

Constrained Classification and Policy Learning*

Toru Kitagawa,[†] Shosei Sakaguchi,[‡] and Aleksey Tetenov[§]

February 28, 2022

Abstract

Modern machine learning approaches to classification, including AdaBoost, support vector machines, and deep neural networks, utilize surrogate loss techniques to circumvent the computational complexity of minimizing empirical classification risk. These techniques are also useful for causal policy learning problems, since estimation of individualized treatment rules can be cast as a weighted (cost-sensitive) classification problem. Consistency of the surrogate loss approaches studied in Zhang (2004) and Bartlett et al. (2006) relies on the assumption of *correct specification*, which means that the specified set of classifiers is rich enough to contain a first-best classifier. This assumption is, however, less credible when the set of classifiers is constrained by interpretability or fairness, leaving the applicability of surrogate loss-based algorithms unknown in such second-best scenarios. This paper studies consistency of surrogate loss procedures under a constrained set of classifiers without assuming correct specification. We show that in settings where the constraint restricts the classifier’s prediction set only, hinge losses (i.e., ℓ_1 -support vector machines) are the only surrogate losses that preserve consistency in second-best scenarios. If the constraint additionally restricts the functional form of the classifier, consistency of a surrogate loss approach is not guaranteed, even with hinge loss. We therefore characterize conditions on the constrained set of classifiers that can guarantee consistency of hinge risk minimizing classifiers. Exploiting our theoretical results, we develop robust and computationally attractive hinge loss-based procedures for a monotone classification problem.

Keywords: Surrogate loss, support vector machine, monotone classification, fairness in machine learning, statistical treatment choice, personalized medicine

*We thank Ashesh Rambachan, Jörg Stoye, and Max Tabord-Meehan for valuable discussions and comments. We also thank participants at the 2021 Cowles Foundation Econometrics Conference, 2021 NASMES, 2021 SEA conference, and seminar participants at Bristol, Chicago, Cornell, Glasgow, Penn State, SciencesPo, UC-Irvine, UC-Riverside, and Zürich for beneficial comments. The authors gratefully acknowledge financial support from ERC grant (number 715940) and the ESRC Centre for Microdata Methods and Practice (CeMMAP) (grant number RES-589-28-0001).

[†]Department of Economics, Brown University, and Department of Economics, University College London. Email: toru.kitagawa@brown.edu

[‡]Department of Economics, University College London. Email: s.sakaguchi@ucl.ac.uk

[§]Geneva School of Economics and Management, University of Geneva. Email: aleksey.tetenov@unige.ch

1 Introduction

Binary classification, the prediction of a binary dependent variable $Y \in \{-1, +1\}$ based upon covariate information $X \in \mathcal{X}$, is one of the most fundamental problems in statistics and econometrics. Many modern machine learning algorithms build on statistically and computationally efficient classification algorithms, and their application has had a sizeable impact on various fields of study and in society in general, e.g., pattern recognition, credit approval systems, personalized recommendation systems, to list but a few examples. Since estimation of an optimal treatment assignment policy can be cast as a weighted (cost-sensitive) classification problem (Zadrozny (2003)), methodological advances in the study of the classification problem apply to the causal problem of designing individualized treatment assignment policies. As the allocation of resources in both business and public policy settings has become more evidence-based and dependent upon algorithms, so too has there been increasingly active debate on how to make allocation algorithms respect societal preferences for interpretability and fairness (Dwork et al. (2012)). Understanding the theoretical performance guarantee and efficient implementation of classification algorithms under interpretability or fairness constraints is a problem of fundamental importance with a strong connection to real life.

In the supervised binary classification problem, the typical objective is to learn a classification rule that minimizes the probability of false prediction. We denote the distribution of (Y, X) by P , and a (non-randomized) classifier that predicts $Y \in \{-1, +1\}$ based upon $\text{sign}(f(X))$ by $f : \mathcal{X} \rightarrow \mathbb{R}$, where $\text{sign}(\alpha) = 1\{\alpha \geq 0\} - 1\{\alpha < 0\}$. We denote the 0-level set of f by $G_f \equiv \{x \in \mathcal{X} : f(x) \geq 0\} \subset \mathcal{X}$, and refer to G_f as the *prediction set* of f . The goal is to learn a classifier that minimizes *classification risk*:

$$R(f) \equiv P(\text{sign}(f(X)) \neq Y) = E_P[1\{Y \cdot \text{sign}(f(X)) \leq 0\}]. \quad (1)$$

Given a training sample $\{(Y_i, X_i) \sim_{iid} P : i = 1, \dots, n\}$, the empirical risk minimization principle of Vapnik (1998) recommends estimating the optimal classifier by minimizing empirical classification risk,

$$\begin{aligned} \hat{f} &\in \arg \inf_{f \in \mathcal{F}} \hat{R}(f), \\ \hat{R}(f) &\equiv \frac{1}{n} \sum_{i=1}^n 1\{Y_i \cdot \text{sign}(f(X_i)) \leq 0\}, \end{aligned} \quad (2)$$

over a class of classifiers $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathbb{R}\}$. If the complexity of \mathcal{F} is properly constrained, the empirical risk minimizing (ERM) classifier \hat{f} has statistically attractive properties including risk consistency and minimax rate optimality. See, for example, Devroye et al. (1996) and Lugosi (2002).

Despite the desirable performance guarantee of the ERM classifier, the computational complexity of solving the optimization in (2) becomes a serious hurdle to practical implementation, especially when the dimension of covariates is moderate to large. To get around this issue, the existing literature has offered various alternatives to the ERM classifier, including support vector machines (Cortes and Vapnik (1995)), AdaBoost (Freund and Schapire (1997)), and neural networks. Focusing on optimization, each of these algorithms can be viewed as targeting the minimization of *surrogate risk*,

$$R_\phi(f) \equiv E_P[\phi(Yf(X))], \quad (3)$$

where $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is called the *surrogate loss* function, a different specification of which corresponds to a different learning algorithm. Convex functions make for a desirable choice of surrogate loss function as, combined with some functional form specification for f , the minimization problem for the empirical analogue of the surrogate risk in (3) is a convex optimization problem. This insight and the computational benefit that it yields has been pivotal to learning algorithms being able to handle large scale problems with high-dimensional features.

Can surrogate risk minimization lead to an optimal classifier in terms of the original classification risk? The seminal works of Zhang (2004) and Bartlett et al. (2006) provide theoretical justification for the use of surrogate losses by clarifying the conditions under which surrogate risk minimization also minimizes the original classification risk. A crucial assumption for this important result is *correct specification* of the classifiers, requiring that the class of classifiers \mathcal{F} over which the surrogate risk is minimized contains a classifier that globally minimizes the original classification risk, i.e., a classifier that is identical to or performs as well as the Bayes classifier $f_{Bayes}^*(x) \equiv 2P(Y = 1|X = x) - 1$ in terms of its classification risk.

The credibility of the assumption of correct specification is, however, limited if the set of implementable classifiers is constrained exogenously, independently of any belief concerning the underlying data generating process. Such a situation is becoming more prevalent due to the increasing need for interpretability or fairness of classification algorithms. Given that f determines the classification rule only through G_f , such constraints can be represented by shape restrictions on the prediction set of f , i.e., the class of feasible f is represented by $\mathcal{F}_{\mathcal{G}} \equiv \{f \in \mathcal{F} : G_f \in \mathcal{G}\}$, where \mathcal{G} is a restricted class of sets in \mathcal{X} satisfying the requirements for interpretability and fairness. To the best of our knowledge, how the validity of a surrogate loss approach is affected if $\mathcal{F}_{\mathcal{G}}$ misses the first-best classifier is not known.

The main contribution of this paper is to establish conditions under which a surrogate loss approach is valid without assuming correct specification. We first characterize those

conditions on surrogate loss such that minimization of the surrogate risk can lead to a second-best rule (i.e., constrained optimum) in terms of the original classification risk. Specifically, we show that hinge losses $\phi_h(\alpha) = c \max\{0, 1 - \alpha\}$, $c > 0$, are the only surrogate losses that guarantee consistency of the surrogate risk minimization for a second-best classifier. An important implication of this result is that ℓ_1 -support vector machines are the only surrogate loss-based methods that are robust to misspecification.

The computational attractiveness of a surrogate loss approach crucially depends not only upon the convexity of the surrogate loss function ϕ but also upon the functional form restrictions on the classifier f that lead to a convex \mathcal{F} . We therefore investigate how additional constraints on f on top of $G_f \in \mathcal{G}$ can affect the consistency of the hinge risk minimization. As a second contribution of this paper, we characterize a simple-to-check sufficient condition for consistency of the hinge risk minimization in terms of the additional functional form restrictions we can impose on $\mathcal{F}_{\mathcal{G}}$. We term a subclass of classifiers of $\mathcal{F}_{\mathcal{G}}$ satisfying the sufficient condition a *classification-preserving reduction* of $\mathcal{F}_{\mathcal{G}}$.

Exploiting our main theoretical results, we develop novel procedures for monotone classification. In monotone classification, prediction sets are constrained to

$$\mathcal{G}_M \equiv \{G \subset \mathcal{X} : x \in G \Rightarrow x' \in G \forall x' \leq x\},$$

where $x' \leq x$ is an element-wise weak inequality. Since \mathcal{G}_M coincides with the class of prediction sets spanned by the class of monotonically decreasing bounded functions $\mathcal{F}_M \equiv \{f : f \text{ decreasing in } x, -1 \leq f \leq 1\}$, hinge loss-based estimation for monotone classification can be performed by solving

$$\begin{aligned} \hat{f}_M &\in \arg \inf_{f \in \mathcal{F}_M} \hat{R}_{\phi_h}(f), \\ \hat{R}_{\phi_h}(f) &\equiv \frac{1}{n} \sum_{i=1}^n \phi_h(y_i f(x_i)). \end{aligned} \tag{4}$$

We show that the class of monotone classifiers \mathcal{F}_M is a constrained classification-preserving reduction of $\mathcal{F}_{\mathcal{G}_M}$, guaranteeing consistency of the hinge-risk minimizing classifier \hat{f}_M . Furthermore, we show that convexity of \mathcal{F}_M reduces the optimization of (4) to a finite dimensional linear programming problem and hence delivers significant computational gains relative to minimization of the original empirical classification risk. We also consider approximating \mathcal{F}_M using a sieve of Bernstein polynomials and estimating a monotone classifier by solving (4) over the Bernstein polynomials. Adopting either approach, the application of our main theorems guarantees

$$R(\hat{f}_M) - \inf_{f \in \mathcal{F}_M} R(f) \xrightarrow{p} 0,$$

as $n \rightarrow 0$, and this convergence is valid regardless of whether \mathcal{F}_M attains the first-best risk, i.e., $\inf_{f \in \mathcal{F}_M} R(f) = \inf_{f \in \bar{\mathcal{F}}} R(f)$, or not, where $\bar{\mathcal{F}}$ is the class of measurable functions $f : \mathcal{X} \rightarrow \mathbb{R}$. We also derive the uniform upper bound of the mean of $R(\hat{f}_M) - \inf_{f \in \mathcal{F}_M} R(f)$ to characterize the regret convergence rate attained by \hat{f}_M .

1.1 Connection and contributions to causal policy learning

For simplicity of exposition, this paper mainly focuses on the prototypical setting of binary classification. The main theoretical results can easily be extended to weighted (cost-sensitive) classification, where the canonical representation of the population risk criterion is given by

$$R^\omega(f) \equiv E_P[\omega \cdot 1\{Y \cdot \text{sign}(f(X)) \leq 0\}]. \quad (5)$$

Here, ω is a non-negative random variable defining the cost of misclassifying Y that typically depends on (Y, X) . The cost of misclassification ω may represent the decision-maker's economic cost (Lieli and White (2010)) or welfare weights over the individuals to be classified, as considered in Rambachan et al. (2020) and Babii et al. (2020). The surrogate risk for weighted classification can be defined similarly to (3), as

$$R_\phi^\omega(f) = E_P[\omega \cdot \phi(Yf(X))]. \quad (6)$$

As discussed in Kitagawa and Tetenov (2018), there are fundamental conceptual differences between the prediction problem of classification and the causal problem of treatment choice. Nevertheless, if the training sample is obtained from a randomized control trial (RCT) or an observational study satisfying unconfoundedness (selection on observables), we can view minimization of the weighted classification risk in (5) as being equivalent to the maximization of the additive welfare criterion commonly specified in treatment choice problems. To see this equivalence, let $\{(Z_i, D_i, X_i) : i = 1, \dots, n\}$ be an independent and identically distributed RCT sample of n experimental subjects, where $Z_i \in \mathbb{R}$ is subject i 's observed outcome, $D_i \in \{-1, +1\}$ is an indicator for his assigned treatment, and $X_i \in \mathcal{X}$ is a vector of pretreatment covariates, and let $(Z_i(d) : d \in \{-1, +1\})$ be i 's potential outcomes satisfying $Z_i = Z_i(+1) \cdot 1\{D_i = +1\} + Z_i(-1) \cdot 1\{D_i = -1\}$. We denote the propensity score in the RCT sample by $e(x) \equiv P(D = +1|X = x)$ and assume that $e(x)$ is bounded away from 0 and 1 for all $x \in \mathcal{X}$. We denote the joint distribution of $(Z_i(+1), Z_i(-1), D_i, X_i)$ by P and assume P satisfies unconfoundedness, $(Z(+1), Z(-1)) \perp D|X$.

Similar to our consideration of classification, we represent a (non-randomized) treatment assignment rule by the sign of $f : \mathcal{X} \rightarrow \mathbb{R}$ –i.e., the 0-level set $G_f = \{x \in \mathcal{X} : f(x) \geq 0\} \subset \mathcal{X}$ specifies the subgroup of the population assigned to the treatment +1. Following

Manski (2004), we consider evaluating the welfare performance of the assignment policy f by the average outcomes attained under its associated assignment rule:

$$W(f) \equiv E_P [Z(+1) \cdot 1\{X \in G_f\} + Z(-1) \cdot 1\{X \notin G_f\}]$$

Relying on unconfoundedness of the experimental data and employing the inverse propensity score weighting technique, we can express this welfare in terms of the observable variables as¹

$$\begin{aligned} W(f) &= E_P \left[\frac{Z}{De(X) + (1 - D)/2} \cdot 1\{D = \text{sign}(f(X))\} \right] \\ &= E_P \left[\max \left\{ 0, \frac{Z}{De(X) + (1 - D)/2} \right\} \right] - E_P [\omega_p \cdot 1\{\text{sign}(Z) \cdot D \cdot \text{sign}(f(X)) \leq 0\}], \end{aligned} \tag{7}$$

where

$$\omega_p \equiv \frac{|Z|}{De(X) + (1 - D)/2} \geq 0.$$

Provided that the first moment of ω_p is finite, maximization of $W(f)$ is equivalent to minimization of the weighted classification risk $R^\omega(f)$ defined in (5) with $\omega = \omega_p$ and $Y = \text{sign}(Z) \cdot D$. As a result, optimal treatment assignment rules can be viewed as optimal classifiers for D in terms of weighted classification risk. This equivalence also holds for other methods of policy learning, such as the offset-tree learning of Beygelzimer and Langford (2009) and the doubly-robust approaches of Swaminathan and Joachims (2015) and Athey and Wager (2021), which correspond to different ways of constructing or estimating the weighting term ω_p .

Due to its equivalence to weighted classification, a surrogate loss approach to policy learning proceeds by minimizing the empirical analogue of (6) with $\omega = \omega_p$ and $Y = \text{sign}(Z) \cdot D$. Section 7 of this paper shows that our main theoretical results established for constrained binary classification carry over to the setting of policy learning in which feasible treatment assignment policies are constrained exogenously due to fairness and legislative considerations. This paper therefore offers valuable and novel contributions to current research and public debate regarding how to make use of machine learning algorithms to design individualized policies. If treatment assignment rules are constrained to be monotone, our concrete proposals for monotone classification algorithms can be applied to policy learning, which yields significant gains in computational efficiency relative to the mixed integer programming approaches considered in Kitagawa and Tetenov (2018)

¹Kitagawa and Lin (2021) makes use of this transformation of the welfare objective function to develop an Adaboost algorithm for treatment choice.

and Mbakop and Tabord-Meehan (2021).

1.2 Related literature

This paper is closely related to the literature of consistency and performance guarantees for surrogate risk minimization. Notable works in this literature include Mannor et al. (2003), Jiang (2004), Lugosi and Vayatis (2004), Zhang (2004), Steinwart (2005, 2007), Bartlett et al. (2006), Nguyen et al. (2009), and Scott (2012). Under the assumption of correct specification, Zhang (2004) and Bartlett et al. (2006) derive quantitative relationships between excess classification risk and excess surrogate risk, and then provide general conditions for surrogate risk minimization to achieve risk consistency. Bartlett et al. (2006) show that the classification-calibration property of surrogate loss, defined in Section 3 below, guarantees risk consistency. Zhang (2004) and Bartlett et al. (2006) show that many commonly used surrogate loss functions, including hinge loss, exponential loss, and truncated quadratic loss, satisfy the conditions needed for risk consistency. In a classification problem different from ours, where a pair comprising a quantizer and a classifier is chosen, Nguyen et al. (2009) study sufficient and necessary conditions for surrogate risk minimization to yield risk consistency. Nguyen et al. (2009) show that only hinge loss functions satisfy the conditions required for risk consistency in their problem. Correct specification of the class of classifiers is an essential condition for consistency in all of the surrogate risk minimization approaches studied in the literature. The key contribution of our paper is to relax the assumption of correct specification and to clarify the conditions that are required for the surrogate loss function to yield a consistent surrogate risk minimization procedure.

Relaxing the assumption of correct specification connects this paper to classification problems with exogenous constraints. Such problems are studied in machine learning and statistics, and include interpretable classification (e.g., Zeng et al. (2017), and Zhang et al. (2018)), fair classification (e.g., Dwork et al. (2012)), and monotone classification (e.g., Cano et al. (2019)). Some works in the existing literature adopt a surrogate loss approach. Donini et al. (2018) use the ℓ_1 -support vector machine in fair classification, where the hinge risk minimization is subject to a statistical fairness constraint. Chen and Li (2014) use the ℓ_1 -support vector machine with a monotonicity constraint, which constrains the class of feasible classifiers to a class of certain monotone functions. However, neither paper shows consistency of their hinge risk minimization procedures in terms of classification risk.

