Adversarial Surrogate Risk Bounds for Binary Classification

Natalie Frank *natalief@uw.edu*
Department of Applied Mathematics
University of Washington

Reviewed on OpenReview: *https://openreview.net/forum?id=Bay1cHLk7h*

Abstract

A central concern in classification is the vulnerability of machine learning models to adversarial attacks. Adversarial training is one of the most popular techniques for training robust classifiers, which involves minimizing an adversarial surrogate risk. Recent work has characterized the conditions under which any sequence minimizing the adversarial surrogate risk also minimizes the adversarial classification risk in the binary setting, a property known as *adversarial consistency*. However, these results do not address the rate at which the adversarial classification risk approaches its optimal value along such a sequence. This paper provides surrogate risk bounds that quantify that convergence rate.

1 Introduction

A central concern regarding regarding sophisticated machine learning models is their susceptibility to adversarial attacks. Prior work (Biggio et al., 2013; Szegedy et al., 2013) demonstrated that imperceptible perturbations can degrade the performance of neural nets. As such models are deployed in security-critical applications, including facial recognition (Xu et al., 2022) and medical imaging (Paschali et al., 2018), training robust models remains a key challenge in machine learning.

In the standard classification setting, the *classification risk* is the proportion of incorrectly classified data. Directly minimizing this quantity is a combinatorial optimization problem, so typical machine learning algorithms instead minimize a more tractable *surrogate* risk via gradient-based methods. A surrogate risk is said to be consistent for a given data distribution if every minimizing sequence also minimizes the classification risk for that distribution. Beyond consistency, a central objective is efficiency: minimizing the surrogate risk should translate into a rapid reduction of the classification risk. This rate can be quantified via surrogate risk bounds, which bound the excess classification risk in terms of the excess surrogate risk.

In the standard binary classification setting, consistency and surrogate risk bounds are well-studied topics (Bartlett et al., 2006; Lin, 2004; Steinwart, 2007; Zhang, 2004). A typical approach reduces the problem to a pointwise analysis of the conditional classification and surrogate risks. In contrast, the adversarial setting is less understood. The adversarial classification risk penalizes instances that can be perturbed into the opposite class, while the adversarial surrogate risk computes the worst-case value over an ϵ-ball. The dependence on the value of a function over an ϵ-ball precludes a pointwise decomposition, rendering the classical analysis inapplicable. Frank & Niles-Weed (2024a) characterized the risks that are consistent for all data distributions, and the corresponding losses are referred to as *adversarially consistent*. Unfortunately, no convex loss function can be adversarially consistent for all data distributions (Meunier et al., 2022). On the other hand, Frank (2025) showed that such situations are rather atypical— when the data distribution is absolutely continuous, a surrogate risk is adversarially consistent so long as the adversarial Bayes classifier satisfies a certain notion of uniqueness. While these results characterize consistency, none describe convergence rates.

Our Contributions:

- We prove a linear surrogate risk bound for adversarially consistent losses (Theorem 9).

- When the "distribution of optimal attacks" satisfies a bounded noise condition, we prove a linear surrogate risk bound under mild conditions on the loss (Theorem 9).
- We establish a distribution-dependent surrogate risk bound that applies whenever a loss is adversarially consistent for the data distribution (Theorem 11).

Notably, the last result applies to convex loss functions. By prior consistency results (Frank, 2025; Frank & Niles-Weed, 2024a; Meunier et al., 2022), one cannot hope for distribution independent surrogate bounds for non-adversarially consistent losses. This work presents a framework for surrogate risk bounds that applies to any supremum-based risk under mild conditions. A detailed comparison with prior work is provided in Section 7.

2 Background and Preliminaries

2.1 Surrogate Risks

We study binary classification on $\mathbb{R}^d$ with labels -1 and $+1$, where $\mathbb{P}_0$ and $\mathbb{P}_1$ denote the class-conditional distributions. For a measurable set A, the classification risk is

$$R(A) = \int 1_{A^C}\, d\mathbb{P}_1 + \int 1_A\, d\mathbb{P}_0,$$

with minimum R^* over all Borel sets. Because the indicator function is nondifferentiable, one instead minimizes a *surrogate risk*

$$R_\phi(f) = \int \phi(f)\, d\mathbb{P}_1 + \int \phi(-f)\, d\mathbb{P}_0,$$

with minimum $R_{\phi,*}$ over all Borel functions. The loss ϕ satisfies:

Assumption 1. *ϕ is continuous, non-increasing, and $\lim_{\alpha\to\infty}\phi(\alpha) = 0$.*

Thresholding f at zero yields the classifier $\{f > 0\}$, whose risk is

$$R(f) = R(\{f > 0\}) = \int 1_{f\le 0}\, d\mathbb{P}_1 + \int 1_{f>0}\, d\mathbb{P}_0.$$

It remains to verify that minimizing the surrogate risk R_ϕ will also minimize the classification risk R.

Definition 1. *The loss function ϕ is* consistent for the distribution $\mathbb{P}_0, \mathbb{P}_1$ *if every minimizing sequence of R_ϕ is also a minimizing sequence of R. The loss function ϕ is* consistent *if it is consistent for all distributions.*

Prior work establishes conditions under which many common loss functions are consistent. For convex ϕ, consistency occurs iff ϕ is differentiable at 0 and $\phi'(0) < 0$ (Bartlett et al., 2006, Theorem 2). Frank & Niles-Weed (2024a, Proposition 3) show that consistency holds if $\inf_\alpha \frac{1}{2}(\phi(\alpha) + \phi(-\alpha)) < \phi(0)$, which is satisfied by losses such as the ρ-margin loss $\phi_\rho(\alpha) = \min(1, \max(1 - \alpha/\rho, 0))$ and the shifted sigmoid loss $\phi_\tau(\alpha) = 1/(1 + \exp(\alpha - \tau))$, $\tau > 0$. However, a convex loss ϕ cannot satisfy this inequality:

$$\frac{1}{2}\left(\phi(\alpha) + \phi(-\alpha)\right) \ge \phi\left(\frac{1}{2}\alpha + \frac{1}{2}\cdot -\alpha\right) = \phi(0). \tag{1}$$

2.2 Surrogate Risk Bounds

In addition to consistency, quantifying convergence rates is a key concern. Specifically, prior work (Bartlett et al., 2006; Zhang, 2004) establishes *surrogate risk bounds* of the form $\Psi(R(f) - R_*) \le R_\phi(f) - R_{\phi,*}$ for some function Ψ, linking excess classification risk to excess surrogate risk. These bounds involve pointwise minima of the *conditional* classification and surrogate risks.

Let $\mathbb{P} = \mathbb{P}_0 + \mathbb{P}_1$ and $\eta(\mathbf{x}) = d\mathbb{P}_1/d\mathbb{P}$. An equivalent formulation of the classification risk is

$$R(f) = \int C(\eta(\mathbf{x}), f(\mathbf{x}))d\mathbb{P}(\mathbf{x}) \tag{2}$$

where $C(\eta, \alpha) = \eta \mathbf{1}_{\alpha \leq 0} + (1-\eta)\mathbf{1}_{\alpha > 0}$, with minimal conditional risk

$$C^*(\eta) = \inf_\alpha C(\eta, \alpha) = \min(\eta, 1-\eta), \tag{3}$$

and thus the minimal classification risk is $R_* = \int C^*(\eta(\mathbf{x}))d\mathbb{P}(\mathbf{x})$. Analogously, the surrogate risk in terms of η and $\mathbb{P}$ is

$$R_\phi(f) = \int C_\phi(\eta(\mathbf{x}), f(\mathbf{x}))d\mathbb{P}, \quad C_\phi(\eta, \alpha) = \eta\phi(\alpha) + (1-\eta)\phi(-\alpha) \tag{4}$$

and the minimal surrogate risk is $R_{\phi,*} = \int C^*_\phi(\eta(\mathbf{x}))d\mathbb{P}(\mathbf{x})$ with the minimal conditional risk $C^*_\phi(\eta)$ defined by

$$C^*_\phi(\eta) = \inf_\alpha C_\phi(\eta, \alpha). \tag{5}$$

Prior work on consistency typically establishes surrogate risk bounds via pointwise analysis of the conditional risks, relating the excess conditional surrogate risk $C_\phi(\eta, \alpha) - C^*_\phi(\eta)$ to the excess conditional classification risk $C(\eta, \alpha) - C^*(\eta)$.

The consistency of ϕ can be fully characterized by the properties of the function $C^*_\phi(\eta)$.

Theorem 1. *A loss ϕ is consistent iff $C^*_\phi(\eta) < \phi(0)$ for all $\eta \neq 1/2$.*

Surprisingly, this criterion has not appeared in prior work. See Appendix A for a proof. In terms of the function C^*_ϕ, Frank & Niles-Weed (2024a, Proposition 3) states that any loss ϕ with $C^*_\phi(1/2) < \phi(0)$ is consistent. The function C^*_ϕ is a key component of surrogate risk bounds from prior work. Specifically, Bartlett et al. (2006) shows:

Theorem 2 (Tewari & Bartlett (2007))**.** *Let ϕ be any loss satisfying Assumption 1 with $C^*_\phi(1/2) = \phi(0)$ and define*

$$\Psi(\theta) = \phi(0) - C^*_\phi\left(\frac{1+\theta}{2}\right).$$

Then

$$\Psi(C(\eta, f) - C^*(\eta)) \leq C_\phi(\eta, f) - C^*_\phi(\eta) \tag{6}$$

and consequently

$$\Psi(R(f) - R_*) \leq R_\phi(f) - R^*_\phi. \tag{7}$$

The inequality (7) is a consequence of (6) and Jensen's inequality. Theorem 1 implies that this bound is non-vacuous iff ϕ is consistent— compare with Theorem 1. Moreover, (6) yields a distribution-dependent linear surrogate bound when η is bounded away from $1/2$. If *Massart's noise condition* (Massart & Nédélec, 2006) holds— namely, there exists a $\alpha \in [0, 1/2]$ for which $|\eta - 1/2| \geq \alpha$ $\mathbb{P}$-a.e., then the distribution admits a linear surrogate bound.

Proposition 1. *Let η, $\mathbb{P}$ be a distribution that satisfies $|\eta - 1/2| \geq \alpha$ $\mathbb{P}$-a.e. with a constant $\alpha \in [0, 1/2]$, and let ϕ be a loss with $\phi(0) > C^*_\phi(1/2 - \alpha)$. Then for all $|\eta - 1/2| \geq \alpha$,*

$$C(\eta, f) - C^*(\eta) \leq \frac{1}{\phi(0) - C^*_\phi(\frac{1}{2} - \alpha)}(C_\phi(\eta, f) - C^*_\phi(\eta)) \tag{8}$$

and consequently

$$R(f) - R_* \leq \frac{1}{\phi(0) - C^*_\phi(\frac{1}{2} - \alpha)}(R_\phi(f) - R_{\phi,*}) \tag{9}$$

See Appendix B for a proof of this result. Observe that Theorem 1 guarantees that the linear constant is finite whenever $\alpha \neq 0$ and ϕ is consistent. This bound is distribution-independent when $\phi(0) > C_\phi^*(1/2)$ with $\alpha = 0$, and will later be generalized to adversarial risks. Although the constant in Proposition 1 is not optimal, further refinement offers no improvement to our adversarial bounds, so we opt to retain the simpler form.

2.3 Adversarial Risks

The adversarial classification risk incurs a penalty of 1 whenever a point $\mathbf{x}$ can be perturbed into the opposite class. This penalty can be expressed in terms of supremums of indicator functions— the adversarial classification risk incurs a penalty of 1 whenever $\sup_{\|\mathbf{x}'-\mathbf{x}\|\leq\epsilon} \mathbf{1}_A(\mathbf{x}') = 1$ or $\sup_{\|\mathbf{x}'-\mathbf{x}\|\leq\epsilon} \mathbf{1}_{A^C}(\mathbf{x}') = 1$. Define

$$S_\epsilon(g)(\mathbf{x}) = \sup_{\|\mathbf{x}-\mathbf{x}'\|\leq\epsilon} g(\mathbf{x}').$$

The adversarial classification and surrogate risks are given respectively by[1]

$$R^\epsilon(A) = \int S_\epsilon(\mathbf{1}_{A^C})d\mathbb{P}_1 + \int S_\epsilon(\mathbf{1}_A)d\mathbb{P}_0, \quad R_\phi^\epsilon(f) = \int S_\epsilon(\phi(f))d\mathbb{P}_1 + \int S_\epsilon(\phi(-f))d\mathbb{P}_0.$$

A minimizer of the adversarial classification risk is called an *adversarial Bayes classifier*. After optimizing the surrogate risk, a classifier is obtained by thresholding the resulting function f at zero. The associated adversarial classification error function f is then

$$R^\epsilon(f) = R^\epsilon(\{f > 0\}) = \int S_\epsilon(\mathbf{1}_{f\leq 0})d\mathbb{P}_1 + \int S_\epsilon(\mathbf{1}_{f>0})d\mathbb{P}_0. \tag{10}$$

Just as in the standard case, one would hope that minimizing the adversarial surrogate risk would minimize the adversarial classification risk.

Definition 2. *The loss ϕ is* adversarially consistent for the distribution $\mathbb{P}_0$, $\mathbb{P}_1$ *if any minimizing sequence of R_ϕ^ϵ is also a minimizing sequence of R^ϵ. We say that ϕ is* adversarially consistent *if it is adversarially consistent for all distributions.*

Theorem 2 of Frank & Niles-Weed (2024a) characterizes the adversarially consistent losses:

Theorem 3 (Frank & Niles-Weed (2024a))**.** *The loss ϕ is adversarially consistent iff $C_\phi^*(1/2) < \phi(0)$.*

Frank & Niles-Weed (2024a, Proposition 3) guarantees that every adversarially consistent loss is also consistent in the standard sense. Unfortunately, (1) shows that no convex loss is adversarially consistent. However, distributions for which consistency fails are atypical: for absolutely continuous $\mathbb{P}$, adversarial consistency holds provided the adversarial Bayes classifier is *unique up to degeneracy*.

Definition 3. *Two adversarial Bayes classifiers A_1, A_2 are* equivalent up to degeneracy *if any set A with $A_1 \cap A_2 \subset A \subset A_1 \cup A_2$ is also an adversarial Bayes classifier. The adversarial Bayes classifier is* unique up to degeneracy *if any two adversarial Bayes classifiers are equivalent up to degeneracy.*

See Figure 1 for an illustration of non-equivalent adversarial Bayes classifiers in a distribution where adversarial consistency fails. Theorem 4 of Frank (2025) relates uniqueness of the adversarial Bayes classifier to the consistency of ϕ.

Theorem 4 (Frank (2025))**.** *Let ϕ be a loss with $C_\phi^*(1/2) = \phi(0)$ and assume that $\mathbb{P}$ is absolutely continuous with respect to Lebesgue measure. Then ϕ is adversarially consistent for the data distribution given by $\mathbb{P}_0$, $\mathbb{P}_1$ iff the adversarial Bayes classifier is unique up to degeneracy.*

Any extension of surrogate risk bounds to the adversarial setting must account for the conditions of Theorems 3 and 4.

[1] In order to define the risks R_ϕ^ϵ and R^ϵ, one must argue that $S_\epsilon(g)$ is measurable. Theorem 1 of Frank & Niles-Weed (2024b) proves that whenever g is Borel, $S_\epsilon(g)$ is always measurable with respect to the completion of any Borel measure.

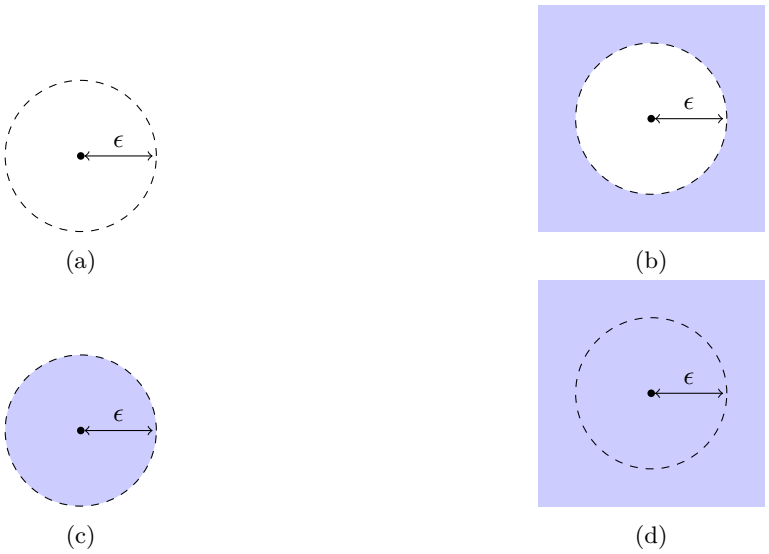

Figure 1: Adversarial Bayes classifiers for the distribution where $\mathbb{P}_0 = \mathbb{P}_1$ are uniform distributions on $\overline{B_\epsilon(\mathbf{0})}$, the counterexample from Meunier et al. (2022). The classifiers in (a) and (b) are equivalent up to degeneracy, as are those in (c) and (d), but the classifiers in (a) and (c) are not. A sequence minimizing R^ϵ_ϕ but not R^ϵ is provided in (33).

2.4 Minimax Theorems

A central tool in analyzing the adversarial consistency of surrogate risks is minimax theorems, which enable a 'pointwise'-style representation of adversarial risks analogous (4). This section reviews the minimax representation for both adversarial classification and surrogate risks, which underlie the bounds in Section 3.

These minimax theorems utilize the ∞-Wasserstein (W_∞) metric from optimal transport. Informally, this metric quantifies the smallest radius ϵ such that the mass of one distribution can be transported to match another without moving any point more than ϵ.

Formally, let $\mathbb{Q}$ and $\mathbb{Q}'$ be finite positive measures with equal total mass. A Borel measure γ on $\mathbb{R}^d \times \mathbb{R}^d$ is a coupling between $\mathbb{Q}$ and $\mathbb{Q}'$ if its first marginal is $\mathbb{Q}$ and its second marginal is $\mathbb{Q}'$, or in other words, $\gamma(A \times \mathbb{R}^d) = \mathbb{Q}(A)$ and $\gamma(\mathbb{R}^d \times A) = \mathbb{Q}'(A)$ for all Borel sets A. Denote the set of couplings between $\mathbb{Q}$ and $\mathbb{Q}'$ by $\Pi(\mathbb{Q}, \mathbb{Q}')$. Then the W_∞ distance is

$$W_\infty(\mathbb{Q}, \mathbb{Q}') = \inf_{\gamma \in \Pi(\mathbb{Q}, \mathbb{Q}')} \operatorname*{ess\,sup}_{(\mathbf{x},\mathbf{y}) \sim \gamma} \|\mathbf{x} - \mathbf{y}\|. \tag{11}$$

Theorem 2.6 of Jylhä (2014) proves that the infimum in (11) is always attained. The ϵ-ball around $\mathbb{Q}$ in the W_∞ metric is $\mathcal{B}^\infty_\epsilon(\mathbb{Q}) = \{\mathbb{Q}' : W_\infty(\mathbb{Q}', \mathbb{Q}) \leq \epsilon\}$.

The next lemma is a standard observation linking adversarial perturbations to W_∞-balls. We include a proof in Appendix C for completeness; it is a known result and not new to this work (see for instance Matthew Staib (2017, Proposition 3.1)).

