}$$

where the second inequality follows from the fact that effectively, there are at most $\mathcal{N}_{\Pi}(\mathbb{F}_{\ell}, \{X_i, Y_i\}_{i=1}^n)$ number of π_f 's. Using the same derivation for (10) (Lemma B.1), we have

$$\exp(\lambda \mathcal{R}(n, \mathbb{F})) \leq 2 \mathcal{N}_{\Pi}(n, \mathbb{F}_{\ell}, \mathbb{P}) \mathbb{E}_{\mathbb{P}, \mathfrak{R}} \left[\max_j \left(\exp \left(\frac{\lambda}{n} \sum_i^j \xi_i \right) \mathbb{1}_{\{\sum_i^j \xi_i \geq 0\}} \right) \right].$$

By Lemma B.2, we have

$$\exp(\lambda \mathcal{R}(n, \mathbb{F})) \leq 4 \mathcal{N}_{\Pi}(n, \mathbb{F}_{\ell}, \mathbb{P}) \exp \left(\frac{\lambda^2}{2n} \right)$$

Now, taking the log from both sides, and dividing by λ , we have $\mathcal{R}(n, \mathbb{F}) \leq \frac{\log(4 \mathcal{N}_{\Pi}(n, \mathbb{F}_{\ell}, \mathbb{P}))}{\lambda} + \frac{\lambda}{2n}$. Choosing $\lambda = \sqrt{2n \log(2 \mathcal{N}_{\Pi}(n, \mathbb{F}_{\ell}, \mathbb{P}))}$, we obtain the final result

$$\mathcal{R}(n, \mathbb{F}) \leq \sqrt{\frac{\log(4 \mathcal{N}_{\Pi}(n, \mathbb{F}_{\ell}, \mathbb{P}))}{2n}}.$$

□

Corollary 4.1. For a finite hypothesis class \mathbb{F} ,

$$\mathcal{R}(n, \mathbb{F}) \leq \sqrt{\frac{\log(4|\mathbb{F}|)}{2n}}.$$

Proof. The result follows since when $|\mathbb{F}| < \infty$, $\mathcal{N}_{\Pi}(n, \mathbb{F}_{\ell}, \mathbb{P}) \leq |\mathbb{F}|$, i.e., we need at most one permutation function to sort the losses for each $f \in \mathbb{F}$.

□

C. Proof of Results in Section 5

Theorem 5.1. For a hypothesis class \mathbb{F} , a bounded loss function ℓ , and $\delta \in (0, 1]$, if $\mathbb{P}(e_n(\mathbb{F}, \ell) \leq \epsilon) \geq 1 - \delta$, then with probability $1 - \delta$, for all $\rho \in \mathbb{T}$, we have

$$\sup_{f \in \mathbb{F}} |\rho(F(\cdot; f)) - \rho(\widehat{F}(\cdot; f))| \leq L(\rho, 1, \mathcal{L}_\infty)\epsilon,$$

where \mathbb{T} is the set of $L(\cdot, 1, \mathcal{L}_\infty)$ Hölder risk functionals on the space of bounded random variables.

Proof. For all $\rho \in \mathbb{T}$, since ρ is $L(\rho, 1, \mathcal{L}_\infty)$ Hölder, we have that $|\rho(F(\cdot; f)) - \rho(\widehat{F}(\cdot; f))| \leq L(\rho, 1, \mathcal{L}_\infty)\|F - \widehat{F}\|_\infty$. The desired result then follows. \square

The error of risk assessment can be bounded using distances other than the sup-norm, such as the Wasserstein distance. For two random variables U and U' with bounded support $[0, D]$, the dual form of the Wasserstein distance $W_1(F_U, F_{U'})$ is given by (Vallender, 1974),

$$W_1(F_U, F_{U'}) = \int_0^D |F_U(t) - F_{U'}(t)| dt \leq D\|F_U - F_{U'}\|_\infty.$$

This inequality suggests the following corollary.

Corollary C.1. Under the setting of Theorem 5.1 where the loss has support $[0, D]$, with probability $1 - \delta$, for all $\rho \in \mathbb{T}$, we have

$$\sup_{f \in \mathbb{F}} |\rho(F(\cdot; f)) - \rho(\widehat{F}(\cdot; f))| \leq L(\rho, p, W_1)D^p \epsilon^p,$$

where \mathbb{T} is the set of $L(\cdot, p, W_1)$ Hölder risk functionals on the space of bounded random variables.