Agarwal et al. (2018) propose an approach to reduce fairness constrained classification problems to weighted classification. Their reduction can accommodate various fairness constraints proposed in the machine learning literature and the monotonicity constraint of

this paper if X is discrete. Our consistency results on surrogate risk minimizing classifiers can apply to an arbitrary class \mathcal{G} regardless of whether their reduction applies or not.

Focusing on optimization, ERM classification and maximum score estimation (Manski (1975), Manski and Thompson (1989)) share the same objective function. Horowitz (1992) proposes smooth maximum score estimation, where kernel smoothing is performed on the 0-1 loss to obtain a differentiable objective function. However, the smoothed objective function remains non-convex and does not offer the computational gains that the surrogate risk minimization approach with convex surrogates can deliver.

This paper also contributes to a growing literature on statistical treatment rules in econometrics, including Manski (2004), Dehejia (2005), Hirano and Porter (2009), Stoye (2009, 2012), Chamberlain (2011), Bhattacharya and Dupas (2012), Tetenov (2012), Kasy (2018), Kitagawa and Tetenov (2018, 2021), Viviano (2019), Kitagawa and Wang (2020), Athey and Wager (2021), Mbakop and Tabord-Meehan (2021), Sakaguchi (2021), among others. As discussed above, the policy learning methods of Kitagawa and Tetenov (2018), Athey and Wager (2021), and Mbakop and Tabord-Meehan (2021) build on the similarity between empirical welfare maximizing treatment choice and ERM classification. Mbakop and Tabord-Meehan (2021) propose penalization methods to control the complexity of treatment assignment rules, and derive relevant finite sample upper bounds on the regret of the estimated treatment rules. Athey and Wager (2021) apply doubly-robust estimators to estimate the weight ω in (6), and show that an $1/\sqrt{n}$ -upper bound on regret can also be achieved in the observational study setting. These works optimize an empirical welfare objective involving an indicator loss function. As a result, the practical implementation of such methods is sometimes discouraging, especially when the sample size or number of covariates is moderate to large.

Estimation of individualized treatment rules is a topic of active research in other fields including medical statistics, machine learning, and computer science. Notable works in these fields include Zdrozny (2003), Beygelzimer and Langford (2009), Qian and Murphy (2011), Zhao et al. (2012), Swaminathan and Joachims (2015), Zhao et al. (2015), and Kallus (2021), among others. Zhao et al. (2012) propose using ℓ_1 -support vector machines to solve the weighted classification with individualized treatment choice problem, and show risk consistency. They specify a rich class of treatment choice rules that is a reproducing kernel Hilbert space, and assume correct specification. Zhao et al. (2015) extend this approach to estimate optimal dynamic treatment regimes.

2 Constrained classification with surrogate loss

Consider the binary classification problem of ascribing a binary label $Y \in \{-1, +1\}$ based upon covariates $X \in \mathcal{X}$, which are collectively distributed according to a joint

distribution P . We let X be a d_x -dimensional vector, $d_x < \infty$, and denote its marginal distribution by P_X . We denote the conditional probability of $Y = +1$ given $X = x$ by $\eta(x) \equiv P(Y = +1|X = x)$ and otherwise maintain the notation introduced in the Introduction. The ultimate objective is to minimize the classification risk of (1).

We study constrained classification problems where an optimal classifier is searched for over a restricted class of functions. Section 2.1 studies the consistency of surrogate risk minimization in the special case that the prespecified class of classifiers contains a classifier whose prediction set agrees with the prediction set of the Bayes classifier. Section 2.2 introduces a classification problem that embeds a constraint on the prediction sets, which is a central problem throughout the paper.

2.1 Misspecification in constrained classification

Let \mathcal{F} be a constrained class of classifiers $f : \mathcal{X} \rightarrow \mathbb{R}$. If the set of classifiers were unconstrained, it is well known that the Bayes classifier defined by

$$f_{Bayes}^* = 2\eta(x) - 1$$

minimizes the classification risk. Due to the constraints on the class of classifiers, however, the minimized classification risk on \mathcal{F} can be strictly larger than the first-best minimal risk $R(f_{Bayes}^*)$. We refer to this situation as R -misspecification of \mathcal{F} , which we formally define in the following definition.

Definition 2.1 (R -misspecification). \mathcal{F} is R -misspecified if

$$\inf_{f \in \mathcal{F}} R(f) > R(f_{Bayes}^*).$$

If the inequality instead holds with equality, we say that \mathcal{F} is R -correctly specified.

Because the 0-1 loss function is neither convex nor continuous, minimizing the empirical analog of $R(f)$ is computationally challenging and often infeasible given the scale of the problems that we encounter in practice. Commonly used classification algorithms, such as boosting and support vector machines, replace the 0-1 loss with a surrogate loss function, $\phi : \mathbb{R} \rightarrow \mathbb{R}$, and aim to minimize the surrogate risk $R_\phi(f) \equiv E_P[\phi(Yf(X))]$. Table 1 below lists some commonly used surrogate loss functions including the *hinge loss* $\phi_h(\alpha) = c \max\{0, 1 - \alpha\}$, which corresponds to ℓ_1 -support vector machines, and the *exponential loss* $\phi_e(\alpha) = \exp(-\alpha)$, which corresponds to AdaBoost.

We also introduce the concept of misspecification of \mathcal{F} in terms of surrogate risk as follows.

Definition 2.2 (R_ϕ -misspecification). Let $f_{\phi,FB}^*$ be a minimizer of R_ϕ over the unconstrained class of classifiers, i.e., the class of all measurable functions $f : \mathcal{X} \rightarrow \mathbb{R}$. A constrained class \mathcal{F} is R_ϕ -misspecified if

$$\inf_{f \in \mathcal{F}} R_\phi(f) > R_\phi(f_{\phi,FB}^*).$$

If the inequality instead holds with equality, we say that \mathcal{F} is R_ϕ -correctly specified.

The seminal theoretical results that guarantee consistency of surrogate-risk classification (Zhang (2004), Bartlett et al. (2006), and Nguyen et al. (2009)) crucially rely on the assumption that \mathcal{F} is both R -correctly specified and R_ϕ -correctly specified in the sense of Definitions 2.1 and 2.2, respectively. The central question that this paper poses is *how is a surrogate loss approach affected if \mathcal{F} is R -misspecified or R_ϕ -misspecified?* This misspecification is a likely scenario, especially when the origins of the constraints have nothing to do with the assumptions on P , as is the case in the examples discussed in the next subsection.

Throughout the paper, we limit our analysis to the class of classification-calibrated loss functions defined in Bartlett et al. (2006).

Definition 2.3 (Classification-calibrated loss functions). For $a \in \mathbb{R}$ and $0 \leq b \leq 1$, define $C_\phi(a, b) \equiv \phi(a)b + \phi(-a)(1 - b)$. A loss function ϕ is classification-calibrated if for any $b \in [0, 1] \setminus \{1/2\}$,

$$\inf_{\{a \in \mathbb{R} : a(2b-1) < 0\}} C_\phi(a, b) > \inf_{\{a \in \mathbb{R} : a(2b-1) \geq 0\}} C_\phi(a, b).$$

Noting that the surrogate risk can be expressed as

$$E_P[\phi(Yf(X))] = E_{P_X}[C_\phi(f(X), \eta(X))], \quad (8)$$

the definition of classification-calibrated loss functions implies that at every $x \in \mathcal{X}$ with $\eta(x) \neq 1/2$, every $f(x)$ that minimizes $C_\phi(f(x), \eta(x))$ has the same sign as the Bayes classifier, $\text{sign}(2\eta(x) - 1)$. Bartlett et al. (2006) shows that many commonly used surrogate loss functions including those listed in Table 1 are classification-calibrated.²

Having introduced two notions of misspecification, we now clarify the relationship between R -misspecification and R_ϕ -misspecification.

²Bartlett et al. (2006) also show that any convex loss function ϕ is classification-calibrated if and only if it is differentiable at 0 and $\phi'(0) < 0$.

Proposition 2.1. *Let \mathcal{F} be a constrained class of classifiers and $f_\phi^* \in \mathcal{F}$ be a minimizer of R_ϕ over \mathcal{F} . Suppose ϕ is a classification-calibrated loss function.*

(i) *For any distribution P on $\{-1, 1\} \times \mathcal{X}$, if \mathcal{F} is R_ϕ -correctly specified, then \mathcal{F} is R -correctly specified and $R(f_\phi^*) = R(f_{\text{Bayes}}^*)$ holds;*

(ii) *If ϕ is, in addition, convex, there exist a distribution P on $\{-1, 1\} \times \mathcal{X}$ and a class of classifiers \mathcal{F} under which \mathcal{F} is R -correctly specified but R_ϕ -misspecified, and $R(f_\phi^*) > R(f_{\text{Bayes}}^*)$ holds.*

Proof. See Appendix A. □

Proposition 2.1 (i), which rephrases Claim 3 of Theorem 1 in Bartlett et al. (2006), implies that surrogate risk minimization on the R_ϕ -correctly specified class \mathcal{F} leads to (first-best) optimal classification in terms of the classification risk. An equivalent statement following Theorem 1 in Bartlett et al. (2006) is that for any P and every sequence of measurable functions $\{f_i : \mathcal{X} \rightarrow \mathbb{R}\}$,

$$R_\phi(f_i) \rightarrow \inf_{f \in \mathcal{F}} R_\phi(f) \text{ implies that } R(f_i) \rightarrow \inf_{f \in \mathcal{F}} R(f).$$

This result justifies the approach of surrogate risk minimization when \mathcal{F} is a sufficiently rich class of classifiers (e.g., the reproducing kernel Hilbert space of functions with a large number of features as used in support vector machines), since R_ϕ -correct specification, which is a credible assumption to make given a rich class of classifiers, guarantees R -correct specification.

Proposition 2.1 (ii), in contrast, shows that R -correct specification of \mathcal{F} does *not* guarantee R_ϕ -correct specification.³ R_ϕ -misspecification of \mathcal{F} can lead to the selection of a suboptimal classifier in \mathcal{F} in terms of the classification risk, which illustrates the pitfall of adopting a surrogate loss approach with constrained classifiers. Even when we are confident that the constrained class \mathcal{F} is R -correctly specified, we cannot justify the use of \mathcal{F} in the surrogate risk minimization.

2.2 \mathcal{G} -constrained classification

In this section, we consider restricting the class of classifiers by requiring that their prediction sets belong to a prespecified class of sets, $\mathcal{G} \subset 2^{\mathcal{X}}$. See Examples 2.4–2.6 below for motivating examples.

³Given a convex classification-calibrated loss function ϕ , our proof of Proposition 2.1 (ii) in Appendix A constructs a pair comprising a R -correctly specified class of classifiers \mathcal{F} and a distribution P that leads to R_ϕ -misspecification. In the construction, we assume that $x_1 \neq x_2 \in \mathcal{X}$ supported by P_X on which $\phi(f(x_1)) < \phi(-f(x_2))$ holds for all $f \in \mathcal{F}$ and that $f(x_1) < 0 \leq f(x_2)$ holds for some $f \in \mathcal{F}$, and consider P that specifies a value of $\eta(x_2)$ close to 1, and a value of $\eta(x_1)$ slightly below 1/2. Such a construction of P is not pathological or limited to the specific class of classifiers considered in the proof.

We denote by

$$\mathcal{F}_{\mathcal{G}} \equiv \{f : G_f \in \mathcal{G}, f(\cdot) \in [-1, 1]\}$$

the class of classifiers whose prediction sets are constrained to \mathcal{G} . In this definition, we restrict f to be bounded and, without loss of generality, normalize its range to $[-1, 1]$. Other than on the shape of the 0-level set and on the range, $\mathcal{F}_{\mathcal{G}}$ does not impose any constraint on the functional form of $f \in \mathcal{F}_{\mathcal{G}}$. The goal of the constrained classification problem is then to find a best classifier, in the sense that it minimizes the classification risk $R(\cdot)$ over $\mathcal{F}_{\mathcal{G}}$. We refer to $\mathcal{F}_{\mathcal{G}}$ as the \mathcal{G} -constrained class of classifiers and to the classification problem over $\mathcal{F}_{\mathcal{G}}$ as \mathcal{G} -constrained classification.

The specification of the class of prediction sets \mathcal{G} represents the fairness, interpretability, and other exogenous requirements that are desired for classification rules. Some examples follow.

Example 2.4 (Interpretable classification). *Decision-makers may prefer simple decision or classification rules that are easily understood or explained even at the cost of harming prediction accuracy. This concept, often referred to as interpretable machine learning, has been pursued, for instance, in the prediction analysis of recidivism (Zeng et al. (2017)) and the decision on medical intervention protocol (Zhang et al. (2018)). An example is a linear classification rule, in which \mathcal{G} is a class of half-spaces with linear boundaries in \mathcal{X} ,*

$$\mathcal{G} = \{x \in \mathbb{R}^{d_x} : x^T \beta \geq 0, \beta \in \mathbb{R}^{d_x}\}.$$

Note that $f \in \mathcal{F}_{\mathcal{G}}$ is not restricted to be a linear function. Any function f , including nonlinear functions, is included in $\mathcal{F}_{\mathcal{G}}$ as long as its prediction set G_f is a hyperplane in \mathcal{X} . A classification tree is another type of classification rule that is interpretable. See, e.g., Breiman et al. (1984).

Example 2.5 (Monotone classification). *The framework we study can accommodate monotonicity constraints on classification. Formally, a monotonicity constraint corresponds to a partial order \preceq on \mathcal{X} , and any prediction set G_f has to respect this partial order in the sense that if $x_1 \preceq x_2$ and $x_1 \in G_f$, then $x_2 \in G_f$. Monotonicity constraints have been utilized in the classification of credit rating (Chen and Li (2014)), and in the assignment of job training in the context of policy learning (Mbakop and Tabord-Meehan (2021)).*

Example 2.6 (Fair classification). *Specification of \mathcal{G} can accommodate some fairness constraints introduced in the literature on fair classification. Let $A = \{0, 1\}$ be an element of X indicating a binary protected group variable (e.g., race, gender). The decision-maker*

wants to ensure fairness of classification by, for instance, equalizing the raw positive classification rate (known as statistical parity): $P_X(f(x) \geq 0 \mid A = 1) = P_X(f(x) \geq 0 \mid A = 0)$. The classification problem embedding this constraint is equivalent to \mathcal{G} -constrained classification with

$$\mathcal{G} = \{G \in 2^{\mathcal{X}} : P_X(X \in G \mid A = 1) = P_X(X \in G \mid A = 0)\},$$

where \mathcal{G} depends on P_X in this case. This fairness constraint is studied by Calders and Verwer (2010), Kamishima et al. (2011), Dwork et al. (2012), Feldman et al. (2015), among others. Some other forms of fairness constraint, such as equalized odds and equalized positive predictive value as reviewed by Chouldechova and Roth (2018), can be accommodated in our framework as well via an appropriate construction of \mathcal{G} .

In the \mathcal{G} -constrained classification problem, R -correct specification of $\mathcal{F}_{\mathcal{G}}$ is necessary and sufficient for the surrogate risk minimizer f_{ϕ}^* to achieve the first-best minimum risk.

Proposition 2.2. *Suppose ϕ is a classification-calibrated loss function. Let $\mathcal{G} \subseteq 2^{\mathcal{X}}$ be a class of measurable subsets of \mathcal{X} and $f_{\phi}^* \in \mathcal{F}_{\mathcal{G}}$ be a minimizer of R_{ϕ} over $\mathcal{F}_{\mathcal{G}}$. Then, for any distribution P on $\{-1, 1\} \times \mathcal{X}$, $R(f_{\phi}^*) = R(f_{\text{Bayes}}^*)$ holds if and only if $\mathcal{F}_{\mathcal{G}}$ is R -correctly specified.*

Proof. See Appendix A. □

Proposition 2.2 shows that if ϕ is classification-calibrated, $f_{\phi}^* \in \mathcal{F}_{\mathcal{G}}$ that minimizes the surrogate risk over $\mathcal{F}_{\mathcal{G}}$ leads to a globally optimal classifier in terms of the classification risk if and only if $\mathcal{F}_{\mathcal{G}}$ is R -correctly specified. A comparison of Proposition 2.1 (ii) and Proposition 2.2 clarifies a special feature of the \mathcal{G} -constrained class of classifiers. Specifically, Proposition 2.1 (ii) establishes that, in general, R -correct specification of a constrained class of classifiers \mathcal{F} does not guarantee $R(f_{\phi}^*) = R(f_{\text{Bayes}}^*)$. In contrast to the seminal results about surrogate risk consistency shown in Zhang (2004) and Bartlett et al. (2006), our claim does not require R_{ϕ} -correct specification of $\mathcal{F}_{\mathcal{G}}$.

If constraints defining \mathcal{G} are motivated by some considerations that are independent of any belief on the underlying data generating process (e.g., Examples 2.4–2.6 above), R -correct specification of $\mathcal{F}_{\mathcal{G}}$ is hard to justify. Therefore, an important question for our analysis to consider is whether or not surrogate risk minimization procedures can yield a classifier achieving $\inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f)$ without requiring R -correct specification of $\mathcal{F}_{\mathcal{G}}$.

3 Calibration of \mathcal{G} -constrained classification

This section investigates the risk consistency of a surrogate risk minimization approach over $\mathcal{F}_{\mathcal{G}}$, where $\mathcal{F}_{\mathcal{G}}$ is now allowed to be R -misspecified. Let f^* be an optimal classifier that minimizes the classification risk over $\mathcal{F}_{\mathcal{G}}$:

$$f^* \in \arg \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f).$$

Similarly, we denote a best classifier among $\mathcal{F}_{\mathcal{G}}$ in terms of the surrogate risk by f_{ϕ}^* ,

$$f_{\phi}^* \in \arg \inf_{f \in \mathcal{F}_{\mathcal{G}}} R_{\phi}(f),$$

To begin our analysis, let us first perform a simple numerical example to assess the influence of misspecification in constrained classification.