Lemma 1. *Let g be a Borel function. Let γ be a coupling between the measures $\mathbb{Q}$ and $\mathbb{Q}'$ supported on $\Delta_\epsilon = \{(\mathbf{x}, \mathbf{x}') : \|\mathbf{x} - \mathbf{x}'\| \leq \epsilon\}$. Then $S_\epsilon(g)(\mathbf{x}) \geq g(\mathbf{x}')$ γ-a.e. and consequently*

$$\int S_\epsilon(g) d\mathbb{Q} \geq \sup_{\mathbb{Q}' \in \mathcal{B}^\infty_\epsilon(\mathbb{Q})} \int g d\mathbb{Q}'.$$

Applying Lemma 1 to R^ϵ shows that $\inf_A R^\epsilon(A)$ can be expressed as an inf-sup problem. The minimax theorem of Pydi & Jog (2021) ensures that the order of the inf and sup can be interchanged. Let $C^*(\eta)$ be as defined in (3) and define

$$\bar{R}(\mathbb{P}'_0, \mathbb{P}'_1) = \inf_{A \text{ Borel}} \int \mathbf{1}_{A^C} d\mathbb{P}'_1 + \int \mathbf{1}_A d\mathbb{P}'_0 = \int C^* \left(\frac{d\mathbb{P}'_1}{d\left(\mathbb{P}'_1 + \mathbb{P}'_0\right)} \right) d\left(\mathbb{P}'_0 + \mathbb{P}'_1\right). \tag{12}$$

Theorem 5 (Frank (2025))**.** *Let $\bar{R}$ be as defined in (12). Then*

$$\inf_{A \text{ Borel}} R^\epsilon(A) = \sup_{\substack{\mathbb{P}'_1 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_1) \\ \mathbb{P}'_0 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_0)}} \bar{R}(\mathbb{P}'_0, \mathbb{P}'_1).$$

with equality attained at some Borel A, $\mathbb{P}^*_0 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_0)$, and $\mathbb{P}^*_1 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_1)$.

See Frank & Niles-Weed (2024a, Theorem 1) for a proof of this statement. The maximizers $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ can be interpreted as optimal adversarial attacks (see discussion following Frank & Niles-Weed (2024b, Theorem 7)). Frank (2024, Theorem 3.4) provide a criterion for uniqueness up to degeneracy in terms of dual maximizers.

Theorem 6 (Frank (2025))**.** *The following are equivalent:*

A) The adversarial Bayes classifier is unique up to degeneracy

B) There are maximizers $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ of $\bar{R}$ for which $\mathbb{P}^*(\eta^* = 1/2) = 0$, where $\mathbb{P}^* = \mathbb{P}^*_0 + \mathbb{P}^*_1$ and $\eta^* = d\mathbb{P}^*_1/d\mathbb{P}^*$

Thus, uniqueness corresponds to the situation in which the set where both classes are equally probable has measure zero under some optimal adversarial attack.

The analogous dual problem to R^ϵ_ϕ uses $C^*_\phi(\eta)$ from (5)

$$\bar{R}_\phi(\mathbb{P}'_0, \mathbb{P}'_1) = \inf_{f \text{ Borel}} \int \phi(f) d\mathbb{P}'_1 + \int \phi(-f) d\mathbb{P}'_0 = \int C^*_\phi \left(\frac{d\mathbb{P}'_1}{d\left(\mathbb{P}'_1 + \mathbb{P}'_0\right)} \right) d\left(\mathbb{P}'_0 + \mathbb{P}'_1\right) \tag{13}$$

and the analogous minimax theorem states (Frank & Niles-Weed, 2024b, Theorem 6):

Theorem 7 (Frank & Niles-Weed (2024b))**.** *Let $\bar{R}_\phi$ be defined as in (13). Then*

$$\inf_{\substack{f \text{ Borel}, \\ \mathbb{R}\text{-valued}}} R^\epsilon_\phi(f) = \sup_{\substack{\mathbb{P}'_1 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_1) \\ \mathbb{P}'_0 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_0)}} \bar{R}_\phi(\mathbb{P}'_0, \mathbb{P}'_1).$$

with maximizers $\mathbb{P}^*_0 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_0)$, $\mathbb{P}^*_1 \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_1)$ attained.

Finally, optimal attacks for the surrogate problem are also optimal for the classification problem:

Theorem 8. *Consider maximizing the dual objectives $\bar{R}_\phi$ and $\bar{R}$ over $\mathcal{B}^\infty_\epsilon(\mathbb{P}_0) \times \mathcal{B}^\infty_\epsilon(\mathbb{P}_1)$.*

1) If ϕ is consistent, then any maximizer $(\mathbb{P}^*_0, \mathbb{P}^*_1)$ of $\bar{R}_\phi$ over $\mathcal{B}^\infty_\epsilon(\mathbb{P}_0) \times \mathcal{B}^\infty_\epsilon(\mathbb{P}_1)$ also maximizes $\bar{R}$.

2) [Frank (2025)] If the adversarial Bayes classifier is unique up to degeneracy, then there exists a maximizer $(\mathbb{P}^*_0, \mathbb{P}^*_1)$ of $\bar{R}_\phi$ with $\mathbb{P}^*(\eta^* = 1/2) = 0$, where $\mathbb{P}^* = \mathbb{P}^*_0 + \mathbb{P}^*_1$ and $\eta^* = d\mathbb{P}^*_1/d\mathbb{P}^*$.

See Appendix D for a proof of Item 1), Item 2) is shown in Theorems 5 and 7 of Frank (2025). This minimax machinery links the adversarial Bayes classifier, optimal attacks, and surrogate risks, establishing the dual formulations used in Section 3 to derive adversarial surrogate risk bounds.

3 Main Results

Prior work has characterized when a loss ϕ is adversarially consistent with respect to a distribution $\mathbb{P}_0$, $\mathbb{P}_1$. Theorem 3 shows that a distribution-independent surrogate risk bound is possible only when $C^*_\phi(1/2) < \phi(0)$. When $C^*_\phi(1/2) = \phi(0)$, Theorem 4 indicates that any such bound must depend on the marginal distribution of η^* under $\mathbb{P}^*$, and moreover, is possible only if $\mathbb{P}^*(\eta^* = 1/2) = 0$.

Compare these statements with Proposition 1: Theorems 3, 4 and 8 together imply if either $C^*_\phi(1/2) < \phi(0)$ or if there exist some maximizers of $\bar{R}_\phi$ that satisfy Massart's noise condition, then ϕ is adversarially consistent for $\mathbb{P}_0$, $\mathbb{P}_1$. Alternatively, due to Theorem 8, one can equivalently assume that there are maximizers of $\bar{R}_\phi$ satisfying Massart's noise condition. Our first result extends Proposition 1 to the adversarial scenario, replacing $\mathbb{P}_0$, $\mathbb{P}_1$ with the distribution of optimal adversarial attacks.

Theorem 9. *Let ϕ be consistent and let $\mathbb{P}_0$, $\mathbb{P}_1$ be a distribution for which there are maximizers $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ of the dual problem $\bar{R}_\phi$ that satisfy $|\eta^* - 1/2| \geq \alpha$ $\mathbb{P}^*$-a.e. for some constant $\alpha \in [0, 1/2]$ with $C^*_\phi(1/2 - \alpha) < \phi(0)$, where $\mathbb{P}^* = \mathbb{P}^*_0 + \mathbb{P}^*_1$, $\eta^* = d\mathbb{P}^*_1/d\mathbb{P}^*$. Then*

$$R^\epsilon(f) - R^\epsilon_* \leq \frac{1}{\phi(0) - C^*_\phi(1/2 - \alpha)} \left(R^\epsilon_\phi(f) - R^\epsilon_{\phi,*}\right) \tag{14}$$

When $C^*_\phi(1/2) < \phi(0)$, setting $\alpha = 0$ in Theorem 9 yields a distribution-independent bound. As noted earlier, two losses satisfying this condition are the ρ-margin loss and the shifted sigmoid loss. Likewise, Theorem 1 ensures that the linear constant is finite whenever $\alpha \neq 0$ and ϕ is consistent.

The constant appearing in Theorem 9 is nearly optimal: Section 4.3 shows that it can be improved by at most a factor of two, and this gap is attained by a known counterexample to consistency. Thus, the result provides a sharp characterization of how tightly the adversarial classification risk can be controlled by the surrogate risk across all consistent convex losses.

Furthermore, the theorem parallels the classical realizable-case guarantee from the non-adversarial setting. If the optimal adversarial risk satisfies $R^\epsilon_* = 0$, then Massart's noise condition holds with $\alpha = 1/2$ (see Lemma 2). In this regime, Theorem 9 yields a linear relationship between adversarial classification and surrogate risks that is directly analogous to the non-adversarial bound in Proposition 1. Zero adversarial risk occurs whenever the supports of $\mathbb{P}_0$ and $\mathbb{P}_1$ are separated by at least 2ϵ (Example 1 and Figure 3a).

Theorem 9 states that if some distribution of *optimal adversarial attacks* satisfies Massart's noise condition, then the excess adversarial surrogate risk is at worst a linear upper bound on the excess adversarial classification risk. However, if $C^*_\phi(1/2) = \phi(0)$, the bound's constant diverges as $\alpha \to 0$, reflecting the failure of adversarial consistency when the adversarial Bayes classifier is not unique up to degeneracy. For $\alpha \neq 1/2$, understanding the assumptions on $(\mathbb{P}_0, \mathbb{P}_1)$ which ensure Massart's condition for the distribution of adversarial attacks $(\mathbb{P}^*_0, \mathbb{P}^*_1)$ remains an open problem. Example 4.6 of Frank (2024) exhibits a distribution that satisfies Massart's noise condition and yet the adversarial Bayes classifier is not unique up to degeneracy. Thus Massart's noise condition for $\mathbb{P}_0, \mathbb{P}_1$ does not guarantee Massart's noise condition for $\mathbb{P}^*_0$, $\mathbb{P}^*_1$. See Example 2 and Figure 3b for an example where Theorem 9 applies with $\alpha > 0$.

One approach to relaxing the distributional restriction is to apply (14) only on the portion of the distribution where $|\eta^* - 1/2| \geq \alpha$ and then add back in the risk on $|\eta^* - 1/2| < \alpha$.

Theorem 10. *Assume that there exist maximizers $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ of $\bar{R}_\phi$ that are induced by transport maps from $\mathbb{P}_0$, $\mathbb{P}_1$, and define $\mathbb{P}^* = \mathbb{P}^*_1 + \mathbb{P}^*_0$, $\eta^* = d\mathbb{P}^*_1/d\mathbb{P}^*$. Let $0 \leq \alpha$, then*

$$R^\epsilon(f) - R^\epsilon_* \leq \frac{1}{\phi(0) - C^*_\phi(1/2 - \alpha)} \left(R^\epsilon_\phi(f) - R^\epsilon_{\phi,*}\right) + \left(\frac{1}{2} + \alpha\right) \mathbb{P}^*(|\eta^* - 1/2| < \alpha)$$

Since this holds for all α, the right-hand side can be minimized over α. Prior work from optimal transport theory verifies the assumption on $\mathbb{P}^*$ under mild conditions: Theorem 3.5 of Jylhä (2014) states that whenever $\mathbb{P}_0, \mathbb{P}_1$ are absolutely continuous with respect to Lebesgue measure and the norm $\|\cdot\|$ is strictly convex, the measures $\mathbb{P}^*_0, \mathbb{P}^*_1$ are induced by a transport map. It is unclear whether this holds for common datasets such as CIFAR-10 or MNIST.

Finally, an alternative approach to removing the distributional restriction is to average bounds of the form (14) over all values of η^* yielding a distribution-dependent surrogate bound, valid whenever the adversarial Bayes classifier is unique up to degeneracy. For a given function f, let the *concave envelope* of f be the smallest concave function larger than f:

$$\text{conc}(f) = \inf\{g :\geq f \text{ on } \text{dom}(f), g \text{ concave and upper semi-continuous}\} \tag{15}$$

(a) (b) (c) (d)

Figure 2: Ambiguous images in the MNIST and CIFAR10 datasets. (a) lies between a '5' and a '3', while (b) is difficult to classify at all, despite being labeled as a '5'. In CIFAR10, image (c) is ambiguous between a ship and an airplane, and image (d) is similarly hard to identify.

Theorem 11. *Assume $\mathbb{P}_0(\mathbb{R}^d)+\mathbb{P}_1(\mathbb{R}^d) \leq 1$, ϕ is a consistent loss with $C^*_\phi(1/2) = \phi(0)$, and the adversarial Bayes classifier is unique up to degeneracy. Let $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ be maximizers of $\bar{R}_\phi$ for which $\mathbb{P}^*(\eta^* = 1/2) = 0$, with $\mathbb{P}^* = \mathbb{P}^*_0 + \mathbb{P}^*_1$ and $\eta^* = d\mathbb{P}^*_1/d\mathbb{P}^*$. Define $H(z) = \mathrm{conc}(\mathbb{P}^*(|\eta^* - 1/2| \leq z))$, Ψ as Theorem 2, and let $\tilde{\Lambda}(z) = \Psi^{-1}(\min(\frac{z}{4}, \phi(0)))$. Then*

$$R^\epsilon(f) - R^\epsilon_* \leq \tilde{\Phi}(R^\epsilon_\phi(f) - R^\epsilon_{\phi,*})$$

with

$$\tilde{\Phi}(z) = 4\left(\mathrm{id} + \min(1, \sqrt{-eH \ln H})\right) \circ \tilde{\Lambda}$$

This theorem is established under the assumption $\mathbb{P}_0(\mathbb{R}^d)+\mathbb{P}_1(\mathbb{R}^d) \leq 1$, which serves as an essential intermediate step for extending the result to case where the adversarial Bayes classifier is not uniquely defined up to degeneracy. See Example 3 and Figure 3c for an example of calculating a distribution-dependent surrogate risk bound.

The function H is always continuous and satisfies $H(0) = 0$, ensuring that this bound is non-vacuous (see Lemma 7 in Section 5). Further notice that $H \ln H$ approaches zero as $H \to 0$.

The map $\tilde{\Phi}$ combines two components: $\tilde{\Lambda}$, a modified version of Ψ^{-1}, and H, a modification of the cdf of $|\eta^* - 1/2|$. The function $\tilde{\Lambda}$ is a scaled version of Ψ^{-1}, where Ψ is the surrogate risk bound in the non-adversarial case of Theorem 2. The domain of Ψ^{-1} is $[0, \phi(0)]$, and thus the role of the min in the definition of $\tilde{\Lambda}$ is to truncate the argument so that it fits into this domain. The factor of $1/4$ in this function appears to be an artifact of our proof, see Section 5 for further discussion. In contrast, the map H translates the distribution of η^* into a surrogate risk transformation. Compare with Theorem 4, which states that consistency fails if $\mathbb{P}^*(\eta^* = 1/2) > 0$; accordingly, the function H becomes a poorer bound when more mass of η^* is near $1/2$.

If $\mathbb{P}^*(\eta^* = 1/2)$ is small, this result can still provide an informative surrogate bound.

Theorem 12. *Assume that there exist maximizers $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ of $\bar{R}_\phi$ that are induced by transport maps from $\mathbb{P}^*_0$, $\mathbb{P}^*_1$, and define $\mathbb{P}^* = \mathbb{P}^*_1 + \mathbb{P}^*_0$, $\eta^* = d\mathbb{P}^*_1/d\mathbb{P}^*$. Let $\tilde{\Phi}$ be the function in Theorem 11, but with H defined as $H(z) = conc(\mathbb{P}^*(0 < |\eta^* - 1/2| \leq z))$. Then*

$$R^\epsilon(f) - R^\epsilon_* \leq \tilde{\Phi}(R^\epsilon_\phi(f) - R^\epsilon_{\phi,*}) + \frac{\mathbb{P}^*(\eta^* = 1/2)}{2}$$

Removing the assumption that $\mathbb{P}^*_0, \mathbb{P}^*_1$ are induced by a transport map from Theorems 10 and 12 remains an open problem. We conjecture that this assumption is, in fact, unnecessary.

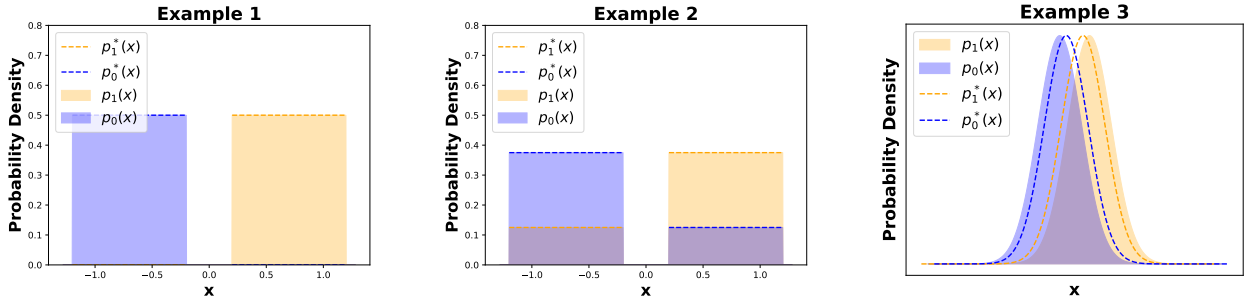

(a) (b) (c)

Figure 3: Distributions from Examples 1 to 3 along with attacks that maximize the dual $\bar{R}_\phi$.

Comparison with real-world datasets

Experimental results from prior work suggest that, in real-world datasets, η^* is typically concentrated near 0 and 1. Bhagoji et al. (2019) compute lower bounds on the adversarial classification risk for binary tasks, focusing on classifying digits '3' and '7' in MNIST under ℓ_2 perturbations. Their lower bound remains close to 0 for $\epsilon \leq 3$ and increases to 0.2 at $\epsilon = 4$. Since $C^*(\eta^*)$ attains its maximum at $\eta^* = 1/2$, a small adversarial risk implies that the distribution places little mass in a neighborhood of $|\eta^* - 1/2| = 0$. Similar trends are observed on Fashion MNIST and CIFAR10. Dai et al. (2023) extend these bounds to the multiclass setting, though extending adversarial surrogate bounds beyond binary classification remains an open problem.

When the optimal adversarial risk is non-zero, the adversarial Bayes classifier may not be unique up to degeneracy. Even without adversarial perturbations, datasets like MNIST and CIFAR10 contain inherently ambiguous examples. Northcutt et al. (2021) identify such cases, four are depicted in Figure 2. One would expect $\eta(\mathbf{x}) = 1/2$ for such examples. Bartoldson et al. (2024) show that similar ambiguity arises in adversarial settings: under ℓ_∞ perturbations of size 8/255, approximately 6% of adversarial examples are ambiguous in the CIFAR10 dataset. In the binary scenario, one would thus expect $\eta^*(x) = 1/2$ for these inputs, and thus one must apply Theorem 10 or Theorem 12. Extending the concept of uniqueness of the adversarial Bayes classifier to multiclass settings remains an open problem.

Examples

Below we present three examples illustrating the applicability of our main theorems. All examples involve one-dimensional distributions, and we denote the pdfs of $\mathbb{P}_0$ and $\mathbb{P}_1$ by p_0 and p_1.

To start, if $R_*^\epsilon = 0$ then $\eta^* \in \{0, 1\}$ $\mathbb{P}^*$-a.e. for any maximizers of $\bar{R}_\phi$. Therefore, for any such distribution, the optimal attack satisfies Massart's noise condition with $\alpha = 1/2$, see Appendix J.1 for a proof.

Lemma 2. *Assume $R_*^\epsilon = 0$, let $(\mathbb{P}_0^*, \mathbb{P}_1^*)$ maximize $\bar{R}_\phi$, and define $\mathbb{P}^* = \mathbb{P}_0^* + \mathbb{P}_1^*$, $\eta^* = d\mathbb{P}_1^*/d\mathbb{P}^*$. Then $\mathbb{P}^*(\eta^* \in \{0, 1\}) = 1$.*

Any distribution for which the supports of $\mathbb{P}_0$, $\mathbb{P}_1$ are more than 2ϵ apart must have zero risk.