Proof. For all $\rho \in \mathbb{T}$, since ρ is $L(\rho, p, W_1)$ Hölder, we have that $|\rho(F(\cdot; f)) - \rho(\widehat{F}(\cdot; f))| \leq L(\rho, p, W_1)W_1(F, \widehat{F})^p \leq L(\rho, p, W_1)D^p\|F - \widehat{F}\|_\infty^p$. The desired result then follows. \square

Corollary 5.1. For a hypothesis class \mathbb{F} , a bounded loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow [0, D]$, and $\delta \in (0, 1]$, if $\mathbb{P}(e_n(\mathbb{F}, \ell) \leq \epsilon) \geq 1 - \delta$, then with probability $1 - \delta$, we have

$$\sup_{f \in \mathbb{F}} |\rho(F(\cdot; f)) - \rho(\widehat{F}(\cdot; f))| \leq L\epsilon,$$

where ρ is a distortion risk with $\frac{L}{D}$ -Lipschitz distortion function.

Proof. As is shown in Huang et al. (2021, Lemma 4.1), ρ is L -Lipschitz. Thus, directly applying Theorem 5.1 concludes the proof. \square

D. Proofs for results in Section 6

D.1. Proof of Lemma 6.1

Before proving Lemma 6.1, we first present two auxiliary lemmas.

Lemma D.1. *For any continuous function $g, h : \mathbb{R}^d \rightarrow \mathbb{R}$, $\min(g, h)$ and $\max(g, h)$ are continuous. Similarly, for any Lipschitz continuous function $g, h : \mathbb{R}^d \rightarrow \mathbb{R}$, $\min(g, h)$ and $\max(g, h)$ are Lipschitz continuous.*

Proof. Denote $f_{\min} = \min(g, h)$ and $f_{\max} = \max(g, h)$.

Continuity: We first consider the case when both g and h are continuous. Define $\bar{\Theta} = \{\theta \in \mathbb{R}^d : g(\theta) = h(\theta)\}$. For $\theta \notin \bar{\Theta}$, f_{\min} and f_{\max} are continuous since g, h are continuous. Consider $\theta \in \bar{\Theta}$. For every $\epsilon > 0$, there exists $\delta_g, \delta_h > 0$ such that for all $\theta' \in \mathbb{R}^d$, $\|\theta - \theta'\| \leq \epsilon \implies |g(\theta) - g(\theta')| \leq \delta_g, |h(\theta) - h(\theta')| \leq \delta_h$. In addition, $f_{\min}(\theta) = g(\theta) = h(\theta)$, $f_{\max}(\theta) = g(\theta) = h(\theta)$ and $f_{\min}(\theta')$, $f_{\max}(\theta')$ can be either $g(\theta')$ or $h(\theta')$. Combining both facts gives us that $|f_{\min}(\theta) - f_{\min}(\theta')| \leq \max(\delta_g, \delta_h)$ and $|f_{\max}(\theta) - f_{\max}(\theta')| \leq \max(\delta_g, \delta_h)$.

Lipschitz Continuity: We next work with the case where both g and h are Lipschitz continuous. When $|f_{\max}(\theta) - f_{\max}(\theta')| = |g(\theta) - h(\theta')|$, we have the following two cases:

1. $g(\theta) > h(\theta')$: $|g(\theta) - h(\theta')| = g(\theta) - h(\theta') \leq g(\theta) - g(\theta')$, since $f_{\max}(\theta') = h(\theta')$.
2. $g(\theta) \leq h(\theta')$: $|g(\theta) - h(\theta')| = h(\theta') - g(\theta) \leq h(\theta') - h(\theta)$, since $f_{\max}(\theta) = g(\theta)$.