Example 3.1 (Numerical example 1). *Let $\mathcal{X} = \{0, 1, 2\}$ and $\mathcal{G} = \{\emptyset, \{2\}, \{2, 1\}, \{2, 1, 0\}\}$. Here, \mathcal{G} imposes monotonicity of the prediction sets in a way that is compatible with Example 2.5. We specify P_X to be uniform on \mathcal{X} and $P(Y = +1 | X = 0) = 0.9$, $P(Y = +1 | X = 1) = 0.3$, and $P(Y = +1 | X = 2) = 0.2$. The Bayes classifier therefore predicts $Y = +1$ at $x = 0$ and $Y = -1$ at $x = 1$ and 2 , but such a prediction set is excluded from \mathcal{G} . That is, $\mathcal{F}_{\mathcal{G}}$ is R -misspecified. Under this specification, the second-best (constrained optimum) classifier f^* has a prediction set equal to \emptyset , and attains the classification risk $R(f^*) = 0.47$.*

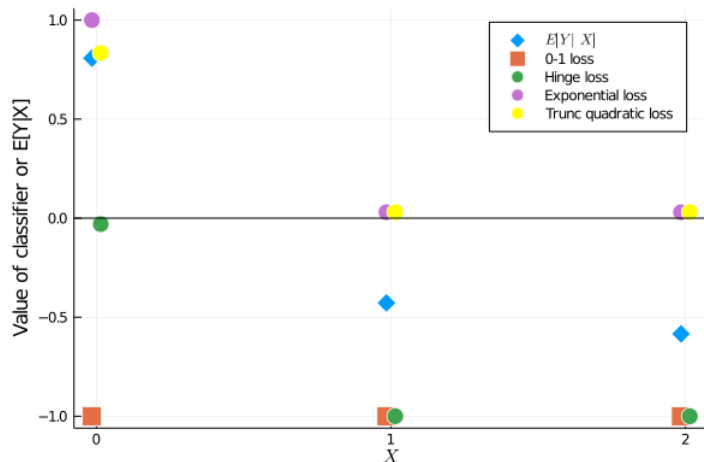
For each of hinge loss ϕ_h with $c = 1$, exponential loss ϕ_e , and truncated quadratic loss ϕ_q , we compute the classifier minimizing the surrogate risk f_{ϕ}^ and the classification risk at the surrogate optimal classifier $R(f_{\phi}^*)$. Figure 1 illustrates each computed classifier with each loss function. We obtain*

$$\begin{aligned} R(f_{\phi_h}^*) &= 0.47 = R(f^*), & R(f_{\phi_e}^*) &= R(f_{\phi_q}^*) = 0.53, \\ G_{f_{\phi_h}^*} &= \emptyset = G_{f^*}, & G_{f_{\phi_e}^*} &= G_{f_{\phi_q}^*} = \{2, 1, 0\}. \end{aligned}$$

In this specification, the hinge risk optimal classifier agrees with the second best optimal classifier, whereas this is not the case for the exponential or truncated quadratic loss.

This example illustrates that hinge loss is robust to R -misspecification of $\mathcal{F}_{\mathcal{G}}$, but exponential and truncated quadratic losses are not. To what extent, can we generalize this finding? What conditions do we need to guarantee that surrogate risk minimizing classifiers are consistent to the second-best (constrained optimal) classification rule in terms of the classification risk? We answer these questions below.

Figure 1: Monotone classifiers minimizing classification and surrogate risks



Notes: The square points correspond to the values of $f^*(x)$ at $x = 0, 1$, and 2 . The circular points correspond to the values of each of $f_{\phi_h}^*(x)$, $f_{\phi_e}^*(x)$, and $f_{\phi_q}^*(x)$ at $x = 0, 1$, and 2 .

For any classifier f , we define the \mathcal{G} -constrained excess risk of f as

$$R(f) - \inf_{f \in \mathcal{F}_G} R(f),$$

which is the *regret* of f relative to a constrained optimum f^* in terms of the classification risk. Similarly, we define the \mathcal{G} -constrained excess ϕ -risk of f as

$$R_\phi(f) - \inf_{f \in \mathcal{F}_G} R_\phi(f).$$

Fix $G \in \mathcal{G}$ and let

$$\mathcal{F}_G \equiv \{f : G_f = G, f(\cdot) \in [-1, 1]\}$$

be the class of classifiers that share the prediction set G . Then $\{\mathcal{F}_G : G \in \mathcal{G}\}$ forms a partition of \mathcal{F}_G indexed by the prediction set, and satisfies $\mathcal{F}_G = \cup_{G \in \mathcal{G}} \mathcal{F}_G$ and $\mathcal{F}_G \cap \mathcal{F}_{G'} = \emptyset$ for $G, G' \in \mathcal{G}$ with $G \neq G'$. With this definition to hand, choosing a classifier from \mathcal{F}_G can be decomposed into two steps: choosing a prediction set G from \mathcal{G} and, then, choosing a classifier f from \mathcal{F}_G .

Denote the classification risk evaluated at a prediction set G by $\mathcal{R}(G) \equiv \inf_{f \in \mathcal{F}_G} R(f)$. Note that any $f \in \mathcal{F}_G$ attains the same level of classification risk, so $\mathcal{R}(G) = R(f)$ holds for all $f \in \mathcal{F}_G$. $\mathcal{R}(G)$ can be written as

$$\mathcal{R}(G) = \int_{\mathcal{X}} [\eta(x)1\{x \notin G\} + (1 - \eta(x))1\{x \in G\}] dP_X(x),$$

$$= \int_{\mathcal{X}} (1 - 2\eta(x)) \cdot 1\{x \in G\} dP_X(x) + P(Y = 1). \quad (9)$$

Similarly, we define the surrogate risk evaluated at G by $\mathcal{R}_\phi(G) \equiv \inf_{f \in \mathcal{F}_G} R_\phi(f)$, which can be written as

$$\begin{aligned} \mathcal{R}_\phi(G) &= \inf_{f \in \mathcal{F}_G} \int_{\mathcal{X}} [\eta(x)\phi(f(x)) + (1 - \eta(x))\phi(-f(x))] dP_X(x) \\ &= \int_G \inf_{0 \leq f(x) \leq 1} C_\phi(f(x), \eta(x)) dP_X(x) + \int_{G^c} \inf_{-1 \leq f(x) < 0} C_\phi(f(x), \eta(x)) dP_X(x), \end{aligned}$$

where the second line follows from the fact that $f \in \mathcal{F}_G$ is unconstrained other than via its prediction set and that the minimization over $f \in \mathcal{F}_G$ can be performed pointwise at each x . For $f \in \mathcal{F}_G$ with $x \in G$, $f(x)$ is constrained to $[0, 1]$, and with $x \in G^c$, $f(x)$ is constrained to $[-1, 0)$. To simplify the notation, we define

$$\begin{aligned} C_\phi^+(\eta(x)) &\equiv \inf_{0 \leq f(x) \leq 1} C_\phi(f(x), \eta(x)), \\ C_\phi^-(\eta(x)) &\equiv \inf_{-1 \leq f(x) < 0} C_\phi(f(x), \eta(x)), \\ \Delta C_\phi(\eta(x)) &\equiv C_\phi^+(\eta(x)) - C_\phi^-(\eta(x)), \end{aligned}$$

where $C_\phi^+(\eta(x))$ and $C_\phi^-(\eta(x))$ are the minimized surrogate risks conditional on $X = x$ under the constraints $f(x) \in [0, 1]$ and $f(x) \in [-1, 0)$, respectively. Using these definitions, the surrogate risk at G can be written as

$$\begin{aligned} \mathcal{R}_\phi(G) &= \int_{\mathcal{X}} [C_\phi^+(\eta(x)) \cdot 1\{x \in G\} + C_\phi^-(\eta(x)) \cdot 1\{x \notin G\}] dP_X(x) \\ &= \int_G \Delta C_\phi(\eta(x)) \cdot 1\{x \in G\} dP_X(x) + \int_{\mathcal{X}} C_\phi^-(\eta(x)) dP_X(x). \quad (10) \end{aligned}$$

By comparing the expressions of the risks in (9) and (10), we obtain the first main theorem that clarifies the condition for the surrogate risk $\mathcal{R}_\phi(G)$ to calibrate the global ordering of the classification risk $\mathcal{R}(G)$ over $G \in \mathcal{G}$.

Theorem 3.2 (Global calibration of the \mathcal{G} -constrained excess risk). *Let P be an arbitrary distribution on $\{-1, 1\} \times \mathcal{X}$ and $\mathcal{G} \subseteq 2^{\mathcal{X}}$ be a class of measurable subsets of \mathcal{X} . For $G, G' \in \mathcal{G}$, the risk ordering $\mathcal{R}(G) \geq \mathcal{R}(G')$ in terms of the classification risk is equivalent to*

$$\int_{G \setminus G'} (1 - 2\eta(x)) dP_X(x) \geq \int_{G' \setminus G} (1 - 2\eta(x)) dP_X(x), \quad (11)$$

while the risk ordering $\mathcal{R}_\phi(G) \geq \mathcal{R}_\phi(G')$ in terms of the surrogate risk is equivalent to

$$\int_{G \setminus G'} \Delta C_\phi(\eta(x)) dP_X(x) \geq \int_{G' \setminus G} \Delta C_\phi(\eta(x)) dP_X(x). \quad (12)$$

Hence, if $\Delta C_\phi(\eta(x))$ is proportional to $1 - 2\eta(x)$ up to a positive constant, i.e.,

$$\Delta C_\phi(\eta(x)) = c(1 - 2\eta(x)) \text{ for some } c > 0, \quad (13)$$

the risk ordering over \mathcal{G} in terms of the surrogate risk $\mathcal{R}_\phi(G)$ agrees with the risk ordering over \mathcal{G} in terms of the classification risk $\mathcal{R}(G)$ for any distribution P on $\{-1, 1\} \times \mathcal{X}$.

In particular, when ϕ is the hinge loss $\phi_h(\alpha) = c \max\{0, 1 - \alpha\}$, $c > 0$,

$$\Delta C_\phi(\eta(x)) = c(1 - 2\eta(x))$$

holds, establishing that hinge risk preserves the risk ordering of the classification risk.

Proof. By equation (9),

$$\begin{aligned} \mathcal{R}(G) - \mathcal{R}(G') &= \int_{\mathcal{X}} (1 - 2\eta(x)) \cdot [1\{x \in G\} - 1\{x \in G'\}] dP_X(x) \\ &= \int_{\mathcal{X}} (1 - 2\eta(x)) \cdot [1\{x \in G \setminus G'\} - 1\{x \in G' \setminus G\}] dP_X(x) \\ &= \int_{G \setminus G'} (1 - 2\eta(x)) dP_X(x) - \int_{G' \setminus G} (1 - 2\eta(x)) dP_X(x). \end{aligned}$$

This proves (11), the first claim of the theorem.

Given the representation of the surrogate risk shown in (10), a similar argument yields (12), the second claim of the theorem.

For the hinge loss $\phi_h(\alpha) = c \max\{0, 1 - \alpha\}$ and $f \in \mathcal{F}_G$, we have

$$C_{\phi_h}(f(x), \eta(x)) = c(1 - 2\eta(x))f(x) + c.$$

Hence, we obtain

$$C_{\phi_h}^+(\eta) = \begin{cases} c(1 - 2\eta) + c & \text{for } \eta > 1/2, \\ c & \text{for } \eta \leq 1/2, \end{cases}$$

$$C_{\phi_h}^-(\eta) = \begin{cases} c & \text{for } \eta > 1/2, \\ 2c\eta & \text{for } \eta \leq 1/2. \end{cases}$$

Hence, $\Delta C_{\phi_h}(\eta) = c(1 - 2\eta)$ holds for all $\eta \in [0, 1]$. □

Theorem 3.2 does not exploit the condition that ϕ is classification-calibrated, but if a surrogate loss function satisfies condition (13), it is automatically classification-calibrated. Another remark follows.

Remark 3.3. *Many commonly used surrogate loss functions do not satisfy condition (13) in Theorem 3.2. Table 1 shows the forms of $\Delta C_\phi(\eta)$ for the hinge loss, exponential loss, logistic loss, quadratic loss, and truncated quadratic loss functions. With the exception of the hinge loss function, none of these functions satisfy condition (13). That is, among the surrogate loss-based algorithms that are commonly used in practice, the ℓ_1 -support vector machine corresponding to hinge loss is the only algorithm whose surrogate risk preserves the classification risk.*

Table 1: Surrogate loss functions and their associated forms for ΔC_ϕ

Loss function	$\phi(\alpha)$	$\Delta C_\phi(\eta)$
0-1 loss	$1\{\alpha \leq 0\}$	$1 - 2\eta$
Hinge loss	$c \max\{0, 1 - \alpha\}$	$c(1 - 2\eta)$
Exponential loss	$e^{-\alpha}$	$\begin{cases} -2\sqrt{\eta(1-\eta)} + 1 & \text{if } 0 \leq \eta < 1/2 \\ 2\sqrt{\eta(1-\eta)} - 1 & \text{if } 1/2 \leq \eta \leq 1 \end{cases}$
Logistic loss	$\log(1 + e^{-\alpha})$	$\begin{cases} \log(2\eta^\eta(1-\eta)^{1-\eta}) & \text{if } 0 \leq \eta < 1/2 \\ -\log(2\eta^\eta(1-\eta)^{1-\eta}) & \text{if } 1/2 \leq \eta \leq 1 \end{cases}$
Quadratic loss	$(1 - \alpha)^2$	$\begin{cases} (1 - 2\eta)^2 & \text{if } 0 \leq \eta < 1/2 \\ -(1 - 2\eta)^2 & \text{if } 1/2 \leq \eta \leq 1 \end{cases}$
Truncated quadratic loss	$(\max\{0, 1 - \alpha\})^2$	$\begin{cases} (1 - 2\eta)^2 & \text{if } 0 \leq \eta < 1/2 \\ -(1 - 2\eta)^2 & \text{if } 1/2 \leq \eta \leq 1 \end{cases}$

The well known inequality by Zhang (2004) relates the excess surrogate risk to the excess classification risk under R -correct specification. As a corollary of Theorem 3.2, if we set $\phi = \phi_h$, we can generalize Zhang's inequality by allowing R -misspecification of the classifiers. To formally state this generalization, we let $G^* \in \arg \inf_{G \in \mathcal{G}} \mathcal{R}(G)$, and set $G' = G^*$ in Theorem 3.2. Let $f \in \mathcal{F}_{\mathcal{G}}$ be arbitrary and $G_f = \{x \in \mathcal{X} : f(x) \geq 0\} \in \mathcal{G}$. The alignment of the risk ordering between the classification and hinge risks implies that the minimizers of $\mathcal{R}(\cdot)$ also minimize $\mathcal{R}_{\phi_h}(\cdot)$, i.e., $\inf_{f \in \mathcal{F}_{\mathcal{G}}} \mathcal{R}_{\phi_h}(f) = \inf_{G \in \mathcal{G}} \mathcal{R}_{\phi_h}(G) = \mathcal{R}_{\phi_h}(G^*)$. Theorem 3.2 therefore implies that the \mathcal{G} -constrained excess classification risk of f satisfies the following inequality:

$$\begin{aligned}
R(f) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f) &= \mathcal{R}(G_f) - \mathcal{R}(G^*) \\
&= \int_{G_f \setminus G^*} (1 - 2\eta(x)) dP_X(x) - \int_{G^* \setminus G_f} (1 - 2\eta(x)) dP_X(x) \\
&= c^{-1} [\mathcal{R}_{\phi_h}(G_f) - \mathcal{R}_{\phi_h}(G^*)]
\end{aligned}$$

$$\begin{aligned}
&= c^{-1} \left[\inf_{f' \in \mathcal{F}_{G_f}} R_{\phi_h}(f') - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R_{\phi_h}(f) \right] \\
&\leq c^{-1} \left[R_{\phi_h}(f) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R_{\phi_h}(f) \right], \tag{14}
\end{aligned}$$

where the second equality follows by equation (9); and the third equality follows by equation (10) and $\Delta C_{\phi_h}(\eta) = c(1 - 2\eta)$. That is, when $\phi = \phi_h$, Zhang's inequality holds without requiring the R -correct specification of the classifiers.

Corollary 3.4. *For any distribution P on $\{-1, 1\} \times \mathcal{X}$ and class of measurable subsets $\mathcal{G} \subseteq 2^{\mathcal{X}}$, if $\Delta C_{\phi}(\eta(x))$ is proportional to $1 - 2\eta(x)$ with a proportionality constant $c > 0$, i.e., $\Delta C_{\phi}(\eta(x)) = c(1 - 2\eta(x))$, then the following inequality holds*

$$c(R(f) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f)) \leq R_{\phi}(f) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R_{\phi}(f)$$

for any $f \in \mathcal{F}_{\mathcal{G}}$.

Proof. See equation (14). □

Corollary 3.4 shows that if the surrogate loss ϕ satisfies condition (13), then the classifier f_{ϕ}^* that minimizes the surrogate risk over $\mathcal{F}_{\mathcal{G}}$ also minimizes the classification risk over $\mathcal{F}_{\mathcal{G}}$. Importantly, this result holds without assuming the R -correct specification of $\mathcal{F}_{\mathcal{G}}$. It justifies the use of hinge loss in the constrained classification problem irrespective of whether or not $\mathcal{F}_{\mathcal{G}}$ is correctly R -specified. Note, however, that the result relies on the fact that at every $x \in \mathcal{X}$ we can choose any $f(x) \in [-1, 1]$ as long as the prediction set constraint $G_f \in \mathcal{G}$ is satisfied. We relax this requirement in the next section.

Further analysis can show that the condition (13) in Theorem 3.2 is not only sufficient but also necessary. To formally show this, we adopt the concept of universal equivalence of loss functions introduced by Nguyen et al. (2009) to the current setting.

Definition 3.5 (Universal equivalence). *Loss functions ϕ_1 and ϕ_2 are universally equivalent, denoted by $\phi_1 \stackrel{u}{\sim} \phi_2$, if for any distribution P on $\{-1, 1\} \times \mathcal{X}$ and class of measurable subsets $\mathcal{G} \subseteq 2^{\mathcal{X}}$,*

$$\mathcal{R}_{\phi_1}(G_1) \leq \mathcal{R}_{\phi_1}(G_2) \Leftrightarrow \mathcal{R}_{\phi_2}(G_1) \leq \mathcal{R}_{\phi_2}(G_2)$$

holds for any $G_1, G_2 \in \mathcal{G}$.