Example 1 (When $R_*^\epsilon = 0$)**.** Let

$$p_0(x) = \begin{cases} 1 & \text{if } x \in [-1-\delta, -\delta] \\ 0 & \text{otherwise} \end{cases} \qquad p_1(x) = \begin{cases} 1 & \text{if } x \in [\delta, 1+\delta] \\ 0 & \text{otherwise} \end{cases}$$

for some $\delta > 0$. See Figure 3a for a depiction. This distribution satisfies $R_*^\epsilon = 0$ for all $\epsilon \leq \delta$ and thus Lemma 2 implies that the surrogate bound of Theorem 9 applies.

Examples 2 and 3 require computing maximizers to the dual $\bar{R}_\phi$; See Appendices J.2 and J.3 for these calculations. The following example illustrates a distribution for which Massart's noise condition can be verified for a distribution of optimal attacks.

Example 2 (Massart's noise condition)**.** Let $\delta > 0$ and let p be the uniform density on $[-1-\delta, -\delta] \cup [\delta, 1+\delta]$. Define η by

$$\eta(x) = \begin{cases} \frac{1}{4} & \text{if } x \in [-1-\delta, -\delta] \\ \frac{3}{4} & \text{if } x \in [\delta, 1+\delta] \end{cases} \tag{16}$$

see Figure 3b for a depiction of p_0 and p_1. For this distribution and $\epsilon \leq \delta$, the minimal surrogate and adversarial surrogate risks are always equal ($R_{\phi,*} = R^{\epsilon}_{\phi,*}$). This fact together with Theorem 7 imply that optimal attacks on this distribution are $\mathbb{P}^*_1 = \mathbb{P}_1$ and $\mathbb{P}^*_0 = \mathbb{P}_0$, see Appendix J.2 for details. Consequently: the distribution of optimal attacks $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ satisfies Massart's noise condition with $\alpha = 1/4$ and as a result the bounds of Theorem 9 apply. When $\epsilon \in (\delta, 1+\delta)$, pdfs of the distributions that maximize the dual are $p^*_1(x) = p_1(x+\epsilon)$, $p^*_0(x) = p_0(x-\epsilon)$, where $p_1(x) = \eta(x)p(x)$ and $p_0(x) = (1-\eta(x))p(x)$. These distributions satisfy $\mathbb{P}^*(\eta = 1/2) = (\epsilon - \delta)$ while $\mathbb{P}^*(|\eta - 1/2| \geq 1/4) = 1 - (\epsilon - \delta)$. Thus Theorem 10 provides a surrogate bound.

The final example presents a case in which Massart's noise condition fails for the distribution of optimal adversarial attacks, yet the adversarial Bayes classifier remains unique up to degeneracy. Theorem 11 still yields an informative surrogate bound.

Example 3 (Gaussian example)**.** Consider an equal-variance Gaussian mixture with $\mu_0 + 2\epsilon < \mu_1 < \mu_0 + \sqrt{2}\sigma$:

$$p_0(x) = \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu_0)^2}{2\sigma^2}}, \quad p_1(x) = \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu_1)^2}{2\sigma^2}},$$

see Figure 3c for a depiction. The optimal attacks $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ are gaussians centered at $\mu_0 + \epsilon$ and $\mu_1 - \epsilon$ respecively, with pdfs

$$p^*_0(x) = \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-(\mu_0+\epsilon))^2}{2\sigma^2}}, \quad p^*_1(x) = \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-(\mu_1-\epsilon))^2}{2\sigma^2}}. \tag{17}$$

We verify that $\mathbb{P}^*_0$ and $\mathbb{P}^*_1$ are in fact optimal by finding a function f^* for which $R^{\epsilon}_{\phi}(f^*) = \bar{R}_{\phi}(\mathbb{P}^*_0, \mathbb{P}^*_1)$, the strong duality result in Theorem 7 will then imply that $\mathbb{P}^*_0$ and $\mathbb{P}^*_1$ must maximize the dual $\bar{R}_{\phi}$, see Appendix J.3 for details.

Further, when $\mu_1 - \mu_0 \leq \sqrt{2}\sigma$, then the function $h(z) = \mathbb{P}^*(|\eta^* - 1/2| \leq z)$ is concave in z and consequently $H = h$, see Appendix J.4 for details. Although h is unwieldy function, comparison to its linear approximation at zero gives the bound

$$H(z) \leq \min\left(\frac{16\sigma^2}{\mu_1 - \mu_0 - 2\epsilon} z, 1\right). \tag{18}$$

Again, see Appendix J.4 for details.

When $\epsilon \geq (\mu_1 - \mu_0)/2$, Frank (2024, Example 4.1) demonstrates that the adversarial Bayes classifier is not unique up to degeneracy. Notably, the bound in preceding example deteriorates as $(\mu_1 - \mu_0)/2 \to \epsilon$, and then fails entirely when $\epsilon = (\mu_1 - \mu_0)/2$.

4 Proof of Linear Surrogate Bounds

4.1 Proof of Theorem 9

The proof of Theorem 9 relies on decomposing the excess adversarial classification and surrogate risks into non-negative terms, revealing their structural similarity and allowing for a pointwise comparison.

Let $\mathbb{P}^*_0, \mathbb{P}^*_1$ be any maximizers of $\bar{R}_{\phi}$. These distributions also maximize $\bar{R}$ by Theorem 8. Set $\mathbb{P}^* = \mathbb{P}^*_0 + \mathbb{P}^*_1$, $\eta^* = d\mathbb{P}^*_1/d\mathbb{P}^*$. Let γ^*_0, γ^*_1 be couplings between $\mathbb{P}_0$, $\mathbb{P}^*_0$ and $\mathbb{P}_1$, $\mathbb{P}^*_1$ achieving the W_{∞} distances (11). The excess classification risk can be decomposed as

$$R^{\epsilon}(f) - R^{\epsilon}_* = R^{\epsilon}(f) - \bar{R}(\mathbb{P}^*_0, \mathbb{P}^*_1) = \int i_1(f) d\gamma^*_1 + \int i_0(f) d\gamma^*_0 \tag{19}$$

with

$$i_1(f) = \big(S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x}) - \mathbf{1}_{f\le 0}(\mathbf{x}')\big) + \big(C(\eta^*, f) - C^*(\eta^*)\big)$$
$$i_0(f) = \big(S_\epsilon(\mathbf{1}_{f>0})(\mathbf{x}) - \mathbf{1}_{f>0}(\mathbf{x}')\big) + \big(C(\eta^*, f) - C^*(\eta^*)\big).$$

The first term measures the discrepancy between the worst-case attack on f and the attack induced by $\mathbb{P}_0^*, \mathbb{P}_1^*$, the optimal attack for the distribution $\mathbb{P}_0, \mathbb{P}_1$. The second term measures the excess conditional risk under the optimal attack $\mathbb{P}_0^*, \mathbb{P}_1^*$. Lemma 1 implies that $S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x}) - \mathbf{1}_{f\le 0}(\mathbf{x}')$ must be positive, while the definition of C^* implies that $C(\eta^*, f) - C^*(\eta^*) \ge 0$.

Similarly, one can express the excess surrogate risk as

$$R_\phi^\epsilon(f) - R_{\phi,*}^\epsilon = \int i_1^\phi(f) d\gamma_1^* + \int i_0^\phi(f) d\gamma_0^* \tag{20}$$

with

$$i_1^\phi(f) = \big(S_\epsilon(\phi(f))(\mathbf{x}) - \phi(f)(\mathbf{x}')\big) + \big(C_\phi(\eta^*, f) - C_\phi^*(\eta^*)\big)$$
$$i_0^\phi(f) = \big(S_\epsilon(\phi(-f))(\mathbf{x}) - \phi(-f)(\mathbf{x}')\big) + \big(C_\phi(\eta^*, f) - C_\phi^*(\eta^*)\big)$$

The following lemma is the core inequality linking i_k to i_k^ϕ under Massart's noise condition. It shows that each classification-risk term can be bounded by a constant multiple of its surrogate risk counterpart.

Lemma 3. *Define i_0^ϕ, i_1^ϕ as in* (20) *and assume that the distribution of optimal adversarial attacks $\mathbb{P}_0^*$, $\mathbb{P}_1^*$ satisfies Massart's noise condition. Then*

$$i_0(f) \le \frac{1}{\phi(0) - C_\phi^*(1/2-\alpha)} i_0^\phi(f). \tag{21}$$
$$i_1(f) \le \frac{1}{\phi(0) - C_\phi^*(1/2-\alpha)} i_1^\phi(f). \tag{22}$$

hold γ_0^*-a.e. and γ_1^*-a.e. respectively.

Lemma 3 directly implies Theorem 9 by integration over couplings γ_1^*, γ_0^*.

Proof of Theorem 9. Combine (19), Lemma 3, and (20). □

4.2 Proof of Lemma 3

The proof proceeds by partitioning the domain $\mathbb{R}^d \times \mathbb{R}^d$ into regions where the supremum-based classification either matches (D_k) or exceeds (E_k) the decision under the optimal attack. On each region, we derive a separate bound relating i_k and i_k^ϕ. Define the sets D_k, E_k,

$$D_0 = \{(\mathbf{x}, \mathbf{x}') : S_\epsilon(\mathbf{1}_{f>0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')} = 0\} \tag{23}$$
$$E_0 = \{(\mathbf{x}, \mathbf{x}') : S_\epsilon(\mathbf{1}_{f>0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')>0} = 1\} \tag{24}$$
$$D_1 = \{(\mathbf{x}, \mathbf{x}') : S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\le 0} = 0\} \tag{25}$$
$$E_1 = \{(\mathbf{x}, \mathbf{x}') : S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\le 0} = 1\} \tag{26}$$

By construction, $D_1 \cup E_1 = \mathbb{R}^d \times \mathbb{R}^d$ and $D_0 \cup E_0 = \mathbb{R}^d \times \mathbb{R}^d$.

The following lemma records a simple but useful structural property of E_0 and E_1, which allows us to bound the surrrogate loss terms from below on these sets.

Lemma 4. *Let E_k be as in Equations* (24) *and* (26)*. Then $S_\epsilon(\mathbf{1}_{f>0})(\mathbf{x}) = \mathbf{1}_{f>0}(\mathbf{x}') = 1$ γ_1^*-a.e. on E_1 while $S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x}) = \mathbf{1}_{f\le 0}(\mathbf{x}') = 1$ γ_0^*-a.e. on E_0.*

Proof. We'll prove the statement for E_1, the argument for E_0 is analogous. Specifically, we will show that one cannot simultaneously have $S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x}) - \mathbf{1}_{f\le 0}(\mathbf{x}') = 1$ and $S_\epsilon(\mathbf{1}_{f>0})(\mathbf{x}) - \mathbf{1}_{f>0}(\mathbf{x}') = 1$.

Consider $(\mathbf{x}, \mathbf{x}') \in E_1$: as both $S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x})$ and $\mathbf{1}_{f\le 0}(\mathbf{x}')$ are 0-1 valued, the relation $S_\epsilon(\mathbf{1}_{f\le 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\le 0} = 1$ implies that $\mathbf{1}_{f(\mathbf{x}')\le 0} = 0$ and thus $\mathbf{1}_{f(\mathbf{x}')>0} = 1$. The fact that $S_\epsilon(\mathbf{1}_{f>0})(\mathbf{x}) \ge \mathbf{1}_{f>0}(\mathbf{x}')$ on $\operatorname{supp}\gamma_1^*$ and $\operatorname{supp}\gamma_1^* \subset \Delta_\epsilon$ implies that $S_\epsilon(\mathbf{1}_{f>0})(\mathbf{x}) = 1$ γ_1^*-a.e. on E_1.

□

The next result bounds the terms i_k^ϕ from below.

Lemma 5. *The relations (27) and (28) hold on E_0 and E_1 respectively.*

$$i_0^\phi(f) \geq \phi(0) - C_\phi^*(\eta^*). \tag{27}$$

$$i_1^\phi(f) \geq \phi(0) - C_\phi^*(\eta^*). \tag{28}$$

Proof. We will prove the statement for E_1, the argument for E_0 is analogous. Observe that

$$i_1^\phi(f) = S_\epsilon(\phi(f))(\mathbf{x}) + (1-\eta^*)(\phi(-f(\mathbf{x}')) - \phi(f(\mathbf{x}'))) - C_\phi^*(\eta^*)$$

Now as $S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) = 1$, one can conclude that there is a point in $\mathbf{z} \in \overline{B_\epsilon(\mathbf{x})}$ for which $f(\mathbf{z}) \leq 0$, and thus $S_\epsilon(\phi(f))(\mathbf{x}) \geq \phi(0)$. Next, Lemma 4 implies that $f(\mathbf{x}') > 0$ and hence $\phi(-f(\mathbf{x}')) - \phi(f(\mathbf{x}')) \geq 0$. Therefore, one can conclude (28).

□

Furthermore, a simple calculation bounds the i_k from above.

Lemma 6. *On the set D_k*

$$i_k(f) = C(\eta^*, f) - C^*(\eta^*) \tag{29}$$

while on E_k

$$i_k(f) = 1 + C(\eta^*, f) - C^*(\eta^*) \tag{30}$$

Proof. We will show the statement $k = 1$ the argument for $k = 0$ is analogous. On D_1, $S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0} = 0$, implying (29). Similarly, on E_1, $S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0} = 1$, implying (30).

□

Comparing the upper and lower bounds present in Lemmas 4 and 6 proves Lemma 3.

Proof of Lemma 3. We will discuss (22), the argument for (21) is analogous. We prove the bound separately on D_1 and E_1, whose union is $\mathbb{R}^d$. First, notice that (8) implies that

$$C(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C^*(\eta^*(\mathbf{x}')) \leq \frac{1}{\phi(0) - C_\phi^*(1/2-\alpha)} \left(C_\phi(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C_\phi^*(\eta^*(\mathbf{x}'))\right) \quad \mathbb{P}^*\text{-a.e.} \tag{31}$$

On the set D_1:
Lemma 6 implies that

$$i_1(f) = C(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C^*(\eta^*(\mathbf{x}'))$$

and thus the desired inequality follows from (31) and the fact that $S_\epsilon(\phi \circ f)(\mathbf{x}) - \phi \circ f(\mathbf{x}') \geq 0$ γ_1^*-a.e.

On the set E_1:
On E_1,

$$i_1(f) = 1 + C(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C^*(\eta^*(\mathbf{x}'))$$

However, due to Lemma 5,

$$S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f\leq 0}(\mathbf{x}') = 1 = \frac{\phi(0) - C_\phi^*(\eta^*)}{\phi(0) - C_\phi^*(\eta^*)} \leq \frac{1}{\phi(0) - C_\phi^*(\eta^*)}(S_\epsilon(\phi \circ f)(\mathbf{x}) - \phi(f(\mathbf{x}')) \tag{32}$$

The last inequality is a consequence of the assumption $|\eta^* - 1/2| \leq \alpha$. Summing this relation with (31) shows (22). □

4.3 Lower Bounds

The bound in Theorem 9 provides a general guarantee relating the adversarial classification and surrogate risks. We now show that this bound cannot be substantially improved.

This example describes functions in which the worst-case attack and the attack induced by $\mathbb{P}_0^*$, $\mathbb{P}_1^*$ differ substantially. This function sequence is the counterexample to consistency proposed in prior work (Meunier et al., 2022; Li & Telgarsky, 2023). Intuitively, the sign flip causes the classifier to misclassify both classes, even though a constant function would achieve lower risk for this distribution.

Example 4 (Lower bound for Theorem 9)**.** Let ϕ satisfy $C_\phi^*(1/2) = \phi(0)$, and consider a distribution supported on $[-\epsilon, \epsilon]$ with $\mathbb{P}^*(\eta = 1/2 + \alpha) = 1$. Define the sequence of functions

$$f_n = \begin{cases} \frac{1}{n} & \mathbf{x} \neq 0 \\ -\frac{1}{n} & \mathbf{x} = 0 \end{cases} \tag{33}$$

For this sequence, $R^\epsilon(f_n) - R_*^\epsilon = 1/2 + \alpha$ while the adversarial surrogate risk converges to $\lim_{n\to\infty} R_\phi^\epsilon(f_n) = \phi(0) - C_\phi^*(1/2 + \alpha)$. Consequently,

$$\lim_{n\to\infty} \frac{R^\epsilon(f_n) - R_*^\epsilon}{R_\phi^\epsilon(f_n) - R_{\phi,*}^\epsilon} = \frac{\frac{1}{2} + \alpha}{\phi(0) - C_\phi^*(1/2 - \alpha)}.$$

It follows that for any $\delta > 0$ there exists f such that

$$R^\epsilon(f) - R_*^\epsilon \geq \frac{1/2 + \alpha}{\phi(0) - C_\phi^*(1/2 - \alpha)} \left(R_\phi^\epsilon(f) - R_{\phi,*}^\epsilon\right) - \delta.$$

In particular, the constant in Theorem 9 is overestimated by factor of at most $1/(1/2 + \alpha) \leq 2$. However, this example demonstrates that Theorem 9 is tight when $\alpha = 1/2$.

The constant in Theorem 9 is known to be sub-optimal when $\alpha < 1/2$. In particular, Theorem 4 of Frank (2025) proves that $R^\epsilon(f) - R_*^\epsilon \leq (1/2)/(\phi(0) - C_{\phi_\rho}^*(1/2))(R_{\phi_\rho}^\epsilon(f) - R_{\phi_\rho,*}^\epsilon)$ for the ρ-margin loss $\phi_\rho(\alpha) = \min(1, \max(0, 1 - \alpha/\rho))$. We conjecture that the tight constant in Theorem 9 is in fact $(1/2 + \alpha)/(\phi(0) - C_\phi^*(1/2 - \alpha))$. Together, these observations indicate that the bound in Theorem 9 captures the correct order of dependence on α and ϕ, and that only the numerican constant can potentially be improved.

5 Proof of Theorem 11

Before proving Theorem 11, we will show that this bound is non-vacuous when the adversarial Bayes classifier is unique up to degeneracy. The function $h(z) = \mathbb{P}(|\eta^* - 1/2| \leq z)$ is a cdf, and is thus right-continuous in z. Furthermore, if the adversarial Bayes classifier is unique up to degeneracy, then $h(0) = 0$. The following lemma implies that if $H = \mathrm{conc}(h)$ then H is continuous at 0 with $H(0) = 0$. See Appendix E for a proof. This result implies that the bound in Theorem 11 is non-vacuous.

Lemma 7. *Let $h : [0, 1/2] \to \mathbb{R}$ be a non-decreasing function with $h(0) = 0$ and $h(1/2) = 1$ that is right-continuous at 0. Then $\mathrm{conc}(h)$ is non-decreasing, continuous on $[0, 1/2]$, and $\mathrm{conc}(h)(0) = 0$.*

The first step in proving Theorem 11 is showing an analog of Theorem 9 with $\alpha = 0$ for which the linear function is replaced by an η-dependent concave function.

Proposition 2. *Let Φ be a concave non-decreasing function for which $C(\eta, \alpha) - C^*(\eta) \leq \Phi(C_\phi(\eta, \alpha) - C_\phi^*(\eta))$ for any $\eta \in [0, 1]$ and $\alpha \in \overline{\mathbb{R}}$. Let $\mathbb{P}_0^*$, $\mathbb{P}_1^*$ be any two maximzers of $\bar{R}_\phi$ for which $\mathbb{P}^*(\eta^* = 1/2) = 0$, where $\mathbb{P}^* = \mathbb{P}_0^* + \mathbb{P}_1^*$ and $\eta^* = d\mathbb{P}_1^*/d\mathbb{P}^*$. Let $G : [0, \infty) \to \mathbb{R}$ be any non-decreasing concave function for which the quantity*

$$K = \int \frac{1}{G(\phi(0) - C_\phi^*(\eta^*))} d\mathbb{P}^*$$

is finite. Then $R^\epsilon(f) - R^\epsilon_* \leq \tilde{\Phi}(R^\epsilon_\phi(f) - R^\epsilon_{\phi,*})$*, where*

$$\tilde{\Phi}(z) = 4\sqrt{KG\left(\frac{1}{4}z\right)} + 2\Phi\left(\frac{1}{2}z\right) \tag{34}$$

The proof strategy mirrors that of Theorem 9, but with Φ and G replacing the fixed constant bound.