Since both h and g are Lipschitz continuous, we obtain that

$$|f_{\max}(\theta) - f_{\max}(\theta')| \leq \max_{f \in \{g, h\}} |f(\theta) - f(\theta')| \lesssim \|\theta - \theta'\|,$$

showing that f_{\max} is Lipschitz continuous. The proof completes with the fact that $\min(f, g) = -\max(-f, -g)$. \square

Lemma D.2. *If $\{\ell_\theta(x_j, y_j)\}_{j=1}^n$ are Lipschitz continuous in θ , then for all $i \in [n]$, $\ell_\theta(\pi_\theta(i))$, i.e., the i -th smallest loss evaluated using data points $\{x_j, y_j\}_{j=1}^n$, is Lipschitz continuous in θ .*

Proof. The key observation is that the i -th smallest loss can be defined as $\ell_\theta(\pi_\theta(i)) = \min\{\max\{\ell_\theta(j) : j \in J\} : J \subseteq [n], |J| = i\}$. Since each $\ell_\theta(j)$ is Lipschitz continuous in θ and that for $j' \in J$, $\max\{\ell_\theta(j) : j \in J\} = \max\{\ell_\theta(j'), \max\{\ell_\theta(j) : j \in J \setminus \{j'\}\}\}$, by Lemma D.1, we have $\max\{\ell_\theta(j) : j \in J\}$ to be Lipschitz continuous in θ . Similarly, since $\max\{\ell_\theta(j) : j \in J\}$ is Lipschitz continuous in θ , we have $\min\{\max\{\ell_\theta(j) : j \in J\} : J \subseteq [n], |J| = i\}$ to be Lipschitz continuous. \square

Lemma 6.1. *If $\{\ell_\theta(z_i)\}_{i=1}^n$ are Lipschitz continuous in $\theta \in \Theta$, then for all $i \in [n]$, $\ell_\theta(\pi_\theta(i))$, i.e., the i -th smallest loss, is Lipschitz continuous in θ and $\rho(\widehat{F}_\theta)$ is differentiable in θ almost everywhere.*

Proof. Using Lemma D.2, we obtain that $\ell_\theta(\pi_\theta(i))$ is Lipschitz in $\theta \in \Theta$. Following from a classical result of Rademacher (Rockafellar & Wets, 2009, Theorem 9.60), i.e., a locally Lipschitz function is differentiable almost everywhere, we have that $\rho(\widehat{F}_\theta)$ is differentiable almost everywhere. \square

D.2. Local Convergence

The proofs for Corollary 6.1 are standard (Bottou et al., 2018), which we provide for completeness.

Corollary 6.1. *If $\{\ell_\theta(z_i)\}_{i=1}^n$ are Lipschitz continuous and $\rho(\widehat{F}_\theta)$ is β -smooth in θ , then the following holds almost surely when the learning rate in (7) is $\eta = \frac{1}{\beta\sqrt{T}}$:*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\|\nabla_\theta \rho(\widehat{F}_{\theta_t})\|^2 \right] \leq \frac{2\beta}{\sqrt{T}} \left(\rho(\widehat{F}_{\theta_1}) - \rho(\widehat{F}_{\theta_*}) + \frac{1}{2\beta} \right).$$

Proof. For notation simplicity, we use $h(\theta)$ to denote $\rho(\widehat{F}_\theta)$ and g_t to denote $\nabla_\theta \rho(\widehat{F}_\theta)$ when $\theta = \theta_t$. Since $h(\theta)$ is differentiable almost everywhere, following (7), the sequence $\{h(\theta_t)\}_{t=1}^T$ will be differentiable almost surely. Since h is

β -smooth, we have that

$$h(\theta_{t+1}) - h(\theta_t) \leq \nabla_{\theta} h(\theta_t)^{\top} (\theta_{t+1} - \theta_t) + \frac{\beta}{2} \|\theta_{t+1} - \theta_t\|^2.$$