Universally equivalent loss functions ϕ_1 and ϕ_2 lead to the same risk ordering over \mathcal{G} . Hence, if a loss function ϕ is universally equivalent to the 0-1 loss, the ϕ -risk shares the same risk ordering with the classification risk.

The following theorem establishes a necessary and sufficient condition for two classification-calibrated loss functions to be universally equivalent.

Theorem 3.6. *Let ϕ_1 and ϕ_2 be classification-calibrated loss functions. Then $\phi_1 \stackrel{u}{\sim} \phi_2$ if and only if $\Delta C_{\phi_2}(\eta) = c\Delta C_{\phi_1}(\eta)$ for some $c > 0$ and any $\eta \in [0, 1]$, i.e., ΔC_{ϕ_2} is proportional to ΔC_{ϕ_1} up to a positive constant.*

Proof. See Appendix A. □

The ‘if’ part of the theorem is a generalization of Theorem 3.2 in that it does not assume that either of ϕ_1 or ϕ_2 is the 0-1 loss function.

When we set ϕ_2 to the 0-1 loss function, Theorem 3.6 yields the class of loss functions that are universally equivalent to the 0-1 loss functions. This class exactly coincides with the class of loss functions that satisfy the condition (13) in Theorem 3.2. Hence, the following corollary holds.

Corollary 3.7. *A classification-calibrated loss function ϕ is universally equivalent to the 0-1 loss function if and only if ϕ satisfies condition (13) for any $\eta(x) \in [0, 1]$. That is, the class of hinge loss functions $\{\phi(\alpha) = a \max\{0, 1 - \alpha\} + b : a > 0, b \geq 0\}$ agrees with the class of loss functions that are universally equivalent to the 0-1 loss function.*

In the following sections, without loss of generality, we maintain the assumption that $c = 1$ in the definition of the hinge loss function where it is convenient to do so. We conclude this section with a remark to compare our constrained classification framework to that of Nguyen et al. (2009).

Remark 3.8. *Nguyen et al. (2009) show that, for the classification problem in which an optimal pair comprising a quantizer and a classifier is to be chosen, the hinge loss function is also the only surrogate loss function that preserves the consistency of surrogate loss classification. In their framework, the quantizer is a stochastic mapping $Q \in \mathcal{Q} : \mathcal{X} \mapsto \mathcal{Z}$, where \mathcal{Z} is a discrete space and \mathcal{Q} is a possibly constrained class of conditional distributions of Z given X , $Q(Z | X)$. The classifier is a function $\gamma \in \Gamma : \mathcal{Z} \mapsto \mathbb{R}$, where Γ is the set of all measurable functions on \mathcal{Z} . The motivation for using Z as an input, instead of X , is to reduce the dimension of X , which might be a high-dimensional vector. Nguyen et al. (2009) propose estimating the pair $(Q, \gamma) \in \mathcal{Q} \times \Gamma$ that minimizes the risk $R(\gamma, Q) \equiv P(Y \neq \text{sign}(\gamma(Z)))$, by solving the surrogate loss classification problem: $\inf_{(Q, \gamma) \in \mathcal{Q} \times \Gamma} R_\phi(Q, \gamma)$, where $R_\phi(Q, \gamma) = E\phi(Y\gamma(Z))$. They show that, among the commonly used surrogate loss functions, only hinge loss classification leads to the optimal pair of (Q, γ) .*

The framework we study is different from that of Nguyen et al. (2009), and neither nests the other. The framework Nguyen et al. (2009) study constrains the mapping $Q : \mathcal{X} \mapsto \mathcal{Z}$, whereas the framework we study constrains prediction sets G_f for all classifiers f . Furthermore, the class of classifiers Γ considered in Nguyen et al. (2009) contains the Bayes classifier, whereas the class of classifiers $\mathcal{F}_{\mathcal{G}}$ we consider may not contain the Bayes classifier.

4 Consistency of hinge risk classification with functional form constraints

The previous section considers $\mathcal{F}_{\mathcal{G}}$, the class of all functions whose prediction sets are in \mathcal{G} . The generalized Zhang’s inequality shown in Corollary 3.4 heavily relies on the richness of $\mathcal{F}_{\mathcal{G}}$. This richness, however, limits the computational attractiveness of a surrogate-loss approach, since convexity in optimization of an empirical analogue of the surrogate risk does not directly follow from $\mathcal{F}_{\mathcal{G}}$, and typically requires additional functional form restrictions for f .

Unfortunately, once a functional form restriction on f is imposed on top of the prediction set constraint $G_f \in \mathcal{G}$, the global calibration property of the hinge risk shown in Theorem 3.2 breaks down. The following example illustrates this phenomenon.

Example 4.1 (Numerical example 2). *Maintain $\mathcal{X} = \{0, 1, 2\}$ and $\mathcal{G} = \{\emptyset, \{2\}, \{2, 1\}, \{2, 1, 0\}\}$ as in Example 3.1. We here consider choosing a classifier from the following class of non-decreasing linear functions:*

$$\mathcal{F}_L = \{f(x) = c_0 + c_1x : c_0 \in \mathbb{R}, c_1 \in \mathbb{R}_+, f(x) \in [-1, 1] \text{ for all } x \in \mathcal{X}\}.$$


Note that the class of prediction sets $\{G_f : f \in \mathcal{F}_L\}$ agrees with \mathcal{G} ; hence, \mathcal{F}_L is a subclass of $\mathcal{F}_{\mathcal{G}}$. We set X to be uniformly distributed on \mathcal{X} , and Y to have conditional probabilities $P(Y = 1 | X = 0) = 0.6$, $P(Y = 1 | X = 1) = 0.2$, and $P(Y = 1 | X = 2) = 0.8$.

The Bayes classifier predicts positive Y at $x = 0$ and 2 . Hence, no classifier in \mathcal{F}_L shares the prediction set with the Bayes classifier, and \mathcal{F}_L is R -misspecified.

Figure 2 illustrates the computed classifiers, f^ and $f_{\phi_h}^*$, that minimize the classification and hinge risks, respectively, over \mathcal{F}_L . The optimal classification risk $R(f^*)$ over \mathcal{F}_L (equivalently, over $\mathcal{F}_{\mathcal{G}}$ since $\{G_f : f \in \mathcal{F}_L\}$ agrees with \mathcal{G}) is $R(f^*) = 0.33$ with $G_{f^*} = \{2\}$, while the classification risk at $f_{\phi_h}^*$ is $R(f_{\phi_h}^*) = 0.54$ with $G_{f_{\phi_h}^*} = \{2, 1\}$. Thus, in contrast to Example 3.1 where f is unconstrained other than via the constraint $G_f \in \mathcal{G}$, adding the linear functional form constraint to $\mathcal{F}_{\mathcal{G}}$ invalidates the calibration property of*

the hinge risk, and the hinge risk minimization is no longer consistent to the second-best (constrained optimal) classifier in terms of the classification risk.

Figure 2: Linear monotone classifiers minimizing classification and hinge risks



Note: The orange and blue lines are the graphs of the computed classifiers, f^* and $f_{\phi_h}^*$, respectively.

This example illustrates that even with hinge loss, consistency to the second best classifier becomes a fragile property once the functional form of f is constrained in addition to the prediction set constraint $G_f \in \mathcal{G}$. Consequently, it is natural to ask what additional functional form restriction we can safely introduce to $\mathcal{F}_{\mathcal{G}}$ without threatening consistency, i.e., for which subclass $\tilde{\mathcal{F}}_{\mathcal{G}} \subset \mathcal{F}_{\mathcal{G}}$ does minimizing the hinge risk $R_{\phi_h}(f)$ over $f \in \tilde{\mathcal{F}}_{\mathcal{G}}$ lead to a classifier that minimizes the classification risk $R(f)$ over $f \in \mathcal{F}_{\mathcal{G}}$?

Formally, we introduce the following definition of *classification-preserving reduction* of $\mathcal{F}_{\mathcal{G}}$.

Definition 4.2 (Classification-preserving reduction). Let $\tilde{f}^* \in \arg \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}(f)$. A subclass of classifiers $\tilde{\mathcal{F}}_{\mathcal{G}} (\subseteq \mathcal{F}_{\mathcal{G}})$ is a *classification-preserving reduction* of $\mathcal{F}_{\mathcal{G}}$ if

$$R(\tilde{f}^*) = \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f)$$

holds for any P , distribution on $\{-1, 1\} \times \mathcal{X}$.

To start with the heuristic, consider a simple case where $\tilde{\mathcal{F}}_{\mathcal{G}}$ consists of piecewise

constant functions with at most J jumps, $J \geq 1$, of the following form:

$$\tilde{\mathcal{F}}_{\mathcal{G},J} = \left\{ f(\cdot) = 2 \sum_{j=1}^J c_j 1\{\cdot \in G_j\} - 1 : G_j \in \mathcal{G} \text{ and } c_j \geq 0 \text{ for } j = 1, \dots, J; \right. \\ \left. G_J \subseteq \dots \subseteq G_1; \sum_{j=1}^J c_j = 1 \right\}. \quad (15)$$

By construction, any function in $\tilde{\mathcal{F}}_{\mathcal{G},J}$ is a step function bounded in $[-1, 1]$ and its sublevel sets $\{x \in \mathcal{X} : f(x) \leq t\}$ belong to \mathcal{G} for any $t \in [-1, 1]$.

Let

$$\mathcal{G}^* \equiv \arg \inf_{G \in \mathcal{G}} \mathcal{R}(G)$$

be the collection of best prediction sets in \mathcal{G} , and

$$\mathcal{R}^* \equiv \inf_{G \in \mathcal{G}} \mathcal{R}(G)$$

be the optimal classification risk. For any $G \in \mathcal{G}$, we define $\tilde{f}_G(x) \equiv 2 \cdot 1\{x \in G\} - 1$, a step function over \mathcal{X} that indicates $x \in G^*$ and $x \notin G^*$ with values $+1$ and -1 , respectively. The following lemma shows that $\tilde{\mathcal{F}}_{\mathcal{G},J}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$.

Lemma 4.3. *Let $\mathcal{G} \subseteq 2^{\mathcal{X}}$ be a class of measurable subsets of \mathcal{X} . The following two claims hold:*

- (i) $\tilde{\mathcal{F}}_{\mathcal{G},J}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$.
- (ii) For any distribution P on $\{-1, 1\} \times \mathcal{X}$ and $G^* \in \mathcal{G}^*$, \tilde{f}_{G^*} is a minimizer of $R_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G},J}$, and $\inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G},J}} R_{\phi_h}(f) = 2\mathcal{R}^*$ holds.

Proof. See Appendix A. □

Characteristic features of $\tilde{\mathcal{F}}_{\mathcal{G},J}$ are (i) sublevel sets of any $f \in \tilde{\mathcal{F}}_{\mathcal{G},J}$ are in \mathcal{G} , and (ii) $\tilde{\mathcal{F}}_{\mathcal{G},J}$ contains \tilde{f}_G for any $G \in \mathcal{G}$. It transpires that these two features are the key features that need to be maintained for $\tilde{\mathcal{F}}_{\mathcal{G}}$ to generalize Lemma 4.3.

The next theorem is the second main theorem of the paper that extends Lemma 4.3 to a more general class of classifiers that can accommodate continuous ones.

Theorem 4.4 (Consistency under classification-preserving reduction). *Given a class of measurable subsets $\mathcal{G} \subseteq 2^{\mathcal{X}}$ and $\mathcal{F}_{\mathcal{G}} = \{f : G_f \in \mathcal{G}, f(\cdot) \in [-1, 1]\}$, suppose $\tilde{\mathcal{F}}_{\mathcal{G}} \subset \mathcal{F}_{\mathcal{G}}$ satisfies the following two conditions:*

- (A1) For every $f \in \tilde{\mathcal{F}}_{\mathcal{G}}$, $\{x \in \mathcal{X} : f(x) \leq t\} \in \mathcal{G}$ for all $t \in [-1, 1]$;

(A2) For any $G \in \mathcal{G}$, $\tilde{f}_G \in \tilde{\mathcal{F}}_G$.

Then the following claims hold:

(i) $\tilde{\mathcal{F}}_G$ is a classification-preserving reduction of \mathcal{F}_G ;

(ii) For any distribution P on $\{-1, 1\} \times \mathcal{X}$ and $G^* \in \mathcal{G}^*$, \tilde{f}_{G^*} is a minimizer of $R_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_G$, and $\inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}(f) = 2\mathcal{R}^*$ holds.

Proof. See Appendix A. □

The theorem establishes that the two conditions (A1) and (A2) are sufficient for $\tilde{\mathcal{F}}_G$ to be a classification-preserving reduction of \mathcal{F}_G . This result holds regardless of whether \mathcal{F}_G is correctly R -specified or not. Examples 4.6 and 4.7 at the end of this section give examples of classification-preserving reductions for linear classification and monotone classification.

The conditions (A1) and (A2) in Theorem 4.4 are simple to interpret and guarantee the consistency of the hinge risk minimization, but they do not imply that the empirical hinge risk minimization over $\tilde{\mathcal{F}}_G$ can be reduced to a convex optimization. We are unaware of a general way to construct a classification-preserving reduction that makes the empirical hinge risk minimization a convex program. For monotone classification, analyzed in Section 6, we propose two constructions of $\tilde{\mathcal{F}}_{G_M}$, one of which is exactly a classification-preserving reduction of \mathcal{F}_{G_M} while the other is approximately classification-preserving. We show that for both cases, minimization of the empirical hinge risk is a linear programming problem.

Although Theorem 4.4 shows the consistency of the hinge risk minimization over $\tilde{\mathcal{F}}_G$, it does not lead to the generalized Zhang's (2004) inequality in Corollary 3.4. Instead, the following corollary gives proportional equality between the \mathcal{G} -constrained excess classification risk and the \mathcal{F}_G -constrained excess hinge risk with an extra term added.

Corollary 4.5. *Assume $\tilde{\mathcal{F}}_G$ is a subclass of \mathcal{F}_G satisfying conditions (A1) and (A2) in Theorem 4.4. If $\Delta C_\phi(\eta) = c(1 - 2\eta)$ holds for some $c > 0$ and any $\eta \in [0, 1]$,*

$$c(R(f) - \inf_{f \in \tilde{\mathcal{F}}_G} R(f)) = \frac{1}{2} \left(R_\phi(f) - \inf_{f \in \tilde{\mathcal{F}}_G} R_\phi(f) \right) + \frac{1}{2} \left(R_\phi(\tilde{f}_{G_f}) - R_\phi(f) \right) \quad (16)$$

for any classifier $f : \mathcal{X} \mapsto [-1, 1]$. Moreover, the following holds:

$$c(R(f) - \inf_{f \in \mathcal{F}_G} R(f)) \leq \frac{1}{2} \left(R_\phi(f) - \inf_{f \in \mathcal{F}_G} R_\phi(f) \right) + \frac{1}{2} \left(R_\phi(f) - \inf_{f \in \mathcal{F}_G} R_\phi(f) \right) \quad (17)$$

for any $f \in \mathcal{F}_G$.

Proof. See Appendix A. □

The extra term (the right-most term) in (16) measures the difference in the hinge risks between a classifier f and the step function indicating the prediction set of f by the values $+1$ or -1 . Due to the fact that some of the best classifiers are of the form $\tilde{f}^*(\cdot) = 1\{\cdot \in G^*\} - 1\{\cdot \notin G^*\}$ for $G^* \in \mathcal{G}^*$ (Theorem 4.4 (ii)), if f closely approximates such a classifier, the extra term is close to zero. In the following section, we use equation (16) to derive the statistical properties of the hinge risk minimization in terms of the \mathcal{G} -constrained excess classification risk. Equation (17) implies that the \mathcal{G} -constrained excess classification risk is bounded from above by the average of the two $\mathcal{F}_{\mathcal{G}}$ -constrained excess hinge risks. One is over $\tilde{\mathcal{F}}_{\mathcal{G}}$ and the other is over $\mathcal{F}_{\mathcal{G}}$. We are unable to determine if the excess hinge risk over $\mathcal{F}_{\mathcal{G}}$ can be bounded from above by a term that is proportional to the excess hinge risk over $\tilde{\mathcal{F}}_{\mathcal{G}}$. As such, the constrained-classification-preserving reduction $\tilde{\mathcal{F}}_{\mathcal{G}}$ cannot replace $\mathcal{F}_{\mathcal{G}}$ in Zhang's inequality, shown in Corollary 3.4.

We conclude this section by presenting examples of classes of classifiers that approximately or exactly satisfy the conditions for classification-preserving reduction.

Example 4.6 (Linear classification with a class of transformed logistic functions). *Suppose that the prediction sets are subject to the linear index rules:*

$$\mathcal{G}_L = \{x \in \mathbb{R}^{d_x} : x^T \beta \geq 0 : \beta \in \mathbb{R}^{d_x}\},$$

where $\mathcal{X} = \mathbb{R}^{d_x}$. Let $\pi(\alpha, k) \equiv (1 - e^{-k\alpha}) / (1 + e^{-k\alpha}) = 2 / (1 + e^{-k\alpha}) - 1$ be a transformed logistic function and define a class of classifiers

$$\mathcal{F}_{\text{Logit}} = \{\pi(x^T \beta, k) : \beta \in \mathbb{R}^{d_x} \text{ and } k \in \mathbb{R}_+\},$$

where k is a tuning parameter that determines the steepness of the logistic curve. $\mathcal{F}_{\text{Logit}}$ satisfies the condition (A1) in Theorem 4.4.⁴ Since $\mathcal{F}_{\text{Logit}}$ for fixed $k < \infty$ rules out any step functions, the condition (A2) in Theorem 4.4 is not exactly met. Fix $G \in \mathcal{G}$, and let $\tilde{\beta}$ be such that $\{x \in \mathcal{X} : x^T \tilde{\beta} \geq 0\} = G$. Then, as $k \rightarrow \infty$, $\pi(x^T \tilde{\beta}, k)$ approximates $\text{sign}(x^T \tilde{\beta})$, so the condition (A2) is approximately met for large k . Every function in $\mathcal{F}_{\text{Logit}}$ is smooth and depends on a finite number of parameters. Hence, the empirical hinge risk becomes a smooth and continuous function with a finite number of parameters, although it is not generally convex.