Uniqueness up to degeneracy and Theorem 1 guarantee that the denominator $\phi(0) - C^*_\phi(\eta^*)$ is strictly positive $\mathbb{P}^*$-a.e. The function Ψ^{-1} in Theorem 2 and the surrogate bounds of Zhang (2004) provide examples of candidate functions for Φ. As before, this result is proved by dividing the risks R^ϵ_ϕ, R^ϵ as the sum of four terms as in (19), (20) and then bounding these quantities over the sets D_k, E_k defined in (25),(26) separately.

The factor of $1/4$ in (34) arises as an artifact of the proof technique: one factor of 2 from averaging over two sets D_1, E_1, (see (56) in Appendix F), and another factor of 2 from combining the bounds associated with the two integrals corresponding to class 0 and class 1(see Equations (54) and (56) in Appendix F).

We now turn to the problem of identifying functions G for which the constant K in the preceding proposition is guaranteed to be finite when the adversarial Bayes classifier is unique, but distribution dependent. Observe that if h is the cdf of $|\eta - 1/2|$ and h is continuous, then $\int 1/h^r dh$ is always finite for $r \in (0,1)$. This calculation suggests $\Phi = h \circ \Psi^{-1}$, with Ψ defined in Theorem 2. To ensure the concavity of G, we instead select $G = H \circ \Psi^{-1}$ with $H = \mathrm{conc}(h)$.

Lemma 8. *Assume* $C^*_\phi(1/2) = \phi(0)$*. Let* $\mathbb{P}_1$*,* $\mathbb{P}_0$*,* $\mathbb{P}^*_1$*,* $\mathbb{P}^*_0$*,* ϕ*,* H*, and* Ψ *be as in Theorem 11. Let* $\Lambda(z) = \Psi^{-1}(\min(z, \phi(0)))$*. Then for any* $r \in (0,1)$*,*

$$R^\epsilon(f) - R^\epsilon_* \leq \tilde{\Phi}(R^\epsilon_\phi(f) - R^\epsilon_{\phi,*}) \tag{35}$$

with

$$\tilde{\Phi}(z) = 4\sqrt{\frac{1}{1-r}H\left(\frac{1}{2}\Lambda\left(\frac{1}{4}z\right)\right)^r} + 2\Lambda\left(\frac{z}{2}\right). \tag{36}$$

Proof. For convenience, let $G = (H \circ \frac{1}{2}\Lambda)^r$. Then G is concave because it is the composition of a concave function and an increasing concave function. We will verify that K is finite and yields the constant in the bound:

$$K = \int \frac{1}{G(\phi(0) - C^*_\phi(\eta^*))} d\mathbb{P}^* \leq \frac{1}{1-r}$$

First,

$$\int \frac{1}{G(\phi(0) - C^*_\phi(\eta^*))} d\mathbb{P}^* = \int \frac{1}{H(|\eta^* - 1/2|)^r} d\mathbb{P}^* = \int_{[0,\frac{1}{2}]} \frac{1}{H(s)^r} d\mathbb{P}^* \sharp s = \int_{(0,\frac{1}{2}]} \frac{1}{H(s)^r} d\mathbb{P}^* \sharp s$$

with $s = |\eta^* - 1/2|$. The assumption $\mathbb{P}^*(|\eta^* = 1/2|) = 0$ allows us to drop 0 from the domain of integration. Because the function H is continuous on $(0,1]$ by Lemma 7, this last expression can be evaluated as a Riemann-Stieltjes integral with respect to the function $h(s) = \mathbb{P}(|\eta^* - 1/2| \leq s)$:

$$\int_{(0,\frac{1}{2}]} \frac{1}{H(s)^r} d\mathbb{P}^* \sharp s = \int_0^{1/2} \frac{1}{H(s)^r} dh \tag{37}$$

This result is standard when $\mathbb{P}^*$ is Lebesgue measure, (see for instance Theorem 5.46 of Wheeden & Zygmund (1977)). We prove equality in (37) for strictly decreasing functions in Proposition 4 in Appendix G.1.

Finally, the integral in (37) can be bounded as

$$\int \frac{1}{H(s)^r} dh \leq \frac{1}{1-r} \tag{38}$$

If h were differentiable, then the chain rule would imply

$$\int \frac{1}{H(s)^r} dh \leq \int \frac{1}{h(s)^r} dh = \int_0^{\frac{1}{2}} \frac{1}{h(s)^r} h'(s) dz = \frac{1}{1-r} h(s)^{1-r} \Big|_0^{\frac{1}{2}} \leq \frac{1}{1-r}.$$

This calculation is more delicate for non-differentiable H; we formally prove inequality in (38) in Appendix G.2.

This calculation proves the inequality (35) with $\tilde{\Phi}$ as (36) □

To obtain the bound in Theorem 11, first observe that the concavity of Λ together with the fact that $\Lambda(0) = 0$ implies that $2\Lambda(z/2) \leq 4\Lambda(z/4)$. Next, minimizing the bound in Lemma 8 over r then produces Theorem 11, see Appendix H for details.

6 Proof of Theorems 10 and 12

The main insight behind Theorems 10 and 12 is that a transport map that realizes the optimal adversarial perturbations also preserves optimality when restricted to certain subsets of $\mathbb{R}^d$, allowing a reduction from the global to a local problem in both the dual and primal formulations. The following lemma formalizes the fact that under a transport map structure, restricting the primal problem to the pre-image of a set Q corresponds exactly to restricting the dual maximizers to Q itself.

Lemma 9. *Let $\mathbb{P}_0, \mathbb{P}_1$ be a data distribution and let $\mathbb{P}_0^* \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_0), \mathbb{P}_1^* \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_1)$ maximize $\bar{R}_\phi$. Assume there exists transport maps T_0, T_1 for which $\mathbb{P}_i^* = \mathbb{P}_i \sharp T_i$ with $\|T_i(\mathbf{x}) - \mathbf{x}\| \leq \epsilon$. Let Q be any set and define $U_i = T_i^{-1}(Q)$.*

If the data is distributed according to $\mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}$, then $\mathbb{P}_0^*|_Q$, $\mathbb{P}_1^*|_Q$ maximize $\bar{R}_\phi$ over $\mathcal{B}_\epsilon^\infty(\mathbb{P}_0|_{U_0}) \times \mathcal{B}_\epsilon^\infty(\mathbb{P}_1|_{U_1})$.

In the remainder of this section it will be useful to include the data distribution in the notation for the primal problem. Thus, for the remainder of this section, we define

$$R_\phi^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1) = \int S_\epsilon(\phi \circ f) d\mathbb{P}_1 + \int S_\epsilon(\phi \circ -f) d\mathbb{P}_0 \quad R^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1) = \int S_\epsilon(\mathbf{1}_{f \leq 0}) d\mathbb{P}_1 + \int S_\epsilon(\mathbf{1}_{f>0}) d\mathbb{P}_0 \quad (39)$$

Similarly, we'll denote

$$R_{\phi,*}^\epsilon(\mathbb{P}_0, \mathbb{P}_1) = \inf_f R_\phi^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1), \quad R_*^\epsilon(\mathbb{P}_0, \mathbb{P}_1) = \inf_f R^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1) \quad (40)$$

Observe that for any two sets U_0, U_1,

$$R_\phi^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1) = R_\phi^\epsilon(f; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) + R_\phi^\epsilon(f; \mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C})$$

This decomposition reflects the fact that the adversarial surrogate risk is additive over disjoint measurable partitions of the data space. If furthermore these sets are induced by transport maps, then the optimal risks also follow this split.

Lemma 10. *Let $\mathbb{P}_0^*, \mathbb{P}_1^*, T_0, T_1, U_0, U_1$ and Q be as in Lemma 9, and define $\mathbb{P}^* = \mathbb{P}_0^* + \mathbb{P}_1^*$, $\eta^* = d\mathbb{P}_1^*/d\mathbb{P}^*$. Then*

$$R_{\phi,*}^\epsilon(\mathbb{P}_0, \mathbb{P}_1) = R_{\phi,*}^\epsilon(\mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) + R_{\phi,*}^\epsilon(\mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C})$$

and furthermore, $R_{\phi,*}^\epsilon(\mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C}) = \int_{Q^C} C_\phi^*(\eta^*) d\mathbb{P}^*$.

See Appendix I.2 for a proof of Lemmas 9 and 10.

Proof of Theorem 12. Let $Q = \{\mathbf{x}' : \eta^*(\mathbf{x}') = 1/2\}$. Then Lemma 9 applied to Q^C shows that $(\mathbb{P}_0^*|_{Q^C}, \mathbb{P}_1^*|_{Q^C})$ maximize $\bar{R}_\phi$ over $\mathcal{B}_\epsilon^\infty(\mathbb{P}_0|_{U_0}) \times \mathcal{B}_\epsilon^\infty(\mathbb{P}_1|_{U_1})$, with $U_0 = T_0^{-1}(Q^C)$ and $U_1 = T_1^{-1}(Q^C)$. Theorem 11

and Lemma 10 imply that

$$\begin{aligned} R^\epsilon(f;\mathbb{P}_0|_{U_0},\mathbb{P}_1|_{U_1}) - R^\epsilon_*(\mathbb{P}_0|_{U_0},\mathbb{P}_1|_{U_1}) &\leq \tilde{\Phi}\left(R^\epsilon_\phi(f;\mathbb{P}_0|_{U_0},\mathbb{P}_1|_{U_1}) - R^\epsilon_{\phi*}(\mathbb{P}_0|_{U_0},\mathbb{P}_1|_{U_1})\right) \\ &\leq \tilde{\Phi}\left(R^\epsilon_\phi(f;\mathbb{P}_0,\mathbb{P}_1) - R^\epsilon_{\phi*}(\mathbb{P}_0,\mathbb{P}_1)\right). \end{aligned}$$

Next, by Lemma 10, adding $R^\epsilon(f;\mathbb{P}_0|_{U_0^C},\mathbb{P}_1|_{U_1^C}) - R^\epsilon_*(\mathbb{P}_0|_{U_0^C},\mathbb{P}_1|_{U_1^C})$ to both sides of the inequality above results in

$$R^\epsilon(f;\mathbb{P}_0,\mathbb{P}_1) - R^\epsilon_*(\mathbb{P}_0,\mathbb{P}_1) \leq \tilde{\Phi}\left(R^\epsilon_\phi(f;\mathbb{P}_0,\mathbb{P}_1) - R^\epsilon_{\phi*}(\mathbb{P}_0|_{U_0},\mathbb{P}_1)\right) + R^\epsilon(f;\mathbb{P}_0|_{U_0^C},\mathbb{P}_1|_{U_1^C}) - \int_Q C^*(\eta^*)d\mathbb{P}^*.$$

The fact that $C^*(\eta^*) = 1/2$ on Q while $S_\epsilon(\mathbf{1}_{f\leq 0}) \leq 1$, $S_\epsilon(\mathbf{1}_{f>0}) \leq 1$ implies that $R^\epsilon(f;\mathbb{P}_0|_{U_0^C},\mathbb{P}_1|_{U_1^C}) - \int_Q C^*(\eta^*)d\mathbb{P}^* \leq \frac{1}{2}\mathbb{P}(\eta^* = 1/2)$. Thus, the excess risk contribution from the region Q is at most $\mathbb{P}^*(\eta^* = 1/2)/2$. □

The proof of Theorem 10 follows the same steps, except that we take $Q = \{\mathbf{x}' : |\eta(\mathbf{x}') - 1/2| < \alpha\}$, see Appendix I.3 for a proof.

7 Related Works

Surrogate Risk Bounds: The statistical consistency of surrogate risks in both the standard and adversarial context has been widely studied. Bartlett et al. (2006); Zhang (2004) establish surrogate risk bounds that apply to the class of all measurable functions while Lin (2004); Steinwart (2007) prove further results on consistency in the standard setting. Frongillo & Waggoner (2021) study the optimally of such bounds, and Bao (2023) derive bounds using the modulus of convexity of C^*_ϕ to construct surrogate risk bounds. Several works (Philip M. Long, 2013; Mingyuan Zhang, 2020; Awasthi et al., 2022; Mao et al., 2023a;b; Awasthi et al., 2023b) study consistency within a restricted function class; a concept known as $\mathcal{H}$-*consistency*. Mahdavi et al. (2014) combine surrogate risk bounds with surrogate generalization bounds to study the generalization of the classification error.

Adversarial Surrogate Risk Bounds: Most closely related to our results are Li & Telgarsky (2023); Mao et al. (2023a). Li & Telgarsky (2023) derive a surrogate bound for convex losses in which the threshold in (10) is optimized rather than fixed at zero. Mao et al. (2023a) establish an adversarial surrogate bound for a modified ρ-margin loss.

Adversarial Consistency: In the adversarial setting, Meunier et al. (2022); Frank & Niles-Weed (2024a) characterize which losses are adversarially consistent for all data distributions. Frank (2025) show that under reasonable distributional assumptions, a consistent loss is adversarially consistent for a specific distribution iff the adversarial Bayes classifier is unique up to degeneracy. Awasthi et al. (2021) study adversarial consistency for a well-motivated class of linear functions while Awasthi et al. (2023b); Mao et al. (2023a) study $\mathcal{H}$-consistency in the adversarial setting for specific surrogate risks. Standard and adversarial surrogate risk bounds are a central tool in the derivation of the $\mathcal{H}$-consistency bounds in this line of research. Whether the adversarial surrogate bounds presented in this paper could result in improved adversarial $\mathcal{H}$-consistency bounds remains an open problem.

The Adversarial Bayes Classifier: Our proofs draw on prior work that investigates adversarial risks and adversarial Bayes classifiers. Bungert et al. (2021); Pydi & Jog (2021; 2020); Bhagoji et al. (2019); Awasthi et al. (2023a) establish existence results for the adversarial Bayes classifier, while Frank & Niles-Weed (2024b); Pydi & Jog (2020; 2021); Bhagoji et al. (2019); Frank (2025) prove minimax theorems for adversarial surrogate and classification risks. Pydi & Jog (2020) use such results to analyze the adversarial Bayes classifier, and Frank (2024) employ them to study uniqueness.

Sample Complexity and Surrogate Risks: The bound of Theorem 2 can be linear even for convex loss functions. For the hinge loss $\phi(\alpha) = \max(1 - \alpha, 0)$, the function ϕ computes to $\phi(\theta) = |\theta|$. Mahdavi et al. (2014) emphasize the importance of a linear convergence rate in a surrogate risk bound. They note that convex surrogates with favorable sample complexity often fail to satisfy strong surrogate risk bounds, due to Theorem 2 Frongillo & Waggoner (2021): convex losses which are locally strictly convex and Lipschitz achieve at best a square root surrogate risk rate. Thus, Proposition 1 suggests that favorable sample complexity guarantees for convex surrogates may require distributional conditions such as Massart's noise condition, under which Massart & Nédélec (2006) also show improved sample complexity.

8 Conclusion

In conclusion, we prove surrogate risk bounds for adversarial risks. When ϕ is adversarially consistent or the distribution of optimal adversarial attacks satisfies Massart's noise condition, we obtain a linear surrogate risk bound. In the general case, we prove a concave distribution-dependent bound. Understanding the optimality of the concave bound remains an open problem, as does understanding how these bounds interact with the sample complexity of estimating the surrogate risk. While related questions have been studied in the standard setting (Frongillo & Waggoner, 2021; Mahdavi et al., 2014), the adversarial context remains largely unexplored. Advancing these directions could bridge the current gap between theoretical guarantees and practical robustness in adversarial learning.

Acknowledgments

Natalie Frank was supported in part by the Research Training Group in Modeling and Simulation funded by the National Science Foundation via grant RTG/DMS – 1646339, NSF grant DMS-2210583, and NSF TRIPODS II - DMS 2023166.

References

Tom M. Apostol. Mathematical analysis, 1974.

Pranjal Awasthi, Natalie S. Frank, Anqui Mao, Mehryar Mohri, and Yutao Zhong. Calibration and consistency of adversarial surrogate losses. *NeurIps*, 2021.

Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. H-consistency bounds for surrogate loss minimizers. In *Proceedings of the 39th International Conference on Machine Learning*. PMLR, 2022.

Pranjal Awasthi, Natalie S. Frank, and Mehryar Mohri. On the existence of the adversarial bayes classifier (extended version). *arxiv*, 2023a.

Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Theoretically grounded loss functions and algorithms for adversarial robustness. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research. PMLR, 2023b.

Han Bao. Proper losses, moduli of convexity, and surrogate regret bounds. In *Proceedings of Thirty Sixth Conference on Learning Theory*, Proceedings of Machine Learning Research. PMLR, 2023.

Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473), 2006.

Brian R. Bartoldson, James Diffenderfer, Konstantinos Parasyris, and Bhavya Kailkhura. Adversarial robustness limits via scaling-law and human-alignment studies, 2024. URL `https://arxiv.org/abs/2404.09349`.

Arjun Nitin Bhagoji, Daniel Cullina, and Prateek Mittal. Lower bounds on adversarial robustness from optimal transport, 2019.

Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 387–402. Springer, 2013.

Leon Bungert, Nicolás García Trillos, and Ryan Murray. The geometry of adversarial training in binary classification. *arxiv*, 2021.

Sihui Dai, Wenxin Ding, Arjun Nitin Bhagoji, Daniel Cullina, Ben Y. Zhao, Haitao Zheng, and Prateek Mittal. Characterizing the optimal 0-1 loss for multi-class classification with a test-time attacker. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*, 2023. URL `https://arxiv.org/abs/2302.10722`.

Natalie S. Frank. A notion of uniqueness for the adversarial bayes classifier, 2024.

Natalie S. Frank. Adversarial consistency and the uniqueness of the adversarial bayes classifier. *European Journal of Applied Mathematics*, 2025.

Natalie S. Frank and Jonathan Niles-Weed. The adversarial consistency of surrogate risks for binary classification. *NeurIps*, 2024a.

Natalie S. Frank and Jonathan Niles-Weed. Existence and minimax theorems for adversarial surrogate risks in binary classification. *Journal of Machine Learning Research*, 2024b.

Rafael Frongillo and Bo Waggoner. Surrogate regret bounds for polyhedral losses. In *Advances in Neural Information Processing Systems*, 2021.

Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. Fundamentals of convex analysis, 2001.

Heikki Jylhä. The l^∞ optimal transport: Infinite cyclical monotonicity and the existence of optimal transport maps. *Calculus of Variations and Partial Differential Equations*, 2014.

Justin D. Li and Matus Telgarsky. On achieving optimal adversarial test error, 2023.

Yi Lin. A note on margin-based loss functions in classification. *Statistics & Probability Letters*, 68(1):73–82, 2004.

Mehrdad Mahdavi, Lijun Zhang, and Rong Jin. Binary excess risk for smooth convex surrogates, 2014.

Anqi Mao, Mehryar Mohri, and Yutao Zhong. Cross-entropy loss functions: Theoretical analysis and applications, 2023a.

Anqi Mao, Mehryar Mohri, and Yutao Zhong. Structured prediction with stronger consistency guarantees. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2023b.

Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34, 2006.