We denote the filtration for the stochastic process $\{\theta_t\}_{t=1}^T$ to be $\{\mathcal{F}_t\}_{t=1}^T$. Extending the above inequality, we have

$$h(\theta_{t+1}) - h(\theta_t) \leq \nabla_{\theta} h(\theta_t)^{\top} (-\eta(g_t + w_t)) + \frac{\beta}{2} \|\eta(g_t + w_t)\|^2,$$

which suggests that

$$\mathbb{E}[h(\theta_{t+1}) - h(\theta_t)] \leq -\eta \nabla_{\theta} h(\theta_t)^{\top} g_t + \eta^2 \frac{\beta}{2} (\|g_t\|^2 + d\sigma_w^2),$$

where $\sigma_w^2 = 1/d$ is the variance of w_t . Therefore, for the conditional expectation, we have

$$\begin{aligned} \mathbb{E} \left[h(\theta_{t+1}) - h(\theta_t) \middle| \mathcal{F}_t \right] &\leq \mathbb{E} \left[-\eta \nabla_{\theta} h(\theta_t)^{\top} g_t + \eta^2 \frac{\beta}{2} (\|g_t\|^2 + 1) \middle| \mathcal{F}_t \right] \\ &= -(\eta - \eta^2 \frac{\beta}{2}) \mathbb{E} \left[\|\nabla_{\theta} h(\theta_t)\|^2 \middle| \mathcal{F}_t \right] + \eta^2 \frac{\beta}{2}. \end{aligned}$$

Using the telescoping sum and law of total expectation, we obtain

$$\mathbb{E} \left[h(\theta_T) - h(\theta_1) \middle| \mathcal{F}_1 \right] = \mathbb{E} \left[\sum_{t=1}^T h(\theta_t) - h(\theta_{t-1}) \middle| \mathcal{F}_1 \right] \leq (\eta - \eta^2 \frac{\beta}{2}) \mathbb{E} \left[\sum_{t=1}^T \|\nabla_{\theta} h(\theta_t)\|^2 \middle| \mathcal{F}_1 \right] + T\eta^2 \frac{\beta}{2}. \quad (20)$$

Use the fact that $h(\theta_{\star}) \leq h(\theta_T)$, we have $\mathbb{E} [h(\theta_1) - h(\theta_T)] \leq h(\theta_1) - h(\theta_{\star})$, which implies

$$(\eta - \eta^2 \frac{\beta}{2}) \mathbb{E} \left[\sum_{t=1}^T \|\nabla_{\theta} h(\theta_t)\|^2 \right] \leq h(\theta_1) - h(\theta_{\star}) + T\eta^2 \frac{\beta}{2}.$$

Plugging the learning rate $\eta = \frac{1}{\beta\sqrt{T}}$, we have $\eta - \eta^2 \frac{\beta}{2} = \frac{1}{\beta\sqrt{T}} - \frac{1}{2\beta T} \geq \frac{1}{\beta\sqrt{T}} - \frac{1}{2\beta\sqrt{T}} \geq \frac{1}{2\beta\sqrt{T}} > 0$. Therefore we have

$$\frac{1}{2\beta\sqrt{T}} \mathbb{E} \left[\sum_{t=1}^T \|\nabla_{\theta} h(\theta_t)\|^2 \right] \leq h(\theta_1) - h(\theta_{\star}) + \frac{1}{2\beta}.$$

Rearranging the above inequality gives the result. \square

E. Additional Experimental Details

E.1. Risk Assessment on ImageNet Models

Figure 3 shows the CVaR of models presented in Table 1 under different α 's. We note that CVaR_α is the expected value above the top 100α percent losses.




Figure 3. $\text{CVaR}_\alpha(\ell_f(Z))$ across different α 's where ℓ_f is the cross-entropy loss evaluated on the ImageNet validation dataset.

E.2. Empirical Distortion Risk Minimization

Toy Example The data used in this experiment is generated using the `make_blobs` function from `sklearn.datasets` with the following parameters: `n_samples = [1000, 50]`, `centers = [[0.0, 0.0], [1.0, 1.0]]`, `cluster_std = [1.5, 0.5]`, `random_state = 0`, `shuffle = False`.

CIFAR-10 For completeness, we report the average test accuracy for the VGG-16 models obtained through minimizing the empirical risks for expected loss, $\text{CVaR}_{0.05}$, $\text{CVaR}_{0.7}$ and $\text{HRM}_{3,4}$. The models are trained over 150 epochs and the learning rate is chosen to be 0.005. The accuracy is 55.2%, 13.2%, 51.3%, and 53.9% respectively. We note that the goal of this experiment is to illustrate the efficacy of our proposed optimization procedure for minimizing distortion risks instead of arguing for the usage of a particular risk functional for finding a model with the highest test accuracy.