Example 4.7 (Monotonic classification with a class of monotone functions). *Hinge risk minimization embedding a monotonicity restriction remains consistent when we use a*

⁴Fix $\beta \in \mathbb{R}^{d_x}$ and $k \in \mathbb{R}_+$. The condition (A1) is satisfied as, for any $t \in [-1, 1]$, $\{x : \pi(x^T \beta, k) \leq t\} = \{x : x^T \beta \leq \pi^{-1}(t, k)\} \in \mathcal{G}$, where $\pi^{-1}(\cdot, k)$ is an inverse function of $\pi(\cdot, k)$ with fixed k .

class of monotone functions. Let \preceq be a partial order on \mathcal{X} , and let \mathcal{G}_{\preceq} be the collection of all $G \in 2^{\mathcal{X}}$ that respect monotonicity (i.e., if $x_1 \preceq x_2$ and $x_1 \in G$, then $x_2 \in G$). Define $\tilde{\mathcal{F}}_{\mathcal{G}_{\preceq}}$ as a class of functions $f : \mathcal{X} \rightarrow [-1, 1]$ that are weakly monotonic in \preceq (i.e., satisfying $f(x_1) \leq f(x_2)$ if $x_1 \preceq x_2$). Then the prediction set of any $f \in \tilde{\mathcal{F}}_{\mathcal{G}_{\preceq}}$ respects the partial order \preceq (i.e., if $x_1 \preceq x_2$ and $x_1 \in G_f$, then $x_2 \in G_f$). For any $t \in [-1, 1]$ and $f \in \tilde{\mathcal{F}}_{\preceq}$, $\{x : f(x) \leq t\} = \{x : x \preceq \tilde{x} \text{ for any } \tilde{x} \text{ such that } f(\tilde{x}) = t\} \in \mathcal{G}_{\preceq}$ holds, satisfying condition (A1) in Theorem 4.4. In addition, for any $G \in \mathcal{G}_{\preceq}$, since $\tilde{f}_G(x) = 2 \cdot 1\{x \in G\}$ is weakly monotonic in \preceq , $\tilde{f}_G \in \mathcal{G}_{\preceq}$ holds, satisfying condition (A2) in Theorem 4.4. Hence $\tilde{\mathcal{F}}_{\mathcal{G}_{\preceq}}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}_{\preceq}}$. Therefore, according to Theorem 4.4, hinge risk minimization over $\tilde{\mathcal{F}}_{\preceq}$ yields the optimal classifier in terms of the classification risk. Section 6 focuses on monotone classification and investigates its statistical and computational properties.

5 Statistical properties

The analyses presented so far concern the consistency of a surrogate loss approach in terms of the population risk criterion. It is important to note that Theorems 3.2 and 4.4 do not impose any restriction on the underlying distribution of (Y, X) . Accordingly, equivalence of the risk orderings and risk minimizing classifiers between the classification and hinge risks remains valid even if we consider empirical analogues of the risks constructed from the empirical distribution of the sample. Theorems 3.2 and 4.4 hence guarantee that a classifier minimizing the empirical hinge risk over $\mathcal{F}_{\mathcal{G}}$ or over a classification-preserving reduction $\tilde{\mathcal{F}}_{\mathcal{G}}$ also minimizes the empirical classification risk.

In this section, we assess the generalization performance of hinge risk minimizing classifiers, allowing for general misspecification of the constrained class of classifiers. Towards that goal, let \mathcal{G} be fixed and consider $\check{\mathcal{F}}$, a class of classifiers whose members satisfy $-1 \leq f \leq 1$. $\check{\mathcal{F}}$ may or may not be a subclass of $\mathcal{F}_{\mathcal{G}}$, while in our analysis of monotone classification below, $\check{\mathcal{F}}$ corresponds to an approximation of a classification-preserving reduction $\tilde{\mathcal{F}}_{\mathcal{G}}$. Let $\{(Y_i, X_i) : i = 1, \dots, n\}$ be a sample of observations that are independent and identically distributed (i.i.d.) as (Y, X) . We denote the joint distribution of a sample of n observations by P^n and the expectation with respect to P^n by $E_{P^n}[\cdot]$. We define the empirical classification risk and empirical hinge risk, respectively, as

$$\begin{aligned} \hat{R}(f) &\equiv \frac{1}{n} \sum_{i=1}^n 1\{Y_i \cdot \text{sign}(f(X_i)) \leq 0\}, \\ \hat{R}_{\phi_h}(f) &\equiv \frac{1}{n} \sum_{i=1}^n \max\{0, 1 - Y_i f(X_i)\} = \frac{1}{n} \sum_{i=1}^n (1 - Y_i f(X_i)), \end{aligned}$$

where the max operator in the hinge loss is redundant if we constrain $f(\cdot)$ to $[-1, 1]$. Let \hat{f} be a classifier that minimizes $\hat{R}_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}$. We evaluate the statistical properties of \hat{f} in terms of the excess classification risk relative to the minimal risk over $\mathcal{F}_{\mathcal{G}}$, $R(\hat{f}) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f)$. In particular, we later derive a distribution-free upper bound on the mean of the excess classification risk.

Let $\tilde{\mathcal{F}}_{\mathcal{G}}$ be a subclass of $\mathcal{F}_{\mathcal{G}}$ that satisfies conditions (A1) and (A2) in Theorem 4.4. $\tilde{\mathcal{F}}_{\mathcal{G}}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$ (Definition 4.2). Following Corollary 4.5, we have

$$R(\hat{f}) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f) = \frac{1}{2} \left(R_{\phi_h}(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}(f) \right) + \frac{1}{2} \left(R_{\phi_h}(\tilde{f}_{\mathcal{G}_{\hat{f}}}) - R_{\phi_h}(\hat{f}) \right). \quad (18)$$

When $\tilde{\mathcal{F}} = \tilde{\mathcal{F}}_{\mathcal{G}}$, evaluating each term on the right hand side of (18) gives an upper bound on the mean of the \mathcal{G} -constrained excess classification risk of \hat{f} .

Let $H_1^B(\epsilon, \mathcal{F}, P_X)$ be the $L_1(P_X)$ -bracketing entropy of a class of functions \mathcal{F} and $H_1^B(\epsilon, \mathcal{G}, P_X)$ be that of a class of prediction sets \mathcal{G} .⁵ For definitions of these two terms, see Definition B.1 in Appendix B. When $\tilde{\mathcal{F}}$ coincides with $\tilde{\mathcal{F}}_{\mathcal{G}}$, the following theorem gives a non-asymptotic distribution-free upper bound on the mean of the \mathcal{G} -constrained excess classification risk in terms of the bracketing entropy.

Theorem 5.1. *Let $\tilde{\mathcal{F}}_{\mathcal{G}}$ be a subclass of $\mathcal{F}_{\mathcal{G}}$ that satisfies conditions (A1) and (A2) in Theorem 4.4. Suppose that \mathcal{P} is a class of distributions on $\{-1, 1\} \times \mathcal{X}$ such that there exist positive constants C and r for which*

$$H_1^B(\epsilon, \mathcal{G}, P_X) \leq C\epsilon^{-r} \quad (19)$$

holds for any $P \in \mathcal{P}$ and $\epsilon > 0$, or

$$H_1^B(\epsilon, \tilde{\mathcal{F}}_{\mathcal{G}}, P_X) \leq C\epsilon^{-r} \quad (20)$$

holds for any $P \in \mathcal{P}$ and $\epsilon > 0$. Define $\tau_n = n^{-1/2}$ if $r < 1$, $\tau_n = \log(n) / \sqrt{n}$ if $r = 1$, and $\tau_n = n^{-1/(r+1)}$ if $r \geq 2$. Let $q_n = \sqrt{n}\tau_n$. Then, for $\hat{f} \in \arg \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}(f)$, the following holds:

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[R(\hat{f}) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f) \right] \leq L_C(r, n), \quad (21)$$

⁵With a slight abuse of notation, we denote by $H_1^B(\epsilon, \mathcal{G}, P_X)$ the bracketing entropy number of the class of indicator functions, $H_1^B(\epsilon, \mathcal{H}_{\mathcal{G}}, P_X)$, where $\mathcal{H}_{\mathcal{G}} \equiv \{1\{\cdot \in G\} : G \in \mathcal{G}\}$.

where

$$L_C(r, n) = \begin{cases} 2D_1\tau_n + 4D_2 \exp(-D_1^2 q_n^2) & \text{if } r \geq 1 \\ 2D_3\tau_n + 2n^{-1}D_4 & \text{if } r < 1 \end{cases}$$

for some positive constants D_1, D_2, D_3, D_4 , which depend only on C and r .

Proof. See Appendix B. □

The upper bound on the mean of the \mathcal{G} -constrained excess classification risk converges to zero at the rate of τ_n , which depends on r in the bracketing entropy conditions (19) and (20). Dudley (1999) shows many examples that satisfy these bracketing entropy conditions. In particular, the class $\mathcal{G}_{\succ} \subseteq 2^{\mathcal{X}}$ that is compatible with monotone classification and introduced in Example 4.7 satisfies condition (19) with r equal to $d_x - 1$ (see Theorem 8.3.2 in Dudley (1999)).

We next consider the case when $\tilde{\mathcal{F}}$ does not coincide with $\tilde{\mathcal{F}}_{\mathcal{G}}$. This case corresponds to a scenario where minimizing the empirical hinge risk over $\tilde{\mathcal{F}}_{\mathcal{G}}$ is difficult but minimizing over $\tilde{\mathcal{F}}$, a class approximating $\tilde{\mathcal{F}}_{\mathcal{G}}$, is feasible.

A further decomposition of $R_{\phi_h}(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}(f)$ in (18) leads to

$$\begin{aligned} R(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R(f) &= \frac{1}{2} \left(R_{\phi_h}(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}(f) \right) + \frac{1}{2} \left(\inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}(f) - \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}(f) \right) \\ &\quad + \frac{1}{2} \left(R_{\phi_h}(\tilde{f}_{\mathcal{G}_{\tilde{f}}}) - R_{\phi_h}(\hat{f}) \right). \end{aligned} \quad (22)$$

Hence the \mathcal{G} -constrained excess classification risk is decomposed into three terms. We call the first term estimation error, the second term approximation error to a best classifier, and the third term approximation error to a step classifier. Evaluating each error gives an upper bound on the \mathcal{G} -constrained excess classification risk.

The following theorem evaluates the estimation error in terms of bracketing entropy.

Theorem 5.2. *Let $\tilde{\mathcal{F}}_{\mathcal{G}}$ be a subclass of $\mathcal{F}_{\mathcal{G}}$ that satisfies conditions (A1) and (A2) in Theorem 4.4. Suppose that \mathcal{P} is a class of distributions on $\{-1, 1\} \times \mathcal{X}$ such that there exist positive constants C' and r' for which*

$$H_1^B(\epsilon, \tilde{\mathcal{F}}, P_X) \leq C' \epsilon^{-r'} \quad (23)$$

holds for any $P \in \mathcal{P}$ and $\epsilon > 0$. Let $\hat{f} \in \arg \inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}(f)$. Then

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[R(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R(f) \right] \leq L_{C'}(r', n) + \frac{1}{2} \left(\inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}(f) - \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}(f) \right)$$

$$+ \frac{1}{2} \left(R_{\phi_h}(\tilde{f}_{G_{\tilde{f}}}) - R_{\phi_h}(\hat{f}) \right), \quad (24)$$

where $L_{C'}(r', n)$ is as defined in Theorem 5.1.

Proof. See Appendix B. □

Remark 5.3 (Approximation errors). *Evaluating each approximation error (the final two terms on the right-hand side) in (24) depends on the functional form restriction placed on $f \in \tilde{\mathcal{F}}$. If $\tilde{\mathcal{F}}$ grows and approaches $\tilde{\mathcal{F}}_G$ as $n \rightarrow \infty$, each approximation error converges to zero. In Section 6.2 below, where we consider the monotone classification problem and set $\tilde{\mathcal{F}}$ being a sieve of Bernstein polynomials, we characterize convergence of these two approximation errors. We then apply Theorem 5.2 to obtain the regret convergence rate of the estimated monotone classifier.*

6 Applications to monotone classification

This section applies the general theoretical results shown in Sections 3–5 to the monotone classification problem (Example 2.5). By Theorem 3.2, we limit our analysis to hinge loss. We assume that \mathcal{X} is compact in \mathbb{R}^{d_x} , $d_x < \infty$, and without loss of generality, we represent \mathcal{X} as the d_x -dimensional unit hypercube (i.e., $\mathcal{X} = [0, 1]^{d_x}$). To be specific, we consider the class of monotone prediction sets \mathcal{G}_M such that, for any $G \in \mathcal{G}_M$ and $x, \tilde{x} \in \mathcal{X}$, $x \in G$ and $x \leq \tilde{x}$ implies $\tilde{x} \in G$ holds⁶ (i.e., \mathcal{G}_M respects the partial order \leq on \mathcal{X}). Accordingly, the class of monotonically increasing classifiers can be represented as

$$\mathcal{F}_M \equiv \{f : f(x) \leq f(\tilde{x}) \text{ for any } x, \tilde{x} \in \mathcal{X} \text{ with } x \leq \tilde{x} ; f(\cdot) \in [-1, 1]\}.$$

In this section, we first study the monotone classification problem on \mathcal{F}_M . Note that \mathcal{F}_M is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}_M}$ (see Example 4.7). As an alternative to \mathcal{F}_M , we next consider using a sieve of Bernstein polynomials to approximate a hinge risk minimizing classifier on \mathcal{F}_M . The Bernstein polynomial is known for its capability to accommodate bound constraints and various shape constraints on functions (e.g., monotonicity or convexity). The class of Bernstein polynomials becomes a classification-preserving reduction only at the limit with a growing order of polynomials.

⁶We define the partial order \leq on \mathcal{X} as follows. For any $x = (x_1, \dots, x_d)^T$ and $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_d)^T$, we say $x \leq \tilde{x}$ if $x_j \leq \tilde{x}_j$ for every $j = 1, \dots, d$. We further say $x < \tilde{x}$ if $x \leq \tilde{x}$ holds and for some $j \in \{1, \dots, d\}$, $x_j < \tilde{x}_j$ holds.

6.1 Nonparametric monotone classification

We first consider hinge risk minimization given the class of monotonically increasing classifiers \mathcal{F}_M . Let \hat{f}_M be a minimizer of $\hat{R}_{\phi_h}(\cdot)$ over \mathcal{F}_M . Since the hinge risk for classifiers constrained to $-1 \leq f(x) \leq 1$ gives the linear loss $\phi_h(yf(x)) = 1 - yf(x)$, minimization of the empirical hinge risk can be formulated as the following linear programming:

$$\begin{aligned} & \max_{f_1, \dots, f_n} \sum_{i=1}^n Y_i f_i & (25) \\ \text{s.t. } & f_i \geq f_j \text{ for any } X_i \neq X_j \text{ with } X_i \geq X_j \text{ for } 1 \leq i \leq j \leq n; \\ & -1 \leq f_i \leq 1 \text{ for } 1 \leq i \leq n, \end{aligned}$$

where the first inequality constraints correspond to the monotonicity constraint on \mathcal{F}_M , and the second inequality constraints correspond to the range constraint for $f \in \mathcal{F}_M$. Solving this linear program yields the values of \hat{f}_M at the values of x observed in the training sample. Let $(\hat{f}_M(X_1), \dots, \hat{f}_M(X_n))$ be the solution of (25). Then any function in \mathcal{F}_M that passes the points $((X_1, \hat{f}_M(X_1)), \dots, (X_n, \hat{f}_M(X_n)))$ minimizes the empirical hinge risk over \mathcal{F}_M .⁷ Since \mathcal{F}_M is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}_M}$, Theorem 4.4 with P replaced by P^n shows that any solution to (25) exactly minimizes $\hat{R}_{\phi_h}(\cdot)$ over $\mathcal{F}_{\mathcal{G}_M}$.

We investigate the statistical properties of this procedure. Since \mathcal{F}_M is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}_M}$, we can apply Theorem 5.1. Towards this goal, we first characterize an upper bound on the bracketing entropy number of the class of monotone prediction sets \mathcal{G}_M . The next lemma, which we borrow from Theorem 8.3.2 in Dudley (1999), gives an upper bound on the $L_1(P_X)$ -bracketing entropy of \mathcal{G}_M . Here, we assume that X is continuously distributed with bounded density.

Lemma 6.1. *Suppose that P_X is absolutely continuous with respect to the Lebesgue measure on \mathcal{X} and has a density that is bounded from above by a finite constant $A > 0$. Then there exists a constant C , which depends only on A , such that*

$$H_1^B(\epsilon, \mathcal{G}_M, P_X) \leq C\epsilon^{1-d_X}.$$

holds for all $\epsilon > 0$.

⁷All classifiers obtained from this procedure predict a unique label at each point x observed in the training sample, whereas they may not give a unique prediction at a point x not observed in the training sample. One possible way to predict a label at an unobserved point x without violating the monotonicity constraint is to predict its label by the largest label among those predicted by all classifiers in $\arg \inf_{f \in \mathcal{F}_M} \hat{R}_{\phi_h}(f)$. Let $\tilde{\mathcal{X}}$ be a set of x observed in the training sample. Given any $\hat{f}_M \in \arg \inf_{f \in \mathcal{F}_M} \hat{R}_{\phi_h}(f)$, this way of predicting a label is equivalent to predicting the label of $x \in \mathcal{X} \setminus \tilde{\mathcal{X}}$ as the sign of $\min\{\hat{f}_M(\tilde{x}) : \tilde{x} \in \tilde{\mathcal{X}}, \tilde{x} \geq x\}$ if there exists $\tilde{x} \in \tilde{\mathcal{X}}$ such that $\tilde{x} \geq x$, and as 1 otherwise.

Proof. See Appendix C. □

With this lemma to hand, setting $r = 1 - d_x$ in Theorem 5.1 yields a finite sample uniform upper bound on the \mathcal{G}_M -constrained excess classification risk of \hat{f}_M . The following theorem shows that the excess risk of \hat{f}_M obtained from the linear program in (25) attains the same convergence rate as the welfare regret of monotone treatment rules shown by Mbakop and Tabord-Meehan (2021).