Stefanie Jegelka Matthew Staib. Distributionally robust learning as a generalization of adversarial training. *NeurIps*, 2017.

Laurent Meunier, Raphaël Ettedgui, Rafael Pinot, Yann Chevaleyre, and Jamal Atif. Towards consistency in adversarial classification. *arXiv*, 2022.

Shivani Agarwal Mingyuan Zhang. Consistency vs. h-consistency: The interplay between surrogate loss functions and the scoring function class. *NeurIps*, 2020.

Curtis G. Northcutt, Anish Athalye, and Jonas Mueller. Pervasive label errors in test sets destabilize machine learning benchmarks. In *Proceedings of the 35th Conference on Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets & Benchmarks Track)*, December 2021. Includes the companion website `https://labelerrors.com`.

Magdalini Paschali, Sailesh Conjeti, Fernando Navarro, and Nassir Navab. Generalizability vs. robustness: Adversarial examples for medical imaging. *Springer*, 2018.

Rocco A. Servedio Philip M. Long. Consistency versus realizable h-consistency for multiclass classification. *ICML*, 2013.

Muni Sreenivas Pydi and Varun Jog. Adversarial risk via optimal transport and optimal couplings. *ICML*, 2020.

Muni Sreenivas Pydi and Varun Jog. The many faces of adversarial risk. *Neural Information Processing Systems*, 2021.

Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 2007.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

Ambuj Tewari and Peter L. Bartlett. On the consistency of multiclass classification methods. *Journal of Machine Learning Research*, 8(36), 2007.

Richard L. Wheeden and Antoni Zygmund. *Measure and Integral*. Pure and Applied Mathematics. Marcel Dekker Inc., 1977.

Ying Xu, Kiran Raja, Raghavendra Ramachandra, and Christoph Busch. *Adversarial Attacks on Face Recognition Systems*, pp. 139–161. Springer International Publishing, Cham, 2022.

Tong Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *The Annals of Statistics*, 2004.

Contents of Appendix

A Proof of Theorem 1 **21**

B Linear Surrogate Risk Bounds—Proof of Proposition 1 **21**

C Proof of Lemma 1 **22**

D Proof of Item 1), Theorem 8 **23**

E Proof of Lemma 7 **25**

F Proof of Proposition 2 **27**

G Technical Integral Lemmas **28**

G.1 The Lebesgue and Riemann–Stieltjes integral of an increasing function 29

G.2 Proof of the last equality in (38) 30

H Optimizing the Bound of Lemma 8 over r **30**

I Deferred proofs from Section 6 **31**

I.1 Existence of Minimizers and Complementary slackness 31

I.2 Proof of Lemmas 9 and 10 31

I.3 Proof of Theorem 10 33

J Further details from Examples 1 to 3 **33**

J.1 Proof of Lemma 2 33

J.2 Calculating the optimal $\mathbb{P}_0^*$, $\mathbb{P}_1^*$ for Example 2 33

J.3 Calculating the optimal $\mathbb{P}_0^*$ and $\mathbb{P}_1^*$ for Example 3— Proof of (17) 35

J.4 Showing (18) 35

A Proof of Theorem 1

Lemma 11. *Assume ϕ is consistent. Then $C_\phi^*(\eta) = \phi(0)$ implies that $\eta = 1/2$.*

This result appeared as Lemma 7 of Frank (2025).

Proof. If ϕ is consistent and 0 minimizes $C_\phi(\eta, \alpha)$, then 0 must also minimize $C(\eta, \alpha) = \eta \mathbf{1}_{\alpha \leq 0} + (1-\eta)\mathbf{1}_{\alpha > 0}$ and consequently $\eta \leq 1/2$. However $C_\phi(\eta, \alpha) = C_\phi(1-\eta, -\alpha)$ so that 0 must minimize $C(1-\eta, -\alpha)$ as well. Consequently, $1 - \eta \leq 1/2$ and thus η must actually equal $1/2$. □

Proof of Theorem 1. **Forward direction:** Assume that ϕ is consistent. Note that $C_\phi^*(\eta) \leq C_\phi(\eta, 0) = \phi(0)$ for any η. Thus Lemma 11 implies that $C_\phi^*(\eta) < \phi(0)$ for $\eta \neq 1/2$.

Backward direction: Assume that $C_\phi^*(\eta) < \phi(0)$ for all $\eta \neq 1/2$. Notice that if $\eta = 1/2$, $C(1/2, \alpha)$ is constant in α so *any* sequence α_n minimizes $C(1/2, \cdot)$. We will show if $\eta > 1/2$ and α_n is a minimizing sequence of $C_\phi(\eta, \cdot)$, then $\alpha_n > 0$ for sufficiently large n, and thus must also minimize $C(\eta, \cdot)$. An analogous argument will imply that if $\eta < 1/2$, any minimizing sequence of $C_\phi(\eta, \cdot)$ must also minimize $C(\eta, \cdot)$ as well.

Assume $\eta > 1/2$ and let α_n be any minimizing sequence of $C_\phi(\eta, \cdot)$. Let α^* be a limit point of the sequence α_n in the extended real number line $\overline{\mathbb{R}}$. Then α^* is a minimizer of $C_\phi(\eta, \alpha)$. Next, observe that one of $\phi(\alpha^*)$, $\phi(-\alpha^*)$ is larger that or equal to $\phi(0)$ and the other is less than or equal to $\phi(0)$. As $\eta > 1/2$ and α^* is a minimizer of $C_\phi(\eta, \cdot)$ and $C_\phi(\eta, \alpha^*) < \phi(0)$, one can conclude that $\phi(\alpha^*) < \phi(0)$ and consequently $\alpha^* > 0$.

Therefore, every limit point of the sequence $\{\alpha_n\}$ is strictly positive. Consequently, one can conclude that $\alpha_n > 0$ for sufficiently large n.

□

B Linear Surrogate Risk Bounds—Proof of Proposition 1

In this appendix, we will find it useful to study the function

$$C_\phi^-(\eta) = \inf_{z(2\eta-1) \leq 0} C_\phi(\eta, z)$$

introduced by Bartlett et al. (2006). This function maps η to the smallest value of the conditional ϕ-risk assuming an incorrect classification. The symmetry $C_\phi(\eta, \alpha) = C_\phi(1-\eta, -\alpha)$ implies $C_\phi^-(\eta) = C_\phi^-(1-\eta)$. Further, the function C_ϕ^- is concave on each of the intervals $[0, 1/2]$ and $[1/2, 1]$, as it is an infimum of linear functions on each of these regions. The next result examines the monotonicity properties of C_ϕ^* and C_ϕ^-.

Lemma 12. *The function C_ϕ^* is non-decreasing on $[0, 1/2]$ and non-increasing on $[1/2, 1]$. In contrast, C_ϕ^- is non-increasing on $[0, 1/2]$ and non-decreasing on $[1/2, 1]$*

Proof. The symmetry $C_\phi^*(\eta) = C_\phi^*(1-\eta)$ and $C_\phi^-(\eta) = C_\phi^-(1-\eta)$ implies that it suffices to check monotonicity on $[0, 1/2]$. Observe that

$$C_\phi(\eta, \alpha) - C_\phi(\eta, -\alpha) = \eta(\phi(\alpha) - \phi(-\alpha)) + (1-\eta)(\phi(-\alpha) - \phi(\alpha)) = (2\eta - 1)(\phi(\alpha) - \phi(-\alpha)).$$

If $\eta \leq 1/2$, then this quantity is non-negative when $\alpha \leq 0$. Therefore, when computing C_ϕ^* over $[0, 1/2]$, it suffices to minimize $C_\phi(\eta, \alpha)$ over $\alpha \leq 0$. In other words, for $\eta \leq 1/2$,

$$C_\phi^*(\eta) = \inf_\alpha C_\phi(\eta, \alpha) = \inf_{\alpha \leq 0} C_\phi(\eta, \alpha)$$

For any fixed $\alpha \leq 0$, the quantity $C_\phi(\eta, \alpha)$ is non-increasing in η and thus $C_\phi^*(\eta_1) \leq C_\phi^*(\eta_2)$ when $\eta_1 \leq \eta_2 \leq 1/2$.

In contrast, for any $\alpha \geq 0$, the quantity $C_\phi(\eta, \alpha)$ is non-decreasing in η and thus $C_\phi^-(\eta_1) \geq C_\phi^-(\eta_2)$ when $\eta_1 \leq \eta_2 \leq 1/2$.

□

Next we'll prove a useful lower bound on C_ϕ^-.

Lemma 13. *For all $\eta \in [0,1]$,*

$$C_\phi^-(\eta) \geq |1-2\eta|\phi(0) + 2\min(\eta, 1-\eta)C_\phi^*(\eta) \tag{41}$$

Proof. First, assume that $\eta \leq 1/2$ and observe that η is the convex combination $\eta = 2\eta \cdot 1/2 + (1-2\eta)\cdot 0$. By the concavity of C_ϕ^- on $[0,1/2]$,

$$C_\phi^-(\eta) = C_\phi^-\left(2\eta \cdot \frac{1}{2} + (1-2\eta)\cdot 0\right) \geq (1-2\eta)C_\phi^-(0) + 2\eta C_\phi^-\left(\tfrac{1}{2}\right)$$

However, $C_\phi^-(0) = \phi(0)$ while $C_\phi^-(1/2) = C_\phi^*(1/2)$. Further, Lemma 12 implies that $C_\phi^*(1/2) \geq C_\phi^*(\eta)$, yielding the inequality

$$C_\phi^-(\eta) \geq (1-2\eta)\phi(0) + 2\eta C_\phi^*(\eta)$$

Symmetry $C_\phi^-(\eta) = C_\phi^-(1-\eta)$ then implies (41).

□

Proof of Proposition 1. If $C(\eta,f) - C^*(\eta) = 0$ then (8) holds trivially. Otherwise, $C(\eta,f) - C^*(\eta) = |2\eta - 1|$. If $C(\eta,f) = |2\eta-1|$, then

$$\begin{aligned} C(\eta,f) - C^*(\eta) = |2\eta-1| &= |2\eta-1| \cdot \frac{\phi(0) - C_\phi^*(\eta)}{\phi(0) - C_\phi^*(\eta)} \\ &\leq \frac{1}{\phi(0) - C_\phi^*(\eta)} \left(\left(|2\eta-1|\phi(0) + (1-|2\eta-1|)C_\phi^*(\eta)\right) - C_\phi^*(\eta)\right) \end{aligned} \tag{42}$$

At the same time, because $|\eta - 1/2| \geq \alpha$ $\mathbb{P}$-a.e. Lemma 12 implies that $C_\phi^*(\eta) \leq C_\phi^*(1/2-\alpha)$ $\mathbb{P}$-a.e. Furthermore, the relation $2\min(\eta, 1-\eta) = 1 - |1-2\eta|$ together with (41) shows that

$$|2\eta-1|\phi(0) + (1-|2\eta-1|)C_\phi^*(\eta) \leq C_\phi^-(\eta).$$

Therefore, (42) is bounded above by

$$\leq \frac{1}{\phi(0) - C_\phi^*\left(\frac{1}{2}-\alpha\right)} \left(C_\phi^-(\eta) - C_\phi^*(\eta)\right) \leq \frac{1}{\phi(0) - C_\phi^*\left(\frac{1}{2}-\alpha\right)} \left(C_\phi(\eta,f) - C_\phi^*(\eta)\right). \tag{43}$$

The last equality follows from the supposition $C(\eta,f) - C^*(\eta) = |2\eta-1|$, as it implies $(2\eta-1)f \leq 0$, and thus $C_\phi(\eta,f) \geq C_\phi^-(\eta)$. Consequently, (43) implies (8).

Integrating (8) with respect to $\mathbb{P}$ then produces the surrogate bound (9).

□

C Proof of Lemma 1

Proof of Lemma 1. If $\mathbf{x}' \in \overline{B_\epsilon(\mathbf{x})}$ then $S_\epsilon(g)(\mathbf{x}) \geq g(\mathbf{x}')$. Thus if γ is a coupling between $\mathbb{Q}$ and $\mathbb{Q}'$ supported on Δ_ϵ, then $S_\epsilon(g)(\mathbf{x}) \geq g(\mathbf{x}')$ γ-a.e. Integrating this inequality in γ produces

$$\int S_\epsilon(g) d\mathbb{Q} \geq \int g d\mathbb{Q}'.$$

Taking the supreumum over all $\mathbb{Q} \in \mathcal{B}_\epsilon^\infty(\mathbb{Q})$ then proves the result. □

D Proof of Item 1), Theorem 8

We will work with an alternative primal problem from Frank & Niles-Weed (2024b) that will make it easier to study the dual. Consider minimizing

$$\Theta(h_0, h_1) = \int S_\epsilon(h_1) d\mathbb{P}_1 + \int S_\epsilon(h_0) d\mathbb{P}_0$$

over the convex set

$$S_\phi = \begin{Bmatrix} (h_0, h_1) \colon h_0, h_1 \colon K^\epsilon \to \overline{\mathbb{R}} \text{ Borel, } 0 \leq h_0, h_1 \text{ and for} \\ \text{all } \mathbf{x} \in \mathbb{R}^d,\ \eta \in [0,1],\ \eta h_1(\mathbf{x}) + (1-\eta) h_0(\mathbf{x}) \geq C_\phi^*(\eta) \end{Bmatrix} \tag{44}$$

Then strong duality holds with Θ in place of R_ϕ^ϵ. Furthermore, there exist minimizers over the set of $\overline{\mathbb{R}}$-valued functions, where $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$.

Theorem 13. *Define $\bar{R}_\phi$ as in (13).*

$$\inf_{(h_0,h_1) \in S_\phi} \Theta(h_0, h_1) = \sup_{\substack{\mathbb{P}_0' \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_0) \\ \mathbb{P}_1' \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_1)}} \bar{R}_\phi(\mathbb{P}_0', \mathbb{P}_1')$$

Furthermore, the infimum is attained at some $\overline{\mathbb{R}}$-valued h_0^*, h_1^*.

See (Frank & Niles-Weed, 2024b, Lemma 14,Lemma 21) for a proof of this result. Theorem 7 already implies that the dual problem attains its supremum. Complimentary slackness conditions further characterize minimizers and maximizers.

Theorem 14 (Complementary Slackness)**.** *The pair (h_0^*, h_1^*) minimize Θ over S_ϕ and the measures $(\mathbb{P}_0^*, \mathbb{P}_1^*)$ maximize $\bar{R}_\phi$ over $\mathcal{B}_\epsilon^\infty(\mathbb{P}_0) \times \mathcal{B}_\epsilon^\infty(\mathbb{P}_1)$ iff the following two conditions hold:*

1)

$$\int S_\epsilon(h_1^*) d\mathbb{P}_1 = \int h_1^* d\mathbb{P}_1^* \quad \textit{and} \quad \int S_\epsilon(h_0^*) d\mathbb{P}_0 = \int h_0^* d\mathbb{P}_0^*$$

2)

$$\eta^* h_1^* + (1-\eta^*) h_0^* = C_\phi^*(\eta^*) \quad \mathbb{P}^*\text{-a.e.}$$

See Frank & Niles-Weed (2024b, Lemma 15) for a proof. Theorems 13 and 14 apply to the conditional risk $C^*(\eta)$ as $C^*(\eta) = C_\phi^*(\eta)$ for the hinge $\phi(\alpha) = \frac{1}{2}(1-\alpha)_+$.

We will use a characterization of consistency similar to Theorem 1 in the proof of Item 1), Theorem 8.

Theorem 15. *A loss function ϕ is consistent iff $C_\phi^*(\eta)$ has a strict maximum at $1/2$.*

Proof. If $C_\phi^*(1/2) = \phi(0)$, this statement is exactly Theorem 1. If $C_\phi^*(1/2) < \phi(0)$, Frank & Niles-Weed (2024a, Proposition 3) implies that ϕ is consistent. It remains to show that if $C_\phi^*(1/2) < \phi(0)$, then $C_\phi^*(\eta)$ has a strict maximum at $1/2$. As every sequence has a convergent subsequence in $\overline{\mathbb{R}}$, one can assume that $C_\phi(1/2, \cdot)$ has a minimizer α^* and $C_\phi^*(1/2) < \phi(0)$ implies $\alpha^* \neq 0$. Symmetry of $C_\phi(1/2, \cdot)$ implies that we can assume $\alpha^* > 0$, and thus $\phi(\alpha^*) \leq \phi(0)$ and $\phi(-\alpha^*) \geq \phi(0)$. The fact that $C_\phi^*(1/2, \alpha^*) < \phi(0)$ implies that in fact $\phi(\alpha^*) < \phi(0) \leq \phi(-\alpha^*)$. Next, observe that for any α,

$$C_\phi(\eta, \alpha) = \frac{1}{2}(\phi(\alpha) + \phi(-\alpha)) + (\eta - \frac{1}{2})(\phi(\alpha) - \phi(-\alpha))$$

Thus, one can bound $C_\phi^*(\eta)$ by

$$C_\phi^*(\eta) \leq C_\phi(\eta, \alpha^*) = \frac{1}{2}(\phi(\alpha^*) + \phi(-\alpha^*)) + (\eta - 1/2)(\phi(\alpha^*) - \phi(-\alpha^*)) = C_\phi^*(1/2) + (\eta - 1/2)(\phi(\alpha^*) - \phi(-\alpha^*))$$

Thus if $\eta > 1/2$, then $C^*_\phi(\eta) < C^*_\phi(1/2)$. Symmetry implies that $C^*_\phi(\eta) < C^*_\phi(1/2)$ for all η. Thus C^*_ϕ has a strict maximum at $1/2$.

□

Next, Theorem 14 implies that minimizers of Θ assume their suprema. This observation will make it easier to work with these functions.

Lemma 14. *If (h_0^*, h_1^*) minimizes Θ over S_ϕ, then the functions h_0^*, h_1^* assume their suprema $\mathbb{P}_0$-a.e. and $\mathbb{P}_1$-a.e. respectively*

Proof. We will show the statement for h_1^*, the argument for h_0^* is analogous. Let γ_1^* be the coupling between $\mathbb{P}_1$ and $\mathbb{P}_1^*$ that achieves the minimum W_∞ distance. Lemma 1 and Item 1) of Theorem 14 implies that

$$S_\epsilon(h_1)(\mathbf{x}) = h_1(\mathbf{x}') \quad \gamma_1^*\text{-a.e.}$$

and thus h_1^* assumes its maximum over closed ϵ-balls $\mathbb{P}_1^*$-a.e. □

Lemma 15. *If $(h_0^*, h_1^*) \in S_\phi$, then at any $\mathbf{x}$ either $h_1^*(\mathbf{x}) \geq C^*_\phi(\frac{1}{2})$ or $h_0^*(\mathbf{x}) > C^*_\phi(\frac{1}{2})$.*

Proof. If $(h_0^*, h_1^*) \in S_\phi$, then at any point $\mathbf{x}$,

$$\frac{1}{2}h_0^*(\mathbf{x}) + \frac{1}{2}h_1^*(\mathbf{x}) \geq C^*_\phi(\frac{1}{2}).$$

The inequality $h_0^*(\mathbf{x}) \leq C^*_\phi(1/2)$ implies $h_1^*(\mathbf{x}) \geq C^*_\phi(1/2)$. Thus either $h_0^*(\mathbf{x}) > C^*_\phi(1/2)$ or $h_1^*(\mathbf{x}) \geq C^*_\phi(1/2)$ at any point. □

Proof of Item 1) of Theorem 8 . Let $\phi_{\text{hinge}}(\alpha) = \frac{1}{2}(1-\alpha)_+$, then $C^*_{\phi_{\text{hinge}}}(\eta) = C^*(\eta)$.

Let (h_0^*, h_1^*) minimize Θ over S_ϕ and $(\mathbb{P}_0^*, \mathbb{P}_1^*)$ maximize $\bar{R}_\phi$ over $\mathcal{B}^\infty_\epsilon(\mathbb{P}_0) \times \mathcal{B}^\infty_\epsilon(\mathbb{P}_1)$. We will show that the functions defined by

$$\tilde{h}_1^*(\mathbf{x}) = \mathbf{1}_{h_1^*(\mathbf{x}) \geq C^*_\phi(\frac{1}{2})} \quad \tilde{h}_0^*(\mathbf{x}) = \mathbf{1}_{h_0^*(\mathbf{x}) > C^*_\phi(\frac{1}{2})}$$

maximize Θ over $S_{\phi_{\text{hinge}}}$ and $(\mathbb{P}_0^*, \mathbb{P}_1^*)$ maximize $\bar{R}_{\phi_{\text{hinge}}}$ by verifying the constraint $(\tilde{h}_0^*, \tilde{h}_1^*) \in S_{\phi_{\text{hinge}}}$ and the complimentary slackness conditions. The proof thus consists of three steps: verifying $(\tilde{h}_0^*, \tilde{h}_1^*) \in S_{\phi_{\text{hinge}}}$, and checking the two complementary slackness conditions in Theorem 14.

1) **Verifying the constraint defining $S_{\phi_{\text{hinge}}}$:** Observe that Lemma 15 implies that at any $\mathbf{x}$, at least one of $\tilde{h}_0^*(\mathbf{x})$ and $\tilde{h}_1^*(\mathbf{x})$ is 1, and thus

$$\eta h_1^*(\mathbf{x}) + (1-\eta)h_0^*(\mathbf{x}) \geq \min(\eta, 1-\eta) = C^*_{\phi_{\text{hinge}}}(\eta)$$

2) **Verifying Item 1) of Theorem 14:** Observe that Lemma 14 implies that $S_\epsilon(\mathbf{1}_{h_1^* \geq C^*_\phi(1/2)})(\mathbf{x}) = \mathbf{1}_{S_\epsilon(h_1^*)(\mathbf{x}) \geq C^*_\phi(1/2)}$ $\mathbb{P}_1^*$-a.e. Subsequently, the Item 1) of Theorem 14 implies that

$$S_\epsilon(\mathbf{1}_{h_1^* \geq C^*_\phi(1/2)})(\mathbf{x}) = \mathbf{1}_{h_1^*(\mathbf{x}') \geq C^*_\phi(1/2)} \quad \gamma_1^*\text{-a.e.},$$

verifying the first complimentary slackness condition for $\tilde{h}_1^*$. Analogous reasoning shows that

$$S_\epsilon(\mathbf{1}_{h_0^* > C^*_\phi(1/2)})(\mathbf{x}) = \mathbf{1}_{h_0^*(\mathbf{x}') > C^*_\phi(1/2)} \quad \gamma_0^*\text{-a.e.}$$

3) **Verifying Item 2) of Theorem 14:** Theorem 14 implies that $\eta^* h_1^*(\mathbf{x}') + (1-\eta^*)h_0^*(\mathbf{x}') = C^*_\phi(\eta^*) \leq C^*_\phi(1/2)$, and thus Lemma 15 implies that *exactly* one of $h_1^*(\mathbf{x}')$ and $h_0^*(\mathbf{x}')$ equals 1 and the other equals 0. We'll consider the cases $\eta^*(\mathbf{x}') < 1/2$, $\eta^*(\mathbf{x}') = 1/2$, and $\eta^*(\mathbf{x}') > 1/2$ separately. In these three separate cases, we will explicitly use the formula $C^*_{\phi_{\text{hinge}}}(\eta) = \min(\eta, 1-\eta)$.

When $\eta^*(\mathbf{x}') = 1/2$: As exactly one of $h_0^*(\mathbf{x}')$ and $h_1^*(\mathbf{x}')$ is 1:

$$\frac{1}{2}\tilde{h}_0^*(\mathbf{x}') + \frac{1}{2}\tilde{h}_1^*(\mathbf{x}') = \frac{1}{2} = C^*_{\phi_{\text{hinge}}}\left(\frac{1}{2}\right)$$

When $\eta^*(\mathbf{x}') < 1/2$: Observe that if $h_0^*(\mathbf{x}') > h_1^*(\mathbf{x}')$, then

$$\eta^* h_0^*(\mathbf{x}') + (1-\eta^*)h_1^*(\mathbf{x}') < \eta^* h_1^*(\mathbf{x}') + (1-\eta^*)h_0^*(\mathbf{x}') = C_\phi^*(\eta^*),$$

which would violate the constraint on S_ϕ. Therefore, $h_0^*(\mathbf{x}') \leq h_1^*(\mathbf{x}')$. Next, Theorem 15 implies that $\eta^* h_1^*(\mathbf{x}') + (1-\eta^*)h_0^*(\mathbf{x}') = C_\phi^*(\eta^*) < C_\phi^*(1/2)$. These two statements together with Lemma 15 imply that $h_0^*(\mathbf{x}') < C_\phi^*(1/2)$ and $h_1^*(\mathbf{x}') \geq C_\phi^*(1/2)$. However, $h_0^*(\mathbf{x}') = C_\phi^*(1/2)$ would still violate $\eta^* h_1^*(\mathbf{x}') + (1-\eta^*)h_0^*(\mathbf{x}') < C_\phi^*(1/2)$ and therefore, $h_0^*(\mathbf{x}') < C_\phi^*(1/2)$. Therefore,

$$\eta^* \tilde{h}_1^* + (1-\eta^*)\tilde{h}_0^* = \eta^* = C^*_{\phi_{\text{hinge}}}(\eta^*)$$

When $\eta^*(\mathbf{x}') > 1/2$: Argument is analogous to the previous case.

□

E Proof of Lemma 7

We define the *concave conjugate* of a function h as

$$h_*(y) = \inf_{x \in \text{dom}(h)} yx - h(x)$$

Recall that $\text{conc}(h)$ as defined in (15) is the biconjugate h_{**}. Consequently, $\text{conc}(h)$ can be expressed as

$$\text{conc}(h)(x) = \inf\{\ell(x) : \ell \text{ linear, and } \ell \geq h \text{ on } \text{dom}(h)\} \tag{45}$$

Lemma 7 is a consequence of the properties of concave conjugates.

Lemma 16. *Let $h : [a,b] \to \mathbb{R}$ be a non-decreasing function. Then $\text{conc}(h)$ is non-decreasing as well.*

Proof. We will argue that if h is non-decreasing, then it suffices to consider the infimum in (45) over non-decreasing linear functions. Observe that if ℓ is a decreasing linear function with $\ell(x) \geq h(x)$ then the constant function $\ell(b)$ satisfies

$$\ell(x) \geq \ell(b) \geq h(b) \geq h(x)$$

for any $x \in [a,b]$. Therefore,

$$\text{conc}(h)(x) = \inf\{\ell(x) : \ell \text{ linear, non-decreasing, and } \ell \geq h\}$$

□

Lemma 17. *Let $h : [0,b] \to \mathbb{R}$ be a non-decreasing function that is right-continuous at zero with $h(0) = 0$. Then $\sup_y h_*(y) = 0$. Furthermore, there is a sequence y_n with $y_n \to \infty$ and $\lim_{n\to\infty} h_*(y_n) = 0$.*

Proof. First, notice that

$$h_*(y) = \inf_{x \in [0,b]} yx - h(x) \leq y \cdot 0 - h(0) = 0 \tag{46}$$

for any $y \in \mathbb{R}$. It remains to show a sequence y_n for which $\lim_{n\to\infty} h_*(y_n) = 0$.

We will argue than any sequence y_n with

$$y_n > nh(b) \geq \sup_{x \in [1/n, b]} \frac{h(x)}{x} \tag{47}$$

satisfies this property.

If $x \in [1/n, b]$ and y_n satisfies (47) then

$$xy_n - h(x) = x\left(y_n - \frac{h(x)}{x}\right) > 0$$

and thus (46) implies that

$$h_*(y_n) = \inf_{x \in [0,1/n)} xy_n - h(x)$$

The monononicity of h then implies that

$$h_*(y_n) \geq -h(1/n)$$

and

$$\lim_{n\to\infty} h_*(y_n) \geq 0$$

because h is right-continuous at zero. This relation together with (46) implies the result.

□

Proof of Lemma 7. Lemma 16 implies that $\operatorname{conc}(h)$ is non-decreasing. Standard results in convex analysis imply that $\operatorname{conc}(h)$ is continuous on $(0, 1/2)$ (Hiriart-Urruty & Lemaréchal, 2001, Lemma 3.1.1) and upper semi-continuous on $[0, 1/2]$ (Hiriart-Urruty & Lemaréchal, 2001, Theorem 1.3.5). Thus monotonicity implies that for all $x \in [0, 1/2]$, $\operatorname{conc}(h)(x) \leq \operatorname{conc}(h)(1/2)$ and thus $\lim_{x\to 1/2} \operatorname{conc}(h)(x) \leq \operatorname{conc}(h)(1/2)$. We will show the opposite inequality, implying that $\operatorname{conc} h$ is continuous at $1/2$.

First, as the constant function $h(1/2)$ is an upper bound on h, one can conclude that $\operatorname{conc}(h)(1/2) = h(1/2) = 1$. Next, recall that $\operatorname{conc}(h)$ can be expressed as an infimum of linear functions as in (45). If $\ell \geq h$, then $\ell(0) \geq 0$ and $\ell(1/2) \geq 1$. Therefore,

$$\ell(\tfrac{1}{2} - \delta) = \ell((1 - 2\delta) \cdot \tfrac{1}{2} + 2\delta \cdot 0) = (1 - 2\delta)\ell(\tfrac{1}{2}) + 2\delta\ell(0) \geq 1 - 2\delta.$$

Therefore, the representation (45) implies that $\operatorname{conc}(h)(1/2 - \delta) \geq 1 - 2\delta$. Taking $\delta \to 0$ proves that $\lim_{x\to 1/2} \operatorname{conc}(h)(x) \geq 1$. Thus, $\operatorname{conc}(h)$ is continuous at $1/2$, if viewed as a function on $[0, 1/2]$.

Next, Lemma 17 implies that $h_{**}(0) = 0$:

$$h_{**}(0) = \inf_{y\in\mathbb{R}} -h_*(y) = -\sup_{y\in\mathbb{R}} h_*(y) = 0.$$

Finally, it remains to show that h_{**} is continuous at 0. The monotonicity of h_{**} implies that $\lim_{y\to 0^+} h_{**}(y) = \inf_{y\in(0,1/2]} h_{**}(y)$ and consequently

$$\lim_{y\to 0^+} h_{**}(y) = \inf_{y\in(0,1/2]} \inf_{x\in\mathbb{R}} yx - h_*(x) = \inf_{x\in\mathbb{R}} \inf_{y\in(0,1/2]} yx - h_*(x) = \inf_{x\in\mathbb{R}} -h_*(x) + \begin{cases} 0 & \text{if } x \geq 0 \\ \frac{x}{2} & \text{if } x < 0 \end{cases}$$

$$= \min\left(\inf_{x\geq 0} -h_*(x), \inf_{x<0} \frac{x}{2} - h_*(x)\right) \tag{48}$$

However, Lemma 17 implies that

$$\inf_{x\geq 0} -h_*(x) = \inf_{x\in\mathbb{R}} -h_*(x) = 0 \tag{49}$$

Notice that if $x \leq 0$,

$$h_*(x) = \inf_{z\in[0,1/2]} xz - h(z) = \frac{x}{2} - h\left(\frac{1}{2}\right) = \frac{x}{2} - 1 \tag{50}$$

Consequently, (49) and (50) implies that (48) evaluates to 0.

□

F Proof of Proposition 2

A modified version of Jensen's inequality will be used at several points in the proof of Proposition 2.

Lemma 18. *Let G be a concave function with $G(0) = 0$ and let ν be a measure with $\nu(\mathbb{R}^d) \leq 1$. Then*

$$\int G(f)d\nu \leq G\left(\int f d\nu\right)$$

Proof. The inequality trivially holds if $\nu(\mathbb{R}^d) = 0$, so we assume $\nu(\mathbb{R}^d) > 0$. Jensen's inequality implies that

$$\int G(f)d\nu = \nu(\mathbb{R}^d)\left(\frac{1}{\nu(\mathbb{R}^d)}\int G(f)d\nu\right) \leq \nu(\mathbb{R}^d)G\left(\frac{1}{\nu(\mathbb{R}^d)}\int f d\nu\right).$$

As $G(0) = 0$, concavity implies that

$$\nu(\mathbb{R}^d)G\left(\frac{1}{\nu(\mathbb{R}^d)}\int f d\nu\right) = \nu(\mathbb{R}^d)G\left(\frac{1}{\nu(\mathbb{R}^d)}\int f d\nu\right) + (1-\nu(\mathbb{R}^d)G(0) \leq G\left(\int f d\nu\right)$$

□

To facilitate the application of Jensen's inequality, the proof will be carried out using integrated quantities. Let $\mathbb{P}_0^*, \mathbb{P}_1^*$ be any maximizers of $\bar{R}_\phi$, which also maximize $\bar{R}$ by Theorem 8. Set $\mathbb{P}^* = \mathbb{P}_0^* + \mathbb{P}_1^*$, $\eta^* = d\mathbb{P}_1^*/d\mathbb{P}^*$. Define

Proof of Proposition 2. Let γ_0^*, γ_1^* be the couplings between $\mathbb{P}_0$, $\mathbb{P}_0^*$ and $\mathbb{P}_1$, $\mathbb{P}_1^*$ respectively that achieve the infimum in (11). Define $I_1(f)$, $I_0(f)$, $I_1^\phi(f)$, and $I_0^\phi(f)$ by

$$I_1(f) = \int i_1(f)d\gamma_1^*, \quad I_1^\phi(f) = \int i_1^\phi(f)d\gamma_1^*, \quad I_0(f) = \int i_0(f)d\gamma_0^*, \quad I_0^\phi(f) = \int i_0^\phi(f)d\gamma_0^*.$$

We will prove $$I_0(f) \leq \frac{1}{2}\tilde{\Phi}\big(2I_0^\phi(f)\big). \tag{51}$$ $$I_1(f) \leq \frac{1}{2}\tilde{\Phi}\big(2I_1^\phi(f)\big) \tag{52}$$

The concavity of $\tilde{\Phi}$ then implies that

$$R^\epsilon(f) - R_*^\epsilon = I_1(f) + I_0(f) \leq \frac{1}{2}\tilde{\Phi}\big(2I_1^\phi(f)\big) + \frac{1}{2}\tilde{\Phi}\big(2I_0^\phi(f)\big) \leq \tilde{\Phi}\Big(\frac{1}{2}2I_1^\phi(f) + \frac{1}{2}2I_0^\phi(f)\Big) = \tilde{\Phi}\big(R_\phi^\epsilon(f) - R_{\phi,*}^\epsilon\big).$$

We will prove (52), the argument for (51) is analogous. Next, let γ_1^* be the coupling between $\mathbb{P}_1$ and $\mathbb{P}_1^*$ supported on Δ_ϵ. The assumption on Φ implies that

$$C(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C^*(\eta^*(\mathbf{x}')) \leq \Phi\big(C_\phi(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C_\phi^*(\eta^*(\mathbf{x}'))\big) \tag{53}$$

and consequently,

$$\int C(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C^*(\eta^*(\mathbf{x}'))d\gamma_1^* \leq \Phi\left(\int C_\phi(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C_\phi^*(\eta^*(\mathbf{x}'))d\gamma_1^*\right) \leq \Phi(I_1^\phi(f)). \tag{54}$$

To bound the term $S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0}$, we consider two different cases for $(\mathbf{x}, \mathbf{x}')$. Define the sets D_1, E_1 as in (25), (26). We will show that if T_1 is any of the sets D_1, E_1, then

$$\begin{aligned}
&\int_{T_1} S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0} d\gamma_1^* \\
&\leq \left(\int \frac{1}{G\Big(\big(\phi(0) - C_\phi^*(\eta^*(\mathbf{x}'))\big)/2\Big)} d\gamma_1^*\right)^{\frac{1}{2}} G\left(\int_{T_1} \big((S_\epsilon(\phi\circ f)(\mathbf{x}) - \phi(f(\mathbf{x}'))\big) + \big(C_\phi(\eta^*(\mathbf{x}'), f(\mathbf{x}')) - C_\phi^*(\eta^*(\mathbf{x}'))\big)\, d\gamma_1^*\right)^{\frac{1}{2}}
\end{aligned} \tag{55}$$

Thus because G is concave and non-decreasing, the composition $\sqrt{G}$ is as well. Thus summing the inequality (55) over $T_1 \in \{D_1, E_1\}$ results in

$$\int S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0} d\gamma_1^* \leq 2 \left(\int \frac{1}{G\Big(\phi(0) - C_\phi^*(\eta^*(\mathbf{x}'))\Big)} d\mathbb{P}^* \right)^{\frac{1}{2}} G\left(\frac{1}{2} I_1^\phi(f)\right)^{\frac{1}{2}} \tag{56}$$

Summing (54) and (56) results in (52).

It remains to show the inequality (55) for the two sets D_1, E_1.

A) **On the set D_1:**

If $S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) = \mathbf{1}_{f(\mathbf{x}')\leq 0}$, then $\int_{D_1} S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0} d\gamma_1^* = 0$ while the left-hand side of (55) is non-negative by Lemma 1, which implies (55) for $T_1 = D_1$.

B) **On the set E_1:**

Lemma 1 then implies that $S_\epsilon(\phi \circ f)(\mathbf{x}) - \phi(f(\mathbf{x}')) \geq 0$ γ_1^*-a.e. and thus Lemma 5 implies

$$S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0} = 1 = \frac{\sqrt{G\left(\phi(0) - C_\phi^*(\eta^*(\mathbf{x}'))\right)}}{\sqrt{G\left(\phi(0) - C_\phi^*(\eta^*(\mathbf{x}'))\right)}} \leq \frac{\sqrt{G\big(S_\epsilon(\phi \circ f)(\mathbf{x}) - \phi(f(\mathbf{x}'))\big)}}{\sqrt{G\left(\phi(0) - C_\phi^*(\eta^*(\mathbf{x}'))\right)}} \quad \gamma_1^*\text{-a.e.} \tag{57}$$

Now the Cauchy-Schwartz inequality and Jensen's inequality(Lemma 18) imply

$$\begin{aligned}
&\int_{E_1} S_\epsilon(\mathbf{1}_{f\leq 0})(\mathbf{x}) - \mathbf{1}_{f(\mathbf{x}')\leq 0} d\gamma_1^* \\
&\leq \left(\int_{E_1} \frac{1}{G\left(\phi(0) - C_\phi^*(\eta^*(\mathbf{x}'))\right)} d\gamma_1^* \right)^{\frac{1}{2}} \left(\int_{E_1} G\big(S_\epsilon(\phi \circ f)(\mathbf{x}) - \phi(f(\mathbf{x}'))\big) d\gamma_1^* \right)^{\frac{1}{2}} \\
&\leq \left(\int \frac{1}{G\left(\phi(0) - C_\phi^*(\eta^*(\mathbf{x}'))\right)} d\gamma_1^* \right)^{\frac{1}{2}} G\left(\int_{E_1} S_\epsilon(\phi \circ f)(\mathbf{x}) - \phi(f(\mathbf{x}')) d\gamma_1^* \right)^{\frac{1}{2}},
\end{aligned} \tag{58}$$

which implies (55).

□

G Technical Integral Lemmas

In this section, we require several technical facts about Riemann–Stieltjes integrals, which we briefly review here.