Theorem 6.2. *Let \mathcal{P} be a class of distributions on $\{-1, 1\} \times \mathcal{X}$ such that the marginal distribution P_X is absolutely continuous with respect to the Lebesgue measure on \mathcal{X} and has a density that is bounded from above by some finite constant $A > 0$. Define $\tau_n = n^{-1/2}$ if $d_x = 1$, $\tau_n = \log(n)/\sqrt{n}$ if $d_x = 2$, and $\tau_n = n^{-1/d_x}$ if $d_x \geq 3$. Let $q_n = \sqrt{n}\tau_n$. Then, for $\hat{f}_M \in \arg \inf_{f \in \mathcal{F}_M} \hat{R}_{\phi_h}(f)$,*

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[R(\hat{f}_M) - \inf_{f \in \mathcal{F}_{\mathcal{G}_M}} R(f) \right] \leq \begin{cases} 2D_1\tau_n + 4D_2 \exp(-D_1^2 q_n^2) & \text{if } d_x \geq 2 \\ 2D_3\tau_n + 2n^{-1}D_4 & \text{if } d_x = 1 \end{cases}$$

for some positive constants D_1, D_2, D_3, D_4 , which depend only on d_x and A .

Proof. Since \mathcal{F}_M satisfies conditions (A1) and (A2) in Theorem 4.4 with \mathcal{G} equal to \mathcal{G}_M (Example 4.7), the result follows from Theorem 5.1 and Lemma 6.1. □

This theorem guarantees the consistency of monotone classification using hinge loss and the class of monotone classifiers \mathcal{F}_M . The rate of convergence corresponds to τ_n .

6.2 Monotone classification with Bernstein polynomials

To illustrate our theoretical results for monotone classification, the second approach we consider is to use multivariate Bernstein polynomials to approximate a best classifier in \mathcal{F}_M . Let $b_{kj}(x) = \binom{k}{j} x^j (1-x)^{k-j}$ be the Bernstein basis. The Bernstein polynomial for a d_x -dimensional function takes the following form:

$$B_{\mathbf{k}}(\theta, x) = \sum_{j_1=0}^{k_1} \cdots \sum_{j_{d_x}=0}^{k_{d_x}} \theta_{j_1 \dots j_{d_x}} \cdot (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x})),$$

where $\mathbf{k} = (k_1, \dots, k_{d_x})^T$ is a vector collecting the orders of the Bernstein polynomial bases specified by the analyst, $\theta \equiv \{\theta_{j_1 \dots j_{d_x}}\}_{j_1=0, \dots, k_1; \dots; j_{d_x}=0, \dots, k_{d_x}}$ is a $(k_1 + 1) \times \cdots \times (k_{d_x} + 1)$ -dimensional vector of the parameters to be estimated, and x_j denotes the j -th element

of the d_x -dimensional vector x . If $-1 \leq \theta_{j_1 \dots j_{d_x}} \leq 1$ for all (j_1, \dots, j_{d_x}) , the range of the function $B_{\mathbf{k}}(\theta, \cdot)$ is bounded in $[-1, 1]$. Moreover, if $\theta_{j_1 \dots j_{d_x}} \geq \theta_{\tilde{j}_1 \dots \tilde{j}_{d_x}}$ for all $(j_1, \dots, j_{d_x}) \geq (\tilde{j}_1, \dots, \tilde{j}_{d_x})$, $B_{\mathbf{k}}(\theta, \cdot)$ is non-decreasing in x .⁸ Hence, to preserve the bound and non-decreasing constraints on \mathcal{F}_M , we constrain the class of Bernstein polynomials to

$$\mathcal{B}_{\mathbf{k}} = \left\{ B_{\mathbf{k}}(\theta, \cdot) : \theta \in \tilde{\Theta} \right\},$$

where $\tilde{\Theta}$ is the set of θ such that $\theta_{j_1 \dots j_{d_x}} \in [-1, 1]$ for all (j_1, \dots, j_{d_x}) and $\theta_{j_1 \dots j_{d_x}} \geq \theta_{\tilde{j}_1 \dots \tilde{j}_{d_x}}$ for all $(j_1, \dots, j_{d_x}) \geq (\tilde{j}_1, \dots, \tilde{j}_{d_x})$. An appropriate choice of \mathbf{k} is discussed later.

Noting that some hinge risk minimizing classifiers on \mathcal{F}_M have the form of step functions taking only the values -1 and 1 (Theorem 4.4), we propose approximating such a step function using the sieve of Bernstein polynomials. To this end, we propose the following two steps:

1. Minimize the empirical hinge risk $\hat{R}_{\phi_h}(f)$ over $f \in \mathcal{B}_{\mathbf{k}}$ and obtain $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_{\mathbf{k}}} \hat{R}_{\phi_h}(f)$.
2. Let $\{\hat{\theta}_{j_1 \dots j_{d_x}}\}_{j_1=0, \dots, k_1; \dots; j_{d_x}=0, \dots, k_{d_x}}$ be the vector of coefficients in \hat{f}_B . Compute a modified classifier

$$\hat{f}_B^\dagger(x) \equiv \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \text{sign}(\hat{\theta}_{j_1 \dots j_{d_x}}) \cdot (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x})),$$

which converts each estimated coefficient $\hat{\theta}_{j_1 \dots j_{d_x}}$ to either -1 or 1 depending on its sign.

Our proposal is to use \hat{f}_B^\dagger rather than \hat{f}_B . Lemma C.3 in Appendix C shows that \hat{f}_B^\dagger also minimizes $R_{\phi_h}(f)$ over $f \in \mathcal{B}_{\mathbf{k}}$. With respect to the first step, since the hinge loss of a classifier f constrained on $[-1, 1]$ has the linear form $\phi_h(yf(x)) = 1 - yf(x)$, any function in $\mathcal{B}_{\mathbf{k}}$ is linear in the parameters θ , and the parameter space $\tilde{\Theta}$ is a polyhedron, minimization of $\hat{R}_{\phi_h}(\cdot)$ over $\mathcal{B}_{\mathbf{k}}$ can be formulated as the following linear program:

$$\begin{aligned} & \max_{\theta} \sum_{i=1}^n Y_i \cdot \left(\sum_{j_1=0}^{k_1} \cdots \sum_{j_{d_x}=0}^{k_{d_x}} \theta_{j_1 \dots j_{d_x}} \cdot (b_{k_1 j_1}(X_{i1}) \times \cdots \times b_{k_{d_x} j_{d_x}}(X_{id_x})) \right) \\ & \text{s.t. } \theta_{j_1 \dots j_{d_x}} \geq \theta_{\tilde{j}_1 \dots \tilde{j}_{d_x}} \text{ for any } (j_1, \dots, j_{d_x}) \geq (\tilde{j}_1, \dots, \tilde{j}_{d_x}); \\ & \quad -1 \leq \theta_{j_1 \dots j_{d_x}} \leq 1 \text{ for all } (j_1, \dots, j_{d_x}), \end{aligned} \tag{26}$$

where X_{ij} denotes the j -th element of X_i .⁹ The first inequality constraints restrict the

⁸See, e.g., Wang and Ghosh (2012) for the bound and shape preserving properties of the multivariate Bernstein polynomials.

⁹The linear program in (25) for the nonparametric monotone classification problem has n -decision

feasible classifiers to a class of non-decreasing functions. The second inequality constraints bound the feasible classifiers to $[-1, 1]$.

We then consider applying the general result for the excess risk bound in Theorem 5.2 with $\tilde{F} = \mathcal{B}_k$. Lemma C.2 in Appendix C gives finite upper bounds on two approximation errors:

$$\inf_{f \in \mathcal{B}_k} R_{\phi_h}(f) - \inf_{f \in \mathcal{F}_M} R_{\phi_h}(f) \text{ and } R_{\phi_h}(1\{\cdot \in G_{\hat{f}_B^\dagger}\} - 1\{\cdot \notin G_{\hat{f}_B^\dagger}\}) - R_{\phi_h}(\hat{f}_B^\dagger)$$

in (24) upon setting $(\tilde{\mathcal{F}}, \tilde{\mathcal{F}}, \hat{f}) = (\mathcal{B}_k, \mathcal{F}_M, \hat{f}_B^\dagger)$. The binarized coefficients in \hat{f}_B^\dagger help us to make the second approximation error shrink to zero. Moreover, Lemma C.1 in Appendix C gives a finite upper bound on the bracketing entropy of \mathcal{B}_k . Combining these results, the following theorem gives a finite sample upper bound on the mean of the \mathcal{G}_M -constrained excess classification risk of \hat{f}_B^\dagger .

Theorem 6.3. *Let \mathcal{P} be a class of distributions on $\{-1, 1\} \times \mathcal{X}$ that satisfy the same conditions as in Theorem 6.2. Let $\tilde{\tau}_n = \log(n) / \sqrt{n}$ if $d_x = 1$ and $\tilde{\tau}_n = n^{-1/d_x}$ if $d_x \geq 2$. Define $\tilde{q}_n = \sqrt{n}\tilde{\tau}_n$. Then the following holds:*

$$\begin{aligned} \sup_{P \in \mathcal{P}} E_{P^n} \left[R(\hat{f}_B^\dagger) - \inf_{f \in \mathcal{F}_{\mathcal{G}_M}} R(f) \right] &\leq 2D_1 \tilde{\tau}_n + 4D_2 \exp(-D_1^2 \tilde{q}_n^2) \\ &+ 4A \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{8}{\sqrt{k_j}}, \end{aligned} \quad (27)$$

where D_1 and D_2 are some positive constants, which depend only on d_x and A .

Proof. From the fact that $\mathcal{B}_k \subseteq \mathcal{F}_M$ and Lemma C.1 in Appendix C, we have $H_1^B(\epsilon, \mathcal{B}_k, P_X) \leq C\epsilon^{-d_x}$ for some positive constant C , which depends only on A , and all $\epsilon > 0$. Then the result follows by combining Theorem 5.2 and Lemma C.2. \square

The upper bound in (27) converges to zero as the sample size n and the order of the Bernstein polynomials k_j ($j = 1, \dots, d_x$) increase. Note that the rate of convergence for the estimation error in this theorem, $\tilde{\tau}_n$, is slower than that in Theorem 6.2, τ_n . The difference in the rates of convergence is due to the different orders of the upper bounds on $H_1^B(\epsilon, \mathcal{G}_M, P_X)$ and $H_1^B(\epsilon, \mathcal{F}_M, P_X)$ in Lemmas 6.1 and C.1. To achieve the convergence rate of $\tilde{\tau}_n$ for the mean of the excess risk of \hat{f}_B^\dagger , Theorem 6.3 suggests the tuning parameters k_j , for $j = 1, \dots, d_x$, should be set sufficiently large so that $\sqrt{\log k_j / k_j} = O(\tilde{\tau}_n)$.

variables, whereas the linear program in (26) has $(k_1 + 1) \times \dots \times (k_{d_x} + 1)$ -decision variables. Hence when the dimension of X is small to moderate relative to the sample size n , the linear programming for the Bernstein polynomials would be easier to compute. The reverse is also true.

In practice, one may want to select the complexity of the Bernstein polynomials by minimizing penalized empirical surrogate risk. The classification and treatment choice literature (Koltchinskii (2006), Mbakop and Tabord-Meehan (2021), and references therein) analyze the regret properties and oracle inequalities for penalized risk minimizing classifiers. We leave theoretical investigation of the applicability of penalization methods to the current hinge risk minimization using Bernstein polynomials for future research.

7 Extension to individualized treatment rules

This section extends the primary results obtained in Sections 3 and 4 for binary classification to the weighted classification introduced in Section 1.1, and to causal policy learning. Extensions of the results in Sections 5 and 6 to weighted classification are presented in Appendix D. We use the same notation and definitions as those introduced in Section 1.1. We term R^ω and R_ϕ^ω , defined in (5) and (6), weighted classification risk and weighted ϕ -risk, respectively. Throughout this section, with some abuse of notation, we denote by P a distribution on $\mathbb{R}_+ \times \{-1, 1\} \times \mathcal{X}$ and suppose that $(\omega, Y, D) \sim P$.

7.1 Consistency of weighted classification with hinge loss

We first show consistency of weighted classification with hinge risk by adapting the analyses in Sections 3 and 4. Given a prespecified \mathcal{G} , let \mathcal{F}_G be as in Section 2. Analogous to $\mathcal{R}(G)$ and $\mathcal{R}_\phi(G)$, we define $\mathcal{R}^\omega(G) \equiv \inf_{f \in \mathcal{F}_G} R^\omega(f)$, the weighted-classification risk evaluated at G , and $\mathcal{R}_\phi^\omega(G) \equiv \inf_{f \in \mathcal{F}_G} R_\phi^\omega(f)$, the weighted ϕ -risk evaluated at G . Note that $\mathcal{R}^\omega(G) = R^\omega(f)$ for all $f \in \mathcal{F}_G$. Let $\mathcal{R}^{\omega*} \equiv \inf_{G \in \mathcal{G}} \mathcal{R}^\omega(G) = \inf_{f \in \mathcal{F}_G} R^\omega(f)$ be the optimal weighted risk, and $\mathcal{G}^* \equiv \arg \inf_{G \in \mathcal{G}} \mathcal{R}^\omega(G)$ be the collection of best prediction sets.

For the non-negative weight variable ω , we define

$$\begin{aligned}\omega_+(x) &\equiv E_P[\omega \mid X = x, Y = +1] \\ \omega_-(x) &\equiv E_P[\omega \mid X = x, Y = -1].\end{aligned}$$

Let $C_\phi(a, b, c, d) \equiv a\phi(c)d + b\phi(-c)(1-d)$, and

$$\begin{aligned}C_\phi^{w+}(\omega_+, \omega_-, \eta) &\equiv \inf_{0 \leq f \leq 1} C_\phi(\omega_+, \omega_-, f, \eta), \\ C_\phi^{w-}(\omega_+, \omega_-, \eta) &\equiv \inf_{-1 \leq f < 0} C_\phi(\omega_+, \omega_-, f, \eta), \\ \Delta C_\phi^\omega(\omega_+, \omega_-, \eta) &\equiv C_\phi^{w+}(\omega_+, \omega_-, \eta) - C_\phi^{w-}(\omega_+, \omega_-, \eta),\end{aligned}$$

which are analogous to C_ϕ^+ , C_ϕ^- and ΔC_ϕ defined in Section 3.

The next theorem generalizes Theorems 3.2, 3.6, and Corollary 3.7 to weighted classification, giving a necessary and sufficient condition for equivalence of the risk ordering among surrogate loss functions. In particular, we show that hinge loss functions share a common risk ordering with the 0-1 loss function.

Theorem 7.1. *Let ϕ_1 and ϕ_2 be classification-calibrated loss functions in the sense of Definition 2.3. Then*

$$\mathcal{R}_{\phi_1}^\omega(G_1) \leq \mathcal{R}_{\phi_1}^\omega(G_2) \Leftrightarrow \mathcal{R}_{\phi_2}^\omega(G_1) \leq \mathcal{R}_{\phi_2}^\omega(G_2)$$

holds for any distribution P on $\mathbb{R}_+ \times \{-1, 1\} \times \mathcal{X}$, any class of measurable subsets $\mathcal{G} \subseteq 2^{\mathcal{X}}$, and any $G_1, G_2 \in \mathcal{G}$ if and only if there exists $c > 0$ such that $\Delta C_{\phi_2}^\omega(\omega_+, \omega_-, \eta) = c \Delta C_{\phi_1}^\omega(\omega_+, \omega_-, \eta)$ holds for any $(\omega_+, \omega_-, \eta) \in \mathbb{R}_+ \times \mathbb{R}_+ \times [0, 1]$. In particular, the 0-1 loss function, $\phi_{01}(\alpha) = 1\{\alpha \leq 0\}$, satisfies

$$\Delta C_{\phi_{01}}^\omega(\omega_+, \omega_-, \eta) = -\omega_+ \eta + \omega_- (1 - \eta), \quad (28)$$

and the hinge loss function $\phi_h(\alpha) = c \max\{0, 1 - \alpha\}$ satisfies

$$\Delta C_{\phi_h}^\omega(\omega_+, \omega_-, \eta) = c(-\omega_+ \eta + \omega_- (1 - \eta)) = c \Delta C_{\phi_{01}}^\omega(\omega_+, \omega_-, \eta).$$

Proof. See Appendix E. □

Theorem 7.1 and inequalities similar to (14) lead to a generalized Zhang's (2004) inequality for weighted classification, as shown in the next corollary.

Corollary 7.2. *For any distribution P on $\mathbb{R}_+ \times \{-1, 1\} \times \mathcal{X}$ and any surrogate loss function ϕ satisfying $\Delta C_\phi^\omega(\omega_+, \omega_-, \eta) = c(-\omega_+ \eta + \omega_- (1 - \eta))$,*

$$c(R^\omega(f) - \inf_{f \in \mathcal{F}_G} R^\omega(f)) \leq R_\phi^\omega(f) - \inf_{f \in \mathcal{F}_G} R_\phi^\omega(f) \quad (29)$$

holds for any $f \in \mathcal{F}_G$.

Proof. See Appendix E. □

Remark 7.3. *Table 2 shows the forms of $\Delta C_\phi^\omega(\omega_+, \omega_-, \eta)$ for the hinge loss, exponential loss, logistic loss, quadratic loss, and truncated quadratic loss functions, where $\mu_+ \equiv \omega_+ \eta$ and $\mu_- \equiv \omega_- (1 - \eta)$. With the exception of the hinge loss function, none of these functions satisfy $\Delta C_\phi^\omega(\omega_+, \omega_-, \eta) = c(-\mu_+ + \mu_-)$ for some positive constant $c > 0$. That is, similar*

Table 2: Surrogate loss functions and their associated forms for ΔC_ϕ^ω

Loss function	$\phi(\alpha)$	$\Delta C_\phi^\omega(\omega_+, \omega_-, \eta)$
0-1 loss	$1\{\alpha \leq 0\}$	$-\mu_+ + \mu_-$
Hinge loss	$c \max\{0, 1 - \alpha\}$	$c(-\mu_+ + \mu_-)$
Exponential loss	$e^{-\alpha}$	$\begin{cases} (\sqrt{\mu_+} - \sqrt{\mu_-})^2 & \text{if } \mu_+ \leq \mu_- \\ -(\sqrt{\mu_+} - \sqrt{\mu_-})^2 & \text{if } \mu_+ > \mu_- \end{cases}$
Logistic loss	$\log(1 + e^{-\alpha})$	$\begin{cases} -\mu_+ \log\left(\frac{2\mu_+}{\mu_+ + \mu_-}\right) - \mu_- \log\left(\frac{2\mu_-}{\mu_+ + \mu_-}\right) & \text{if } \mu_+ \leq \mu_- \\ \mu_+ \log\left(\frac{2\mu_+}{\mu_+ + \mu_-}\right) + \mu_- \log\left(\frac{2\mu_-}{\mu_+ + \mu_-}\right) & \text{if } \mu_+ > \mu_- \end{cases}$
Quadratic loss	$(1 - \alpha)^2$	$\begin{cases} \frac{(\mu_+ - \mu_-)^2}{\mu_+ + \mu_-} & \text{if } \mu_+ \leq \mu_- \\ -\frac{(\mu_+ - \mu_-)^2}{\mu_+ + \mu_-} & \text{if } \mu_+ > \mu_- \end{cases}$
Truncated quadratic loss	$(\max\{0, 1 - \alpha\})^2$	$\begin{cases} \frac{(\mu_+ - \mu_-)^2}{\mu_+ + \mu_-} & \text{if } \mu_+ \leq \mu_- \\ -\frac{(\mu_+ - \mu_-)^2}{\mu_+ + \mu_-} & \text{if } \mu_+ > \mu_- \end{cases}$

Note: $\mu_+ = \omega_+ \eta$ and $\mu_- = \omega_-(1 - \eta)$.

to the standard binary classification, hinge losses also have a special status in weighted classification, since they are the only surrogate losses that preserve classification risk.

Similar to the analysis in Section 4, we consider adding functional form restrictions to the class of classifiers \mathcal{F}_G . Let $\tilde{\mathcal{F}}_G$ be a subclass of \mathcal{F}_G . We suppose that the non-negative weight variable ω is bounded from above.

Condition 7.4 (Bounded weight variable). *There exists $M < \infty$ such that $0 \leq \omega \leq M$ a.s.*

In causal policy learning, Condition 7.4 holds if the outcome variable Z has bounded support and the propensity score $e(x)$ satisfies a strict overlap condition. For example, if the support of Z is contained in $[-\tilde{M}, \tilde{M}]$ for some $\tilde{M} < \infty$, and the propensity score satisfies $\kappa < e(x) < 1 - \kappa$ for some $\kappa \in (0, 1/2)$ and all $x \in \mathcal{X}$, then the weight variable for the causal policy learning ω_p defined in (7) is bounded from above by \tilde{M}/κ a.s.

The following theorem, which is analogous to Theorem 4.4, shows that the two conditions (A1) and (A2) in Theorem 4.4 remain sufficient for $\tilde{\mathcal{F}}_G$ to guarantee the consistency of the hinge risk minimization approach to weighted classification.

Theorem 7.5. *Given a distribution P on $\mathbb{R}_+ \times \{-1, 1\} \times \mathcal{X}$ and a class of measurable subsets $\mathcal{G} \subseteq 2^{\mathcal{X}}$, suppose that $\tilde{\mathcal{F}}_G \subset \mathcal{F}_G$ satisfy the conditions (A1) and (A2) in Theorem 4.4 and that the weight variable ω satisfies Condition 7.4. Then the following claims hold:*

- (i) $\tilde{f}^* \in \arg \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}^\omega(f)$ minimizes the weighted-classification risk $R^\omega(\cdot)$ over \mathcal{F}_G .
- (ii) For $G^* \in \mathcal{G}^*$, $\tilde{f}_{G^*}^*$ is a minimizer of $R_{\phi_h}^\omega(\cdot)$ over $\tilde{\mathcal{F}}_G$.

Proof. See Appendix E. □

8 Empirical illustration

To illustrate the hinge risk minimizing approach in a causal policy learning setting, we apply our weighted classification methods with a monotone constraint to experimental data from Karlan et al. (2019). Karlan et al. (2019) conducted three experiments in India and the Philippines in which they paid off the high-interest moneylender debt of market vendors and gave them brief financial training. Though the main focus of Karlan et al. (2019) is to understand why a debt-trap occurs (i.e., why some individuals repeatedly take on high-interest rate loans), we regard their treatment of paying off debt as policy intervention and study effective treatment allocation maximizing the value of household business in the population.

We use the data of Karlan et al. (2019), collected from an RCT experiment conducted in Cagayan de Oro, the Philippines, in 2010. The observations are divided into two groups: a treatment group (debt paid off and received financial training) and a control group. The data was collected over 5 periods, comprising a baseline period before the policy intervention and 4 follow-up periods after the policy intervention was implemented. The follow-up periods correspond to the 1st, 4th, 8th, and 18-19th months after the policy intervention was implemented. We label these follow-up periods as periods 1 through 4, respectively. After dropping observations with missing values, our main sample consists of 411 observations, of which 289 (122) observations belong to the treatment (control) group.

We focus on the effect of treatment on the present value (PV) of business at the time of the policy intervention. To define the PV of business, we introduce some notation. Let $D \in \{-1, 1\}$ denote a treatment indicator, with 1 indicating treatment and -1 indicating control. Let C be the amount of moneylender debt paid off, which we assume to be a cost of treatment. Let P_t denote monthly take-home profit in month t after the policy intervention. For $s \in \{1, 2, 3, 4\}$, let $P_{(s)}$ denote an average of the monthly take-home profit observed in the follow-up period s . Since the data was collected only for 4 follow-up periods, we assume that $P_t = P_{(1)}$ for $t \leq 2$, $P_t = P_{(2)}$ for $3 \leq t \leq 6$, $P_t = P_{(3)}$ for $7 \leq t \leq 11$, and $P_t = P_{(4)}$ for $t \geq 12$. Let WC_{19} denote the total working capital of business in the 19th month after the policy intervention (at the end of the follow-up survey), which corresponds to the total working capital of business observed in the final follow-up period.¹⁰ All variables P_t , WC_{19} , and C are measured in USD, using the average exchange rate during September and October, 2010. For $T \geq 1$, we define the outcome

¹⁰The total working capital of a business, as defined in Karlan et al. (2019), is the worth of current business assets plus the amount spent on an average restocking trip minus any current or daily loans owed.

variable Z as the PV of business minus the cost of the treatment as follows:

$$Z = \sum_{t=1}^T \frac{1}{(1+r)^t} P_t + \frac{1}{(1+r)^{19}} WC_{19} - D \cdot C,$$

where r is the monthly discount rate of the business value, to $0.037/12$, the average of the annual real interest rate in force in the Philippines between 2010 and 2019 divided by 12. In our analysis, we set $T = 19$ (the duration of the follow-up survey), 60, 120, and 240.

The covariates that we use for treatment assignment are the amount of moneylender debt in the baseline period, financial literacy index, and food expenditure ratio. Specifically, we consider the following three sets of covariates¹¹:

$$\begin{aligned} X_1 &= \{\text{amount of moneylender debt in the baseline period, financial literacy index}\}, \\ X_2 &= \{\text{amount of moneylender debt in the baseline period, food expenditure ratio}\}, \\ X_3 &= \{\text{amount of moneylender debt in the baseline period, financial literacy index,} \\ &\quad \text{food expenditure ratio}\}. \end{aligned}$$

All of the covariates are observed in the baseline period. The food expenditure ratio is the proportion of expenditure on food and drink to total expenditure, and is used to gauge living standards, i.e., the higher the ratio, the poorer the household is considered to be. The amount of moneylender debt in the baseline period is measured in USD, using the same average exchange rate used to define Z . For each set of covariates, we constrain the class of feasible treatment rules to the class of monotonically increasing treatment rules, i.e.,

$$\mathcal{F}_M = \{f : f(x) \geq f(\tilde{x}) \text{ for any } x, \tilde{x} \in \mathcal{X} \text{ with } x \geq \tilde{x} \text{ and } f(\cdot) \in [-1, 1]\},$$

where $x \geq \tilde{x}$ is an element-wise weak inequality. Any treatment rule in \mathcal{F}_M is more likely to award treatment to individuals with more baseline debt, higher financial literacy and a higher food expenditure ratio. This class of monotone rules is intended to represent the planner's (hypothetical) objective of prioritizing those households that are more financially-strained and debt-trapped, and those that have a lower standard of living. At the same time, the monotonicity of assignment in financial literacy disciplines the allocation of the policy by prioritizing those who are more likely to escape a debt-trap, assuming that financial literacy is a good predictor for this.

Let \hat{e} be the empirical probability of treatment, which we use as an estimated propen-

¹¹Aside from these three covariates, Karlan et al. (2019) use vendors' time inconsistent preferences, possession of savings at a bank, math skills index, and predicted probability of household income shock for estimation of the conditional causal effects.

sity score. Setting $\omega = |Z|/(D\hat{e} + (1 - D)/2)$ and $Y = \text{sign}(Z) \cdot D$ in the definitions of the weighted classification and surrogate risks ((5) and (6)), let \hat{f}_M be a classifier that minimizes the empirical weighted hinge risk over \mathcal{F}_M (see Appendix D.2 for details of this procedure). Denote f_{-1} and f_{+1} the never-treating and always-treating rules (i.e., $f_{-1}(x) = -1$ and $f_{+1}(x) = 1$ for all x).

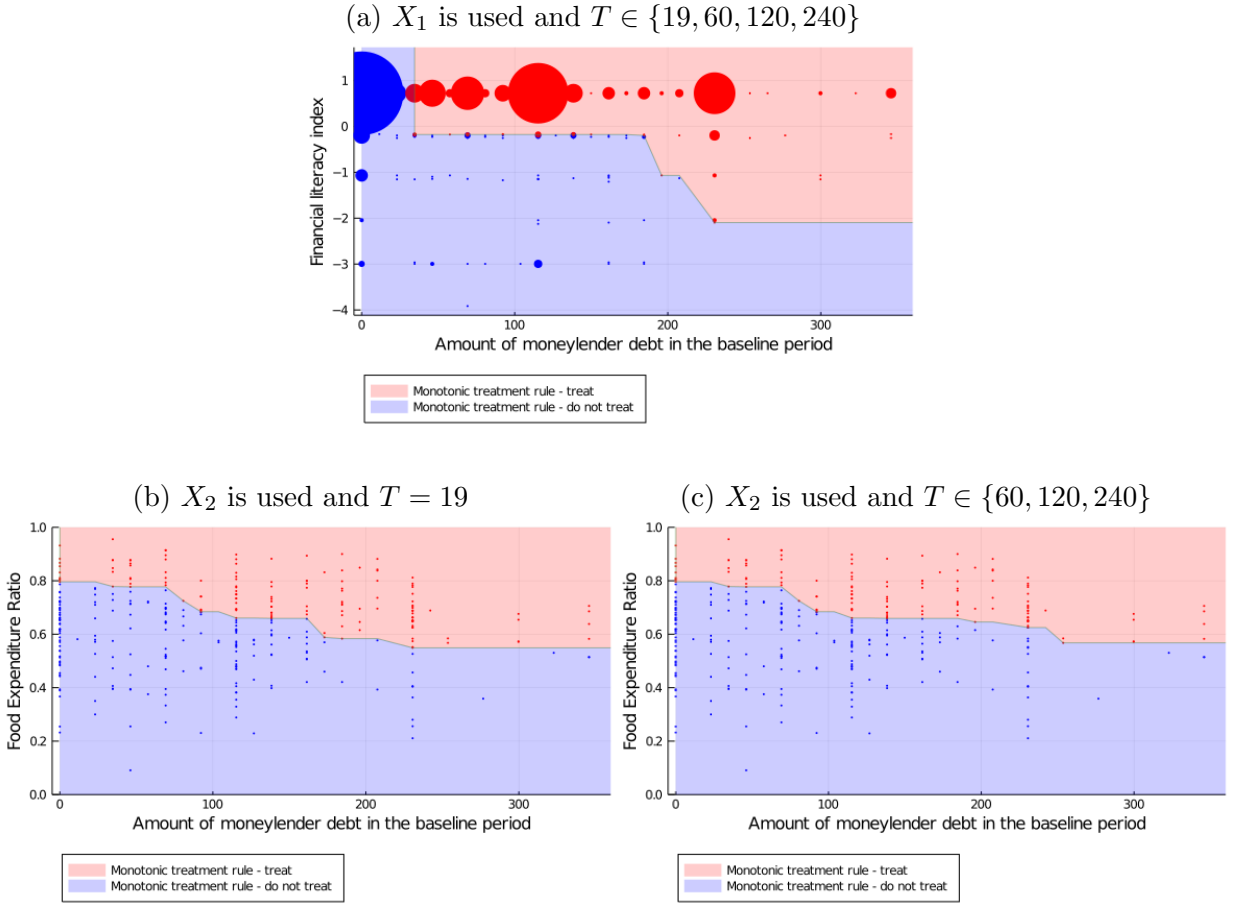
Table 3 shows the estimated welfare gains of \hat{f}_M relative to the never-treating rule (i.e., $W(\hat{f}_M) - W(f_{-1})$). Figure 3 illustrates the resulting treatment allocation $G_{\hat{f}_M}$ for the covariate sets X_1 and X_2 , for each T . The welfare gains in Table 3 are estimated using the same sample as is used to estimate \hat{f}_M ; hence, the estimated welfare gains of \hat{f}_M in Table 3 are positively biased.¹² Using X_2 rather than X_1 leads to higher estimates of the welfare gain from using \hat{f}_M . Figure 3 shows the estimated treatment allocation rules under the monotonicity constraint. When we use X_1 , we obtain an optimal treatment rule that is identical for all $T \in \{19, 60, 120, 240\}$. In contrast, the estimated treatment rules differ between $T = 19$ and $T \in \{60, 120, 240\}$ when we use X_2 .

Table 3: Estimates of welfare gains, probability of treatment, and $E[Z(1)]$ and $E[Z(0)]$

X	T	Sample size	Welfare gain of		Probability of treatment	$E[Z(1)]$	$E[Z(0)]$
			\hat{f}_M	f_{+1}			
X_1	19	411	627.0	495.6	0.43	4406.8	3911.2
	60		1724.0	1223.4	0.43	12516.6	11293.2
	120		3093.8	2132.2	0.43	22642.3	20510.1
	240		5164.4	3506.1	0.43	37949.7	34443.6
X_2	19	337	839.5	518.6	0.43	4483.7	3965.1
	60		2214.2	1031.1	0.40	12513.3	11482.2
	120		3991.9	1670.9	0.40	22538.8	20867.9
	240		6679.5	2638.3	0.40	37694.7	35056.5
X_3	19	337	1226.9	518.6	0.42	4483.7	3965.1
	60		3182.7	1031.1	0.44	12513.3	11482.2
	120		5631.2	1670.9	0.45	22538.8	20867.9
	240		9333.5	2638.3	0.45	37694.7	35056.5

¹²To our knowledge, there is no bias-free estimation method available for the welfare gain of the policy that maximizes the full-sample objective function over the large class of monotone assignment rules. Out-of-sample methods based upon sample-splitting is a simple approach for inferring welfare gain, but the policy estimated from a subsample can sacrifice welfare performance. We leave development of bias-free estimation applicable to the current context for future research.

Figure 3: Monotone treatment rules obtained from empirical hinge risk minimization



Notes: In each figure, every circle represents the sample density at a given observation.

9 Conclusion

This paper studies the consistency of surrogate risk minimization approaches to classification and weighted classification given a constrained set of classifiers, where weighted classification subsumes policy learning for individualized treatment assignment rules. Our focus is on how surrogate risk minimizing classifiers behave if the constrained class of classifiers violates the assumption of correct specification. Our first main result shows that when the constraint restricts classifiers' prediction sets only, hinge losses are the only loss functions that secure consistency of the surrogate-risk minimizing classifier without the assumption of correct specification. When the constraint additionally restricts the functional form of the classifiers, the surrogate risk minimizing classifier is not generally consistent even with hinge loss. Our second main result is to show that, in this case, the set of conditions (A1) and (A2) in Theorem 4.4 becomes a sufficient condition for the consistency of the hinge risk minimizing classifier.

This paper also investigates the statistical properties of hinge risk minimizing classifiers

in terms of uniform upper bounds on the excess regret. Exploiting hinge loss and the class of monotone classifiers in the monotone classification problem, we show that the empirical surrogate-risk minimizing classifier can be computed using linear programming. All of the results obtained in the standard classification setting are naturally extended to the weighted classification problem, so that our contributions carry over to causal policy learning and related applications.

Appendix

A Proofs of the results in Sections 2–4

This appendix provides proof of the results in Sections 2–4 alongside some auxiliary lemmas. We first give the proofs of Propositions 2.1 and 2.2.