Let $g : \mathbb{R} \to \mathbb{R}$, $h : \mathbb{R} \to \mathbb{R}$ be functions and let $P = \{z_0, z_1, \ldots, z_K\}$ be a partition of an interval I. Then the *lower* and *upper* sums with respect to g, h, P are defined as

$$L(g, h, P) = \sum_{k=0}^{K-1} \inf_{z \in [z_k, z_{k+1}]} g(z)(h(z_{k+1}) - h(z_k)), \quad U(g, h, P) = \sum_{k=0}^{K-1} \sup_{z \in [z_k, z_{k+1}]} g(z)(h(z_{k+1}) - h(z_k))$$

respectively. When g is non-increasing, these simplify as $\inf_{z\in[z_k,z_{k+1}]} g(z) = g(z_{k+1})$ and $\sup_{z\in[z_k,z_{k+1}]} g(z) = g(z_k)$.

Riemann–Stieltjes integral $\int_I g dh$ can be approximated by upper and lower sums, much as in the classical Riemann case. The following result records the relevant approximation property:

Proposition 3. *Let $\int_I g dh$ be a Riemann-Stieltjes integral. If g is continuous and h is monotone, then the integral exists. Moreover, for any partition P, $L(g,h,P) \leq \int_I g dh \leq U(g,h,P)$. In addition, for any $\delta > 0$, there exists a partition P for which $U(g,h,P) - \delta \leq \int_I g dh \leq L(g,h,P) + \delta$.*

For details, see Apostol (1974, Theorem 7.17) or Theorem 2.24 of Wheeden & Zygmund (1977) for the existence statement and Apostol (1974, Theorem 7.27) for a discussion of upper and lower integrals.

G.1 The Lebesgue and Riemann–Stieltjes integral of an increasing function

The goal of this section is to prove (37), or namely:

Proposition 4. *Let f be a non-increasing, non-negative, continuous function on an interval $[a,b]$ and let $\mathbb{Q}$ be a finite positive measure. Let z be a random variable distributed according to $\mathbb{Q}$ and define $h(\alpha) = \mathbb{Q}(z \leq \alpha)$. Then*

$$\int_{(a,b]} f(z) d\mathbb{Q}(z) = \int_a^b f(\alpha) dh(\alpha)$$

where the integral on the left is defined as the Lebesgue integral in terms of the measure $\mathbb{Q}$ while the integral on the right is defined as a Riemann–Stieltjes integral.

Proof. Recall that when f is monotonic, the Riemann-Stieltjes integral is the value of the limits

$$\int f dh = \lim_{\Delta\alpha_i \to 0} \sum_{i=0}^{I-1} f(\alpha_i)(h(\alpha_{i+1}) - h(\alpha_i)) = \lim_{\Delta\alpha_i \to 0} \sum_{i=0}^{I-1} f(\alpha_{i+1})(h(\alpha_{i+1}) - h(\alpha_i)), \tag{59}$$

where these limits are evaluated as the size of the partition $\Delta\alpha_i = \alpha_{i+1} - \alpha_i$ approaches 0 (Apostol, 1974, Exercise 7.3, Theorem 7.19), while the Lebesgue integral $\int f d\mathbb{Q}$ is defined as

$$\int f d\mathbb{Q} = \sup \left\{ \int g d\mathbb{Q} : g \leq f, g \text{ simple function}, \right\}.$$

The limits in (59) are upper and lower sums because f is monotonic, and thus by Proposition 3, for any $\delta > 0$, one can choose a partition $\{\alpha_i\}_{i=0}^I$ for which each of the sums in (59) is within δ of $\int f dh$.

Next, consider two simple functions g_1, g_2 defined according to

$$g_1(z) = \sum_{i=0}^{I-1} f(\alpha_{i+1}) \chi_{z \in (\alpha_i, \alpha_{i+1}]}, \quad g_2(z) = \sum_{i=0}^{I-1} f(\alpha_i) \chi_{z \in (\alpha_i, \alpha_{i+1}]}.$$

By construction, $g_1(x) \leq f(x) \leq g_2(x)$ for all $x \in (a,b]$. Moreover, since $f(\alpha_i) - f(\alpha_{i+1}) < \delta$, it follows that $f(x) \leq g_2(x) + \delta$ when $x \in (a,b]$. Now applying the definition of the integral of a simple function, we obtain:

$$\begin{aligned} \int f dh - \delta \leq \sum_{i=0}^{I-1} f(\alpha_{i+1}) \big(h(\alpha_{i+1}) - h(\alpha_i)\big) = \int_{(a,b]} g_1 d\mathbb{Q} \leq \int_{(a,b]} f d\mathbb{Q} \leq \int_{(a,b]} g_2 d\mathbb{Q} \\ = \sum_{i=0}^{I-1} f(\alpha_i) \big(h(\alpha_{i+1}) - h(\alpha_i)\big) \leq \int f dh + \delta \end{aligned}$$

As δ is arbitrary, it follows that $\int f dh = \int f d\mathbb{Q}$. □

Notice that because $H(0) = 0$, the integral in the right-hand side of (37) is technically an improper integral. Thus to show (37), one can conclude that

$$\int_{z \in (\delta, 1/2]} \frac{1}{H(z)} d\mathbb{Q}(z) = \int_\delta^{1/2} \frac{1}{H(\alpha)} dh(\alpha)$$

from Proposition 4 and then take the limit $\delta \to 0$.

G.2 Proof of the last equality in (38)

The goal of this appendix is to prove the following inequality:

Lemma 19. *Let $h : [0, 1/2] \to [0, 1]$ be an increasing and right-continuous function with $h(0) = 0$ and $h(1/2) \leq 1$. Let H be any continuous function with $H \geq h$ and let $r \in [0, 1)$. Then one can bound the Riemann–Stieltjes integral $\int 1/H(z)^r dh$ by*

$$\int_0^{1/2} \frac{1}{H(z)^r} dh \leq \frac{1}{1-r}$$

Proof. Let $\delta > 0$, then one can pick a partition $P = \{z_0 = 0, z_1, \ldots, z_K = 1/2\}$ for which $\int_0^{1/2} H^{-r} dh \leq L(H^{-r}, h, P) + \delta$. As H^{-r} is non-increasing, $L(H^{-r}, h, P) = \sum_{k=0}^{K-1} H^{-r}(z_{k+1})(h(z_{k+1}) - h(z_k))$. Therefore, if we define $a_k = h(z_k)$, then

$$\begin{aligned} \int_0^{1/2} H^{-r} dh &\leq \sum_{k=0}^{K-1} H^{-r}(z_{k+1})(h(z_{k+1}) - h(z_k)) + \delta \leq \sum_{k=1}^{K-1} h^{-r}(z_{k+1})(h(z_{k+1}) - h(z_k)) + \delta \\ &= \sum_{k=0}^{K-1} a_{k+1}^{-r}(a_{k+1} - a_k) + \delta \end{aligned} \tag{60}$$

Because the function $y \mapsto y^{-r}$ is decreasing in y, one can bound $a_{k+1}^{-r}(a_{k+1} - a_k) \leq \int_{a_k}^{a_{k+1}} y^{-r} dy$ and consequently the sum in (60) is bounded above as

$$\sum_{k=0}^{K-1} \int_{a_k}^{a_{k+1}} y^{-r} dy = \int_0^{h(1/2)} y^{-r} dy \leq \int_0^1 y^{-r} dy = \frac{1}{1-r}$$

Therefore $\int_0^{1/2} H^{-r} dh \leq 1/(1-r) + \delta$. The result follows as $\delta > 0$ is arbitrary. □

H Optimizing the Bound of Lemma 8 over r

Proof of Theorem 11. Let

$$f(r) = \frac{1}{1-r} a^r$$

Then

$$f'(r) = \frac{1}{(1-r)^2} a^r + \frac{1}{1-r} \ln a a^r$$

solving $f'(r^*) = 0$ produces $r^* = 1 + \frac{1}{\ln a}$, and

$$f\left(1 + \frac{1}{\ln a}\right) = -\ln a a^{1+\frac{1}{\ln a}} = -ea \ln a$$

One can verify that this point is a minimum via the second derivative test:

$$f'(r) = \left(\frac{1}{1-r} + \ln a\right) f(r)$$

and thus

$$f''(r) = \left(\frac{1}{1-r} + \ln a\right) f'(r) + \frac{1}{(1-r)^2} f(r).$$

Consequently, $f''(r^*) = \ln(a)^2 f(1 + \frac{1}{\ln a}) > 0$.

However, the point r^* is in the interval $[0, 1]$ only when $a \in [0, e^{-1}]$. When $a > e^{-1}$, f is minimized over $[0, 1]$ at $r = 0$. Because r^* is a minimizer when $a \in [0, e^{-1}]$, one can bound $f(0) \geq f(r^*)$ over this set and thus

$$f(r) \leq \min(1, -ea \ln a)$$

□

I Deferred proofs from Section 6

I.1 Existence of Minimizers and Complementary slackness

The existence and complimentary slackness theorems of Appendix D extend to R^ϵ_ϕ. Observe that minimizers of R_ϕ may assume values in $\overline{\mathbb{R}}$; for example, with the exponential loss $\phi(\alpha) = e^{-\alpha}$ and the distribution defined by $\eta(\mathbf{x}) \equiv 1$, the unique minimizer of R_ϕ is $+\infty$. Just as in the non-adversarial scenario, R^ϵ_ϕ may fail to attain its infimum over $\mathbb{R}$-valued functions. Nevertheless, Frank & Niles-Weed (2024a, Lemma 8) and Frank (2025, Theorem 6) guarantee the existence of a minimizer over $\overline{\mathbb{R}}$-valued functions.

Theorem 16. *Let ϕ satisfy Assumption 1. Then*

$$\inf_{f\ \mathbb{R}\text{-valued}} R^\epsilon_\phi(f) = \inf_{f\ \overline{\mathbb{R}}\text{-valued}} R^\epsilon_\phi(f).$$

Furthermore, equality is attained at some Borel measurable, $\overline{\mathbb{R}}$-valued function f^*.

Moreover, Theorem 7 of Frank & Niles-Weed (2024b) describes two conditions that characterize minimizers of R^ϵ_ϕ and maximizers $\bar{R}_\phi$.

Theorem 17 (Complementary Slackness)**.** *The function f^* minimizes R^ϵ_ϕ and the measures $(\mathbb{P}^*_0, \mathbb{P}^*_1)$ maximize $\bar{R}_\phi$ over $\mathcal{B}^\infty_\epsilon(\mathbb{P}_0) \times \mathcal{B}^\infty_\epsilon(\mathbb{P}_1)$ iff the following two conditions hold:*

1)

$$\int S_\epsilon(\phi(f^*))d\mathbb{P}_1 = \int \phi(f^*)d\mathbb{P}^*_1 \quad \textit{and} \quad \int S_\epsilon(\phi(-f^*))d\mathbb{P}_0 = \int \phi(-f^*)d\mathbb{P}^*_0$$

2)

$$C_\phi(\eta^*, f^*) = C^*_\phi(\eta^*) \quad \mathbb{P}^*\text{-a.e.}$$

I.2 Proof of Lemmas 9 and 10

As a preliminary step, we establish that if $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ are induced by transport maps, then these maps determine the locations of maximizers of $\phi \circ f$ and $\phi \circ -f$.

Lemma 20. *Let $\mathbb{P}^*_0, \mathbb{P}^*_1$ be maximizers of $\bar{R}_\phi$ induced by the transport maps T_0, T_1 satisfying $\|T_0(\mathbf{x}) - \mathbf{x}\| \le \epsilon$, $\|T_1(\mathbf{x}) - \mathbf{x}\| \le \epsilon$. Then any minimizer f^* of R^ϵ_ϕ satisfies*

$$S_\epsilon(\phi \circ -f^*)(\mathbf{x}) = \phi(-f^*(T_0(\mathbf{x})) \quad \mathbb{P}_0\text{-a.e.} \qquad (61)$$

$$S_\epsilon(\phi \circ f^*)(\mathbf{x}) = \phi(f^*(T_1(\mathbf{x})) \quad \mathbb{P}_1\text{-a.e.} \qquad (62)$$

Proof. We show (62), the argument for (61) is analogous.

Let f^* minimize R^ϵ_ϕ; such a function exists by Theorem 16. The complementary slackness condition Item 1) in Theorem 17 yields

$$\int S_\epsilon(\phi \circ f^*)(\mathbf{x})d\mathbb{P}_1 = \int \phi \circ f^*(\mathbf{x}')d\mathbb{P}^*_1 = \int \phi(f^*(T_1(\mathbf{x})))d\mathbb{P}$$

As the relation $\|T_1(\mathbf{x}) - \mathbf{x}\| \le \epsilon$ implies $S_\epsilon(\phi \circ f^*)(\mathbf{x}) \ge \phi(f^*(T_1(\mathbf{x})))$ one can conclude (62). □

Next, we verify strong duality for these restricted measures, utilizing the notation defined in Equations (39) and (40). The statement below implies Lemma 9.

Lemma 21. *Let $\mathbb{P}_0, \mathbb{P}_1, \mathbb{P}^*_0, \mathbb{P}^*_1, T_0, T_1, U_0, U_1$ and Q be as in Lemma 9, and let f^* minimize $R^\epsilon_\phi(\,\cdot\,; \mathbb{P}_0, \mathbb{P}_1)$. Then*

$$\mathbb{P}^*_0|_Q \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_0|_{U_0}), \quad \mathbb{P}^*_1|_Q \in \mathcal{B}^\infty_\epsilon(\mathbb{P}_1|_{U_1}) \qquad (63)$$
$$R^\epsilon_\phi(f^*; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) = \bar{R}_\phi(\mathbb{P}^*_0|_Q, \mathbb{P}^*_1|_Q).$$

Consequently f^* minimizes $R^\epsilon_\phi(\,\cdot\,; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1})$ while $\mathbb{P}^*_0|_Q$, $\mathbb{P}^*_1|_Q$ maximize $\bar{R}_\phi$ over $\mathcal{B}^\infty_\epsilon(\mathbb{P}_0|_{U_0}) \times \mathcal{B}^\infty_\epsilon(\mathbb{P}_1|_{U_1})$.

Proof. By the definitions of Q, U_0, U_1

$$\mathbb{P}_0^*|_Q = \mathbb{P}_0|_{U_0} \sharp T_0, \quad \mathbb{P}_1^*|_Q = \mathbb{P}_1|_{U_1} \sharp T_1.$$

The relations $\|T_0(\mathbf{x}) - \mathbf{x}\| \leq \epsilon$, $\|T_1(\mathbf{x}) - \mathbf{x}\| \leq \epsilon$ imply (63). Next, let $\tilde{\eta} = d\mathbb{P}_1^*|_Q / d(\mathbb{P}_1^*|_Q + \mathbb{P}_0^*|_Q)$. On the set Q,

$$\tilde{\eta} = \eta^* \quad \mathbb{P}^*\text{-a.e.} \tag{64}$$

Next, let f^* be a minimizer of $R_\phi^\epsilon(\cdot; \mathbb{P}_0, \mathbb{P}_1)$. Lemma 20 and the definitions of T_0, T_1 imply that

$$R_\phi^\epsilon(f^*; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) = \int S_\epsilon(\phi \circ f^*)(\mathbf{x}) \mathbf{1}_{U_1}(\mathbf{x}) d\mathbb{P}_1 + \int S_\epsilon(\phi \circ -f^*)(\mathbf{x}) \mathbf{1}_{U_0}(\mathbf{x}) d\mathbb{P}_0 \tag{65}$$

$$\int \phi(f^*(T_1(\mathbf{x}))) \mathbf{1}_Q(T_1(\mathbf{x})) d\mathbb{P}_1 + \int \phi(-f^*(T_0(\mathbf{x})) \mathbf{1}_Q(T_0(\mathbf{x})) d\mathbb{P}_0 \tag{66}$$

$$= \int \phi(f^*(\mathbf{x}')) \mathbf{1}_Q(\mathbf{x}') d\mathbb{P}_1 \sharp T_1 + \int \phi(-f^*(\mathbf{x}')) \mathbf{1}_Q(\mathbf{x}') d\mathbb{P}_0 \sharp T_0 \tag{67}$$

$$= \int C_\phi(\eta^*, f^*) \mathbf{1}_Q d\mathbb{P}^* \tag{68}$$

Since f^* minimizes $R_\phi^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1)$, the complimentary slackness condition Item 2) of Theorem 17 implies $C_\phi(\eta^*, f^*) = C_\phi^*(\eta^*)$. Equation 64 further implies $C_\phi^*(\eta^*) = C_\phi^*(\tilde{\eta})$ on Q $\mathbb{P}$-a.e. and therefore,

$$\int C_\phi(\eta^*, f^*) \mathbf{1}_Q d\mathbb{P}^* = \int C_\phi^*(\tilde{\eta}) d\mathbb{P}^*|_Q = \bar{R}(\mathbb{P}_0^*|_Q, \mathbb{P}_1^*|_Q)$$

□

These results show that restricted measures in the dual corresponds directly to restricted measures in the primal.

Proof of Lemma 9. Applying Theorem 7 to the restricted measures $\mathbb{P}_1|_{U_1}$, $\mathbb{P}_0|_{U_0}$ and invoking Lemma 21 yields the claim. □

Finally, one can conclude Lemma 10 by comparing $R_\phi^\epsilon(f^*; \mathbb{P}_0|_{T_0^{-1}(Q^C)}, \mathbb{P}_1|_{T_1^{-1}(Q^C)})$ and $R_\phi^\epsilon(f^*; \mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C})$.

Proof of Lemma 10. Observe that $U_0^C = T_0^{-1}(Q^C)$, $U_1^C = T_1^{-1}(Q^C)$ and let f^* be a minimizer of $R_\phi^\epsilon(\,\cdot\,; \mathbb{P}_0, \mathbb{P}_1)$. Then Lemmas 9 and 21 imply

$$\begin{aligned} R_{\phi,*}^\epsilon(\mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) &= R_\phi^\epsilon(f^*; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) = \int_Q C_\phi^*(\eta^*) d\mathbb{P}^* \\ R_{\phi,*}^\epsilon(\mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C}) &= R_\phi^\epsilon(f^*; \mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C}) = \int_{Q^C} C_\phi^*(\eta^*) d\mathbb{P}^* \end{aligned} \tag{69}$$

Summing these:

$$R_\phi^\epsilon(f^*; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) + R_\phi^\epsilon(f^*; \mathbb{P}_0|_{\tilde{U}_0}, \mathbb{P}_1|_{\tilde{U}_1}) = \int C_\phi^*(\eta^*) d\mathbb{P}^* = R_{\phi,*}^\epsilon(f^*; \mathbb{P}_0, \mathbb{P}_1)$$

□

I.3 Proof of Theorem 10

Proof of Theorem 10. Let $Q = \{\mathbf{x}' : |\eta^*(\mathbf{x}') - 1/2| < \alpha\}$. Then Lemma 9 applied to Q^C shows that $(\mathbb{P}_0^*|_{Q^C}, \mathbb{P}_1^*|_{Q^C})$ maximize $\bar{R}_\phi$ over $\mathcal{B}_\epsilon^\infty(\mathbb{P}_0|_{U_0}) \times \mathcal{B}_\epsilon^\infty(\mathbb{P}_1|_{U_1})$, with $U_0 = T_0^{-1}(Q^C)$ and $U_1 = T_1^{-1}(Q^C)$. Next, observe that simply scaling the inequality (14) shows that Theorem 9 applies even when $\mathbb{P}(\mathbb{R}^d) \leq 1$. Consequently, Theorem 9 and Lemma 10 imply that

$$\begin{aligned} R^\epsilon(f; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) - R_*^\epsilon(\mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) &\leq \frac{1}{\phi(0) - C_\phi^*(1/2-\alpha)} \left(R_\phi^\epsilon(f; \mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1}) - R_{\phi*}^\epsilon(\mathbb{P}_0|_{U_0}, \mathbb{P}_1|_{U_1})\right) \\ &\leq \frac{1}{\phi(0) - C_\phi^*(1/2-\alpha)} \left(R_\phi^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1) - R_{\phi*}^\epsilon(\mathbb{P}_0, \mathbb{P}_1)\right) \end{aligned}$$

Next, by Lemma 10, adding $R^\epsilon(f; \mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C}) - R_*^\epsilon(\mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C})$ to both sides of the inequality above results in

$$\begin{aligned} &R^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1) - R_*^\epsilon(\mathbb{P}_0, \mathbb{P}_1) \leq \\ &\frac{1}{\phi(0) - C_\phi^*(1/2-\alpha)} \left(R_\phi^\epsilon(f; \mathbb{P}_0, \mathbb{P}_1) - R_{\phi*}^\epsilon(\mathbb{P}_0|_{U_0}, \mathbb{P}_1)\right) + R^\epsilon(f; \mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C}) - \int_Q C^*(\eta^*) d\mathbb{P}^* \end{aligned}$$

The fact that $C^*(\eta^*) \geq 1/2 - \alpha$ on Q while $S_\epsilon(\mathbf{1}_{f \leq 0}) \leq 1$, $S_\epsilon(\mathbf{1}_{f>0}) \leq 1$ implies that $R^\epsilon(f; \mathbb{P}_0|_{U_0^C}, \mathbb{P}_1|_{U_1^C}) - \int_Q C^*(\eta^*) d\mathbb{P}^* \leq (\frac{1}{2} + \alpha)\mathbb{P}^*(|\eta - 1/2| < \alpha)$. Thus, the excess risk contribution from the region A is at most $(1/2 + \alpha)\mathbb{P}^*(|\eta^* - 1/2| < \alpha)$. □

J Further details from Examples 1 to 3

In Appendices J.2 and J.3, we use an operation analogous to S_ϵ that calculates the infimum of a function over an ϵ-ball. Formally, we define:

$$I_\epsilon(g)(\mathbf{x}) = \inf_{\|\mathbf{x}' - \mathbf{x}\| \leq \epsilon} g(\mathbf{x}'). \tag{70}$$

Next, we define a mapping α_ϕ from $\eta \in [0, 1]$ to minimizers of $C_\phi(\eta, \cdot)$ by

$$\alpha_\phi(\eta) = \inf\{\alpha : \alpha \text{ is a minimizer of } C_\phi(\eta, \cdot)\}. \tag{71}$$

Lemma 25 of Frank & Niles-Weed (2024b) shows that the function α_ϕ defined in (71) maps η to the smallest minimizer of $C_\phi(\eta, \cdot)$ and is non-decreasing. This property will be instrumental in constructing minimizers for R_ϕ^ϵ.

J.1 Proof of Lemma 2

Proof of Lemma 2. If $R_*^\epsilon = 0$, by Theorem 5, for any measures $\mathbb{P}_0' \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_0)$, $\mathbb{P}_1' \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_1)$ we have $\mathbb{P}'(\eta' = 0 \text{ or } 1) = 1$, where $\mathbb{P}' = \mathbb{P}_0' + \mathbb{P}_1'$ and $\eta' = d\mathbb{P}_1'/d\mathbb{P}'$. This statement must also hold for the $\mathbb{P}_0^* \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_0)$, $\mathbb{P}_1^* \in \mathcal{B}_\epsilon^\infty(\mathbb{P}_1)$ that maximize $\bar{R}_\phi$. □

J.2 Calculating the optimal $\mathbb{P}_0^*$, $\mathbb{P}_1^*$ for Example 2

First, notice that a minimizer of R_ϕ is given by $f(x) = \alpha_\phi(\eta(x))$ with $\eta(x)$ as defined in (16). Below, we construct a minimizer f^* for R_ϕ^ϵ. We'll do this construction separately for $\epsilon \leq \delta$ and $\epsilon \in (\delta, 1 - \delta)$.

When $\epsilon \leq \delta$:

Define a function $\tilde{\eta} : [-\delta - \epsilon - 1, 1 + \delta + \epsilon] \to [0, 1]$ by

$$\tilde{\eta}(x) = \begin{cases} \frac{1}{4} & \text{if } x \in [-1 - \delta - \epsilon, 0) \\ \frac{1}{2} & \text{if } x = 0 \\ \frac{3}{4} & \text{if } x \in (0, 1 + \delta + \epsilon] \end{cases}$$

and a function f^* by $f^*(x) = \alpha_\phi(\tilde{\eta}(x))$.

We'll verify that this function is a minimizer by showing that $R^\epsilon_\phi(f^*) = R_\phi(f)$. As the minimal possible adversarial risk is bounded below by $R_{\phi,*}$, one can conclude that f^* minimizes R^ϵ_ϕ. Consequently, $\bar{R}_\phi(\mathbb{P}_0, \mathbb{P}_1) = R^\epsilon_\phi(f)$ and thus the strong duality result in Theorem 7 would imply that $\mathbb{P}_0$, $\mathbb{P}_1$ must maximize the dual problem.

As both $\tilde{\eta}$ and α_ϕ are non-decreasing, the function f^* must be non-decreasing as well. Consequently, $S_\epsilon(\phi(f^*))(x) = \phi(I_\epsilon(f^*)(x)) = \phi(f^*(x-\epsilon))$ and similarly, $S_\epsilon(\phi(-f^*))(x) = \phi(-S_\epsilon(f^*)(x)) = \phi(-f^*(x+\epsilon))$. (Recall the I_ϵ operation was defined in (70).)

Therefore,

$$\begin{aligned} R^\epsilon_\phi(f^*) &= \int S_\epsilon(\phi(f^*))(x)p_1(x)dx + \int S_\epsilon(\phi(-f^*))(x)p_0(x)dx = \int \phi(f^*(x-\epsilon))p_1(x)dx + \int \phi(-f^*(x+\epsilon))p_0(x)dx \\ &= \int \phi(f^*(x))p_1(x+\epsilon)dx + \int \phi(-f^*(x))p_0(x-\epsilon)dx \end{aligned} \tag{72}$$

Consequently,

$$\begin{aligned} \int \phi(f^*(x))p_1(x+\epsilon)dx &= \int_{-1-\delta-\epsilon}^{-\delta-\epsilon} \frac{1}{8}\phi\Big(\alpha_\phi\Big(\frac{1}{4}\Big)\Big)dx + \int_{\delta-\epsilon}^{1+\delta-\epsilon} \frac{1}{8}\phi\Big(\alpha_\phi\Big(\frac{3}{4}\Big)\Big)dx \\ &= \int_{-1-\delta}^{-\delta} \frac{1}{8}\phi\Big(\alpha_\phi\Big(\frac{1}{4}\Big)\Big)dx + \int_{\delta}^{1+\delta} \frac{1}{8}\phi\Big(\alpha_\phi\Big(\frac{3}{4}\Big)\Big)dx = \int \phi(f(x))p_1(x)dx \end{aligned}$$

Analogously, one can show that

$$\int \phi(-f^*(x))p_0(x-\epsilon)dx = \int \phi(-f(x))p_0(x)dx$$

and consequently $R^\epsilon_\phi(f^*) = R_\phi(f)$.

When $\epsilon \in (\delta, 1+\delta)$:

We will show that $R^\epsilon_\phi(f^*) = \bar{R}_\phi(\mathbb{P}^*_0, \mathbb{P}^*_1)$ for the proposed attacks, proving that $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ are dual optimal distributions. This time, define the function $\tilde{\eta} : [-\delta-\epsilon-1, 1+\delta+\epsilon] \to [0,1]$ by

$$\tilde{\eta}(x) = \begin{cases} 0 & \text{if} x \in [-1-\delta-\epsilon, -1-\delta+\epsilon) \\ \frac{1}{4} & \text{if } x \in [-1-\delta+\epsilon, -(\epsilon-\delta)) \\ \frac{1}{2} & \text{if } x \in [-(\epsilon-\delta), (\epsilon-\delta)] \\ \frac{3}{4} & \text{if } x \in ((\epsilon-\delta), 1+\delta-\epsilon] \\ 1 & \text{if } x \in (1+\delta-\epsilon, 1+\delta+\epsilon] \end{cases}$$

and again take $f^*(x) = \alpha_\phi(\tilde{\eta}(\mathbf{x}))$. The function f^* is non-decreasing, so again (72) holds. Further, defining p^*_1, p^*_0 as $p^*_1(x) = p_1(x+\epsilon)$ and $p^*_0(x) = p_0(x-\epsilon)$ implies the relation

$$R^\epsilon_\phi(f^*) = \int C_\phi(\eta^*, f^*)p^*(x)dx,$$

where $p^*(x) = p^*_0(x) + p^*_1(x)$ and $\eta^* = p^*_1(x)/p^*(x)$. The function $\tilde{\eta}$ was defined so that $\tilde{\eta}(x) = \eta^*(x)$ a.e. and hence

$$C_\phi(\eta^*, f^*) = C_\phi(\eta^*, \alpha_\phi(\eta^*)) = C^*_\phi(\eta^*).$$

This relation implies $R^\epsilon_\phi(f^*) = \bar{R}_\phi(\mathbb{P}^*_0, \mathbb{P}^*_1)$, where $\mathbb{P}^*_0$, $\mathbb{P}^*_1$ are the distributions with pdfs p^*_0 and p^*_1.

J.3 Calculating the optimal $\mathbb{P}_0^*$ and $\mathbb{P}_1^*$ for Example 3— Proof of (17)

We will show that the densities in (17) are dual optimal by finding a function f^* for which $R_\phi^\epsilon(f^*) = \bar{R}_\phi(\mathbb{P}_0^*, \mathbb{P}_1^*)$. Theorem 7 will then imply that $\mathbb{P}_0^*$, $\mathbb{P}_1^*$ must maximize the dual. Define $\eta^*(x)$ by

$$\eta^*(x) = \frac{p_1^*(x)}{p_1^*(x) + p_0^*(x)},$$

with $p_0^*(x)$ and $p_1^*(x)$ as in (17). For a given loss ϕ we will prove that the optimal function f^* is given by

$$f^*(x) = \alpha_\phi(\eta^*(x)).$$

The function η^* computes to

$$\eta^*(x) = \frac{1}{1 + e^{\frac{\mu_1 - \mu_0 - 2\epsilon}{\sigma^2}(\frac{\mu_1+\mu_0}{2} - x)}}.$$

If $\mu_1 - \mu_0 > 2\epsilon$, the conditional probability $\eta^*(x)$ is increasing in x and consequently the function f^* is non-decreasing. Therefore, $S_\epsilon(\phi(f^*))(x) = \phi(I_\epsilon(f^*)(x)) = \phi(f^*(x-\epsilon))$ (recall I_ϵ was defined in (70)). Similarly, one can argue that $S_\epsilon(\phi(-f^*))(x) = \phi(-f^*(x+\epsilon))$, and therefore,

$$\begin{aligned} R_\phi^\epsilon(f^*) &= \int S_\epsilon(\phi(f^*))(x)p_1(x)dx + \int S_\epsilon(\phi(-f^*))(x)p_0(x)dx = \int \phi(f^*(x-\epsilon))p_1(x)dx + \int \phi(-f^*(x+\epsilon))p_0(x)dx \\ &= \int \phi(f^*(x))p_1(x+\epsilon)dx + \int \phi(-f^*(x))p_0(x-\epsilon)dx. \end{aligned}$$

Next, notice that $p_1(x+\epsilon) = p_1^*(x)$ and $p_0(x-\epsilon) = p_0^*(x)$. Define $\mathbb{P}^* = \mathbb{P}_0^* + \mathbb{P}_1^*$. Then

$$R_\phi^\epsilon(f^*) = \int \eta^*\phi(\alpha_\phi(\eta^*)) + (1-\eta^*)\phi(-\alpha_\phi(\eta^*))d\mathbb{P}^* = \int C_\phi^*(\eta^*)d\mathbb{P}^* = \bar{R}_\phi(\mathbb{P}_0^*, \mathbb{P}_1^*)$$

Consequently, the strong duality result in Theorem 7 implies that $\mathbb{P}_0^*$ $\mathbb{P}_1^*$ must maximize the dual $\bar{R}_\phi$.

J.4 Showing (18)

Lemma 22. *Consider an equal gaussian mixture with variance σ and means $\mu_0 < \mu_1$, with pdfs given by*

$$p_0(x) = \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu_0)^2}{2\sigma^2}}, \quad p_1(x) = \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu_1)^2}{2\sigma^2}}$$

Let $\eta(x) = p_1(x)/(p_0(x) + p_1(x))$. Then $|\eta(x) - 1/2| \le z$ iff $x \in [\frac{\mu_0+\mu_1}{2} - \Delta(z), \frac{\mu_0+\mu_1}{2} + \Delta(z)]$, where $\Delta(z)$ is defined by

$$\Delta(z) = \frac{\sigma^2}{\mu_1 - \mu_0} \ln\left(\frac{\frac{1}{2} + z}{\frac{1}{2} - z}\right). \tag{73}$$

Proof. The function η can be rewritten as $\eta(x) = 1/(1 + p_0/p_1)$ while

$$\frac{p_0(x)}{p_1(x)} = \exp\left(-\frac{(x-\mu_0)^2}{2\sigma^2} + \frac{(x-\mu_1)^2}{2\sigma^2}\right) = \exp\left(\frac{\mu_1 - \mu_0}{\sigma^2}\left(\frac{\mu_1+\mu_0}{2} - x\right)\right)$$

Consequently, $|\eta(x) - 1/2| \le z$ is equivalent to

$$\frac{1}{2} - z \le \frac{1}{\exp\left(\frac{\mu_1-\mu_0}{\sigma^2}(\frac{\mu_1+\mu_0}{2} - x)\right) + 1} \le \frac{1}{2} + z$$

which is equivalent to

$$\frac{\mu_1+\mu_0}{2} - \frac{\sigma^2}{\mu_1-\mu_0}\ln\left(\frac{1}{\frac{1}{2}-z}-1\right) \le x \le \frac{\mu_1+\mu_0}{2} - \frac{\sigma^2}{\mu_1-\mu_0}\ln\left(\frac{1}{z+\frac{1}{2}}-1\right)$$

Finally, notice that

$$\Delta(z) = \frac{\sigma^2}{\mu_1-\mu_0}\ln\left(\frac{1}{\frac{1}{2}-z}-1\right) \tag{74}$$

while

$$-\Delta(z) = \frac{\sigma^2}{\mu_1-\mu_0}\ln\left(\frac{1}{z+\frac{1}{2}}-1\right)$$

□

Lemma 23. *Let p_0, p_1, and η be as in Lemma 22 and let $h(z) = \mathbb{P}(|\eta - 1/2| \le z)$. Then if $\mu_1 - \mu_0 \le \sqrt{2}\sigma$, then h is concave.*

Proof. To start, we calculate the second derivative of $\Delta(z)$ and the first derivative of p_0.

The first derivative of Δ is

$$\Delta'(z) = \frac{\sigma^2}{\mu_1-\mu_0}\cdot\frac{1}{\frac{1}{4}-z^2}.$$

and the second derivative of $\Delta(z)$ is

$$\Delta''(z) = \frac{\sigma^2}{\mu_1-\mu_0}\cdot\frac{2z}{(\frac{1}{4}-z^2)^2} \tag{75}$$

Next, one can calculate the derivative of p_0 as

$$p_0'(x) = \frac{1}{2}\cdot\frac{1}{\sqrt{2\pi}\sigma}\cdot\frac{-(x-\mu_0)}{\sigma^2}e^{-\frac{(x-\mu_0)^2}{2\sigma^2}} = -\frac{(x-\mu_0)}{\sigma^2}p_0(x) \tag{76}$$

and similarly

$$p_1'(x) = -\frac{(x-\mu_1)}{\sigma^2}p_1(x) \tag{77}$$

Let $p(x) = p_0 + p_1$. Lemma 22 implies that the function h is given by $h(z) = \int_{\frac{\mu_1+\mu_0}{2}-\Delta(z)}^{\frac{\mu_1+\mu_0}{2}+\Delta(z)} p(z)dz$. The first derivative of h is then

$$h'(z) = \left(p\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big) + p\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)\Delta'(z).$$

Differentiating h twice results in

$$\begin{aligned}
h''(z) &= \left(p\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big) + p\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)\Delta''(z) \\
&+ \left(p'\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big) - p'\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)(\Delta'(z))^2 \\
&= \left(p\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big) + p\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)\left(\Delta''(z) - \frac{\Delta(z)\Delta'(z)^2}{\sigma^2}\right) && (78) \\
&+ \left(p_1\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big) + p_1\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)\frac{\mu_1-\mu_0}{2\sigma^2}(\Delta'(z))^2 && (79) \\
&- \left(p_0\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big) + p_0\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)\frac{\mu_1-\mu_0}{2\sigma^2}(\Delta'(z))^2. && (80)
\end{aligned}$$

where the final equality is a consequence of Equations (76) and (77). Next, we'll argue that the sum of the terms in Equations (79) and (80) is zero:

$$\begin{aligned}
&\left(p_1\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big)+p_1\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)-\left(p_0\Big(\frac{\mu_1+\mu_0}{2}+\Delta(z)\Big)+p_0\Big(\frac{\mu_1+\mu_0}{2}-\Delta(z)\Big)\right)\\
&=\frac{1}{2\sqrt{2\pi}\sigma}\left(\left(e^{-\frac{\left(\frac{\mu_0-\mu_1}{2}+\Delta(z)\right)^2}{2\sigma^2}}+e^{-\frac{\left(\frac{\mu_0-\mu_1}{2}-\Delta(z)\right)^2}{2\sigma^2}}\right)-\left(e^{-\frac{\left(\frac{\mu_1-\mu_0}{2}+\Delta(z)\right)^2}{2\sigma^2}}+e^{-\frac{\left(\frac{\mu_1-\mu_0}{2}-\Delta(z)\right)^2}{2\sigma^2}}\right)\right)\\
&=0
\end{aligned}$$

Next, we'll show that under the assumption $\mu_1-\mu_0\leq\sqrt{2}\sigma$, the term (78) is always negative. Define $k=\sigma^2/(\mu_1-\mu_0)$. Then

$$\Delta''(z)-\Delta(z)\frac{\Delta'(z)^2}{\sigma^2}=\frac{2k}{(\frac{1}{4}-z^2)^2}\left(z-\frac{k^2}{2\sigma^2}\ln\left(\frac{1}{\frac{1}{2}-z}-1\right)\right)\tag{81}$$

The fact that $\Delta''(z)>0$ for all z implies that $\ln(1/(1/2-z)-1)$ is convex, and this function has derivative 4 at zero. Consequently, $\ln(1/(1/2-z)-1)\geq 4z$ and (81) implies

$$\Delta''(z)-\Delta(z)\frac{\Delta'(z)^2}{\sigma^2}\leq\frac{2k}{(\frac{1}{4}-z^2)^2}(z-\frac{k^2}{2\sigma^2}\cdot 4z)=\frac{2kz}{(\frac{1}{4}-z^2)^2}\left(1-\frac{2k^2}{\sigma^2}\right)$$

The condition $\mu_1-\mu_0\leq\sqrt{2}\sigma$ is equivalent to $1-2k^2/\sigma^2<0$.

□

This lemma implies that $h(z)\leq h'(0)z$. Noting also that $h(z)\leq 1$ for all z produces the bound

$$h(z)\leq\min\left(\frac{16\sigma^2}{\mu_1-\mu_0}z,1\right)$$

applying this bound to the gaussians with densities p_0^* and p_1^* results in (18).