Proof of Proposition 2.1. The statement (i) follows from Claim 3 of Theorem 1 of Bartlett et al. (2006). To prove the statement (ii), let $z_1, z_2 \in \mathbb{R}_+$ be such that $\phi(z_2) < \phi(z_1)$. Such a pair (z_1, z_2) exists in a neighborhood of zero because, from Theorem 2 of Bartlett et al. (2006), ϕ is differentiable at 0 and $\phi'(0) < 0$ if ϕ is classification-calibrated and convex. Then the pair (z_1, z_2) satisfies

$$\arg \min_{(b_1, b_2) \in \{(-1, 0), (0, 1)\}} \phi(-b_1 z_1) + \phi(b_2 z_2) = (0, 1).$$

Suppose that \mathcal{F} is a constrained class such that $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$ with $\mathcal{F}_1 = \{f(x) = x^T b : (b_1, b_2) = (-1, 0), b \in \mathbb{R}^{d_x}\}$ and $\mathcal{F}_2 = \{f(x) = x^T b : (b_1, b_2) = (0, 1), b \in \mathbb{R}^{d_x}\}$, where b_j denotes the j -th element of b . Let $x_1 = (z_1, 0, \dots, 0)$ and $x_2 = (0, z_2, 0, \dots, 0)$ be two points in \mathcal{X} , and let P be a distribution such that $\eta(x_1) = 0$, $\eta(x_2) = 1$, and $P_X(x_1) = P_X(x_2) = 1/2$. Given the pair (P, \mathcal{F}) , any classifier f_1 in \mathcal{F}_1 has the same sign as the Bayes classifier, P_X -almost everywhere, because $f_1(x_1) < 0$ and $f_1(x_2) = 0$ while $\eta(x_1) < 1/2$ and $\eta(x_2) \geq 1/2$. This means that \mathcal{F} , as well as \mathcal{F}_1 , is R -correctly specified for such P . On the other hand, any classifier f_2 in \mathcal{F}_2 does not have the same sign as the Bayes classifier at x_1 because $f_2(x_1) = 0$ while $\eta(x_1) < 1/2$. f_ϕ^* must be in \mathcal{F}_2 because for any $f_1 \in \mathcal{F}_1$ and $f_2 \in \mathcal{F}_2$, $R_\phi(f_2) = \phi(z_2)/2 < \phi(z_1)/2 = R_\phi(f_1)$. Hence $R(f_\phi^*) > R(f_{\text{Bayes}}^*)$ holds. \mathcal{F} is also R_ϕ -misspecified because any classifier that minimizes R_ϕ over all measurable functions takes a negative value at x_1 whereas $f_\phi^*(x_1) > 0$. \square

Proof of Proposition 2.2 Assume R -correct specification of $\mathcal{F}_{\mathcal{G}}$. Then $\mathcal{F}_{\mathcal{G}}$ includes a classifier f^* that is identical to or shares the same sign as $f_{Bayes}^*(x) = 2\eta(x) - 1$, P_X -almost everywhere. Since $f \in \mathcal{F}_{\mathcal{G}}$ is unconstrained except for $G_f \in \mathcal{G}$ and $-1 \leq f(\cdot) \leq 1$, the classification-calibrated property of ϕ and the representation of the surrogate risk $R_{\phi}(f) = E_{P_X}[C_{\phi}(f(X), \eta(X))]$ implies

$$f_{\phi}^*(x) \in \arg \min_{a: (2\eta(x)-1)a \geq 0, |a| \leq 1} C_{\phi}(a, \eta(x)),$$

P_X -almost everywhere, because otherwise f^* dominates f_{ϕ}^* in terms of the surrogate risk. This means that $f_{\phi}^*(x)$ has the same sign as $f_{Bayes}^*(x)$, P_X -almost everywhere, i.e., $R(f_{\phi}^*) = R(f_{Bayes}^*)$ holds.

Assume conversely that $\mathcal{F}_{\mathcal{G}}$ is R -misspecified. Then $\text{sign}(f_{\phi}^*)$ has to differ from $\text{sign}(f_{Bayes}^*(x))$ for some x with positive measure in terms of P_X . Failure to find such a value x contradicts the assumption of R -misspecification of $\mathcal{F}_{\mathcal{G}}$. The difference in signs then implies $R(f_{\phi}^*) > R(f_{Bayes}^*)$. \square

We here let ϕ be any surrogate loss function. Before proceeding to the proofs of the results in Sections 3 and 4, we note that if ϕ is classification-calibrated, ΔC_{ϕ} has the same sign as the Bayes classifier: for any $\eta \in [0, 1] \setminus \{1/2\}$,

$$\Delta C_{\phi}(\eta) \begin{cases} > 0 & \text{if } \eta > 1/2 \\ < 0 & \text{if } \eta < 1/2 \end{cases}, \quad (30)$$

which will be used in the following proofs.

Proof of Theorem 3.6.

(‘If’ part)

For any class of measurable subsets $\mathcal{G} \subseteq 2^{\mathcal{X}}$ and any $G_1, G_2 \in \mathcal{G}$, we show in Theorem 3.2 that $\mathcal{R}_{\phi_1}(G_2) \geq \mathcal{R}_{\phi_1}(G_1)$ is equivalent to

$$\int_{G_2 \setminus G_1} \Delta C_{\phi_1}(\eta(x)) dP_X(x) \geq \int_{G_1 \setminus G_2} \Delta C_{\phi_1}(\eta(x)) dP_X(x).$$

This inequality does not change if we replace $\Delta C_{\phi_1}(\eta(x))$ with $\Delta C_{\phi_2}(\eta(x)) = c\Delta C_{\phi_1}(\eta(x))$ with $c > 0$. From Theorem 3.2, $\mathcal{R}_{\phi_2}(G_2) \geq \mathcal{R}_{\phi_2}(G_1)$ if $\Delta C_{\phi_1}(\eta(x))$ is replaced by $\Delta C_{\phi_2}(\eta(x))$ in the above inequality. Therefore, if $\Delta C_{\phi_2}(\cdot) = c\Delta C_{\phi_1}(\cdot)$ with $c > 0$, $\phi_1 \stackrel{u}{\sim} \phi_2$ holds.

(‘Only if’ part)

We prove the ‘only if’ part of the theorem by exploiting a specific class of data generating

processes (DGPs). Suppose $\mathcal{X} = \{1, 2\}$ and $\mathcal{G} = \{\emptyset, G_1, G_2, \mathcal{X}\}$ with $G_1 = \{1\}$ and $G_2 = \{2\}$. Let $\alpha = P(X = 1) (= 1 - P(X = 2))$ and $(\eta_1, \eta_2) = (\eta(1), \eta(2))$. The DGP varies depending on the values of $(\alpha, \eta_1, \eta_2) \in [0, 1]^3$.

In what follows, we will show that

$$\frac{\Delta C_{\phi_1}(\eta_1)}{\Delta C_{\phi_1}(\eta_2)} = \frac{\Delta C_{\phi_2}(\eta_1)}{\Delta C_{\phi_2}(\eta_2)} \quad (31)$$

holds for any $(\eta_1, \eta_2) \in ([0, 1] \setminus \{1/2\})^2$. Then applying Lemma A.1 below proves the ‘only if’ part of the theorem.

In the current setting, for $G \in \mathcal{G}$, $\mathcal{R}_\phi(G)$ can be written as

$$\begin{aligned} \mathcal{R}_\phi(G) &= P(X = 1) \Delta C_\phi(\eta_1) 1\{1 \in G\} + P(X = 2) \Delta C_\phi(\eta_2) 1\{2 \in G\} \\ &\quad + \sum_{x=1}^2 P(X = x) C_\phi^-(\eta(x)) \\ &= \alpha \Delta C_\phi(\eta_1) 1\{1 \in G\} + (1 - \alpha) \Delta C_\phi(\eta_2) 1\{2 \in G\} + C_{\alpha, \eta_1, \eta_2}, \end{aligned}$$

where $C_{\alpha, \eta_1, \eta_2} \equiv \alpha \Delta C_\phi(\eta_1) + (1 - \alpha) \Delta C_\phi(\eta_2)$, which does not depend on G . Thus, we have

$$\begin{aligned} \mathcal{R}_\phi(\emptyset) &= C_{\alpha, \eta_1, \eta_2} \\ \mathcal{R}_\phi(G_1) &= \alpha \Delta C_\phi(\eta_1) + C_{\alpha, \eta_1, \eta_2}, \\ \mathcal{R}_\phi(G_2) &= (1 - \alpha) \Delta C_\phi(\eta_2) + C_{\alpha, \eta_1, \eta_2}, \\ \mathcal{R}_\phi(\mathcal{X}) &= \alpha \Delta C_\phi(\eta_1) + (1 - \alpha) \Delta C_\phi(\eta_2) + C_{\alpha, \eta_1, \eta_2}. \end{aligned}$$

In what follows, we will show that (31) holds, separately, in four cases: (i) $\eta_1 > 1/2$ and $\eta_2 > 1/2$; (ii) $\eta_1 < 1/2$ and $\eta_2 < 1/2$; (iii) $\eta_1 < 1/2$ and $\eta_2 > 1/2$; (iv) $\eta_1 > 1/2$ and $\eta_2 < 1/2$.

First, we consider case (i): $\eta_1 > 1/2$ and $\eta_2 > 1/2$. Because we assume $\phi_1 \stackrel{u}{\sim} \phi_2$,

$$\mathcal{R}_{\phi_1}(G_1) \leq \mathcal{R}_{\phi_1}(G_2) \Leftrightarrow \mathcal{R}_{\phi_2}(G_1) \leq \mathcal{R}_{\phi_2}(G_2),$$

holds for any $(\alpha, \eta_1, \eta_2) \in (0, 1) \times (1/2, 1]^2$. This is equivalent to

$$\alpha \Delta C_{\phi_1}(\eta_1) \leq (1 - \alpha) \Delta C_{\phi_1}(\eta_2) \Leftrightarrow \alpha \Delta C_{\phi_2}(\eta_1) \leq (1 - \alpha) \Delta C_{\phi_2}(\eta_2)$$

for any $(\alpha, \eta_1, \eta_2) \in (0, 1) \times (1/2, 1]^2$. Let $\gamma^+ \equiv (1 - \alpha)/\alpha$, which can take any value in $(0, +\infty)$ by varying α on $(0, 1)$. From the classification-calibrated property (30), both $\Delta C_{\phi_1}(\eta)$ and $\Delta C_{\phi_2}(\eta)$ are positive for $\eta \in (1/2, 1]$. Thus, it follows for any $(\gamma^+, \eta_1, \eta_2) \in$

$(0, +\infty) \times (1/2, 1]^2$ that

$$\frac{\Delta C_{\phi_1}(\eta_1)}{\Delta C_{\phi_1}(\eta_2)} \leq \gamma^+ \Leftrightarrow \frac{\Delta C_{\phi_2}(\eta_1)}{\Delta C_{\phi_2}(\eta_2)} \leq \gamma^+, \quad (32)$$

where both $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2)$ and $\Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ are positive. Since (32) holds for any value of $\gamma^+ \in (0, +\infty)$, $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2) = \Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ holds for any $(\eta_1, \eta_2) \in (1/2, 1]^2$.

Similarly, for case (ii): $\eta_1 < 1/2$ and $\eta < 1/2$, the following equivalence statement holds for any $(\gamma^+, \eta_1, \eta_2) \in (0, +\infty) \times [0, 1/2)^2$:

$$\frac{\Delta C_{\phi_1}(\eta_1)}{\Delta C_{\phi_1}(\eta_2)} \geq \gamma^+ \Leftrightarrow \frac{\Delta C_{\phi_2}(\eta_1)}{\Delta C_{\phi_2}(\eta_2)} \geq \gamma^+, \quad (33)$$

where both $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2)$ and $\Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ are positive. Thus, varying the value of γ^+ on $(0, +\infty)$ in (33) shows that $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2) = \Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ holds for any $(\eta_1, \eta_2) \in [0, 1/2)^2$.

Next, we consider case (iii): $\eta_1 < 1/2$ and $\eta_2 > 1/2$. Because we assume $\phi_1 \stackrel{u}{\sim} \phi_2$, it follows for any $(\alpha, \eta_1, \eta_2) \in (0, 1) \times [0, 1/2) \times (1/2, 1]$ that

$$\mathcal{R}_{\phi_1}(\emptyset) \leq \mathcal{R}_{\phi_1}(\mathcal{X}) \Leftrightarrow \mathcal{R}_{\phi_2}(\emptyset) \leq \mathcal{R}_{\phi_2}(\mathcal{X}),$$

which is further equivalent to

$$0 \leq \alpha \Delta C_{\phi_1}(\eta_1) + (1 - \alpha) \Delta C_{\phi_1}(\eta_2) \Leftrightarrow 0 \leq \alpha \Delta C_{\phi_2}(\eta_1) + (1 - \alpha) \Delta C_{\phi_2}(\eta_2).$$

Let $\gamma^- \equiv (\alpha - 1)/\alpha$, which can take any value in $(-\infty, 0)$ by varying the value of α on $(0, 1)$. Because $\Delta C_{\phi_1}(\eta_1) < 0$ and $\Delta C_{\phi_2}(\eta_2) > 0$ hold for $(\eta_1, \eta_2) \in [0, 1/2) \times (1/2, 1]$ due to the classification-calibrated property (30),

$$\frac{\Delta C_{\phi_1}(\eta_1)}{\Delta C_{\phi_1}(\eta_2)} \geq \gamma^- \Leftrightarrow \frac{\Delta C_{\phi_2}(\eta_1)}{\Delta C_{\phi_2}(\eta_2)} \geq \gamma^- \quad (34)$$

for any $(\gamma^-, \eta_1, \eta_2) \in (-\infty, 0) \times [0, 1/2) \times (1/2, 1]$, where both $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2)$ and $\Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ are negative. Thus, varying the value of γ^- on $(-\infty, 0)$ in (34) shows that $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2) = \Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ holds for any $(\eta_1, \eta_2) \in [0, 1/2) \times (1/2, 1]$.

Similarly, in case (iv): $\eta_1 > 1/2$ and $\eta_2 < 1/2$, the following equivalence statement holds for any $(\gamma^-, \eta_1, \eta_2) \in (-\infty, 0) \times (1/2, 1] \times [0, 1/2)$:

$$\frac{\Delta C_{\phi_1}(\eta_1)}{\Delta C_{\phi_1}(\eta_2)} \leq \gamma^- \Leftrightarrow \frac{\Delta C_{\phi_2}(\eta_1)}{\Delta C_{\phi_2}(\eta_2)} \leq \gamma^-, \quad (35)$$

where both $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2)$ and $\Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ are negative. Therefore, varying the value of γ^- in (35) shows that $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2) = \Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ for any $(\eta_1, \eta_2) \in (1/2, 1] \times [0, 1/2)$.

Combining these four results, we have $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2) = \Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ for any $(\eta_1, \eta_2) \in ([0, 1] \setminus \{1/2\})^2$. Then the proof follows from Lemma A.1 below. \square

Lemma A.1. *Let ϕ_1 and ϕ_2 be classification-calibrated loss functions. If $\Delta C_{\phi_1}(\eta_1)/\Delta C_{\phi_1}(\eta_2) = \Delta C_{\phi_2}(\eta_1)/\Delta C_{\phi_2}(\eta_2)$ holds for any $(\eta_1, \eta_2) \in ([0, 1] \setminus \{1/2\})^2$, then there exists some constant $c > 0$ such that $\Delta C_{\phi_2}(\eta) = c\Delta C_{\phi_1}(\eta)$ for any $\eta \in [0, 1]$.*

Proof. For $\eta \in [0, 1] \setminus \{1/2\}$, let $c(\eta)$ be a value such that

$$\Delta C_{\phi_2}(\eta) = c(\eta) \Delta C_{\phi_1}(\eta). \quad (36)$$

Because ϕ_1 and ϕ_2 are classification-calibrated, $c(\eta)$ must be positive from (30). We will show that $c(\eta)$ is constant over $\eta \in [0, 1] \setminus \{1/2\}$ by contradiction.

To obtain a contradiction, suppose there exists $(\eta_1, \eta_2) \in ([0, 1] \setminus \{1/2\})^2$ such that $c(\eta_1) \neq c(\eta_2)$. From the assumption, the following equations hold

$$\begin{aligned} \Delta C_{\phi_2}(\eta_1) &= \left(\frac{\Delta C_{\phi_2}(\eta_2)}{\Delta C_{\phi_1}(\eta_2)} \right) \Delta C_{\phi_1}(\eta_1), \\ \Delta C_{\phi_2}(\eta_2) &= \left(\frac{\Delta C_{\phi_2}(\eta_1)}{\Delta C_{\phi_1}(\eta_1)} \right) \Delta C_{\phi_2}(\eta_2). \end{aligned}$$

Combining these equations with equation (36), we have $\Delta C_{\phi_2}(\eta_2) = c(\eta_1) \Delta C_{\phi_1}(\eta_2)$ and $\Delta C_{\phi_2}(\eta_2) = c(\eta_2) \Delta C_{\phi_1}(\eta_2)$. However, this contradicts the assumption that $c(\eta_1) \neq c(\eta_2)$. Therefore, $c(\eta)$ must be constant over $\eta \in [0, 1] \setminus \{1/2\}$.

When $\eta = 1/2$, $\Delta C_{\phi_1}(\eta) = \Delta C_{\phi_2}(\eta) = 0$ holds by definition. In this case, $\Delta C_{\phi_2}(\eta) = c\Delta C_{\phi_1}(\eta)$ holds for any c . \square

The following expression of the hinge risk will be used in the proofs of Lemma 4.3 and Theorem 4.4:

$$\begin{aligned} R_{\phi_h}(f) &= \int_{\mathcal{X}} [\eta(x)(1 - f(x)) + (1 - \eta(x))(1 + f(x))] dP_X(x) \\ &= \int_{\mathcal{X}} (1 - 2\eta(x)) f(x) dP_X(x) + 1. \end{aligned} \quad (37)$$

Proof of Lemma 4.3. Fix a distribution P on $\{-1, 1\} \times \mathcal{X}$, and let $\tilde{f} \in \tilde{\mathcal{F}}_{\mathcal{G}, J}$. \tilde{f} has the form

$$\tilde{f}(x) = 2 \sum_{j=1}^J c_j 1\{x \in G_j\} - 1 \quad (38)$$

for some $G_1, \dots, G_J \in \mathcal{G}$ such that $G_J \subseteq \dots \subseteq G_1$ and some $c_j \geq 0$ for $j = 1, \dots, J$ such that $\sum_{j=1}^J c_j = 1$. Thus, substituting \tilde{f} into (37) yields

$$R_{\phi_h}(\tilde{f}) = 2 \sum_{j=1}^J \left[c_j \int_{G_j} (1 - 2\eta(x)) dP_X(x) \right] + 2P(Y = 1). \quad (39)$$

Comparing (39) with equation (9) leads to

$$R_{\phi_h}(\tilde{f}) = 2 \sum_{j=1}^J c_j \mathcal{R}(G_j). \quad (40)$$

From this expression and the assumption that $\sum_{j=1}^J c_j = 1$, $R_{\phi_h}(\tilde{f}) \geq 2\mathcal{R}^*$ holds for any $\tilde{f} \in \tilde{\mathcal{F}}_{\mathcal{G}, J}$.

For $G^* \in \mathcal{G}^*$, a function $\tilde{f}_{G^*}(x) = 2 \cdot 1\{x \in G^*\} - 1$ can be extracted from $\tilde{\mathcal{F}}_{\mathcal{G}, J}$ by setting $G_1 = G^*$ and $c_1 = 1$. Then, from equation (40), $R_{\phi_h}(\tilde{f}_{G^*}) = 2\mathcal{R}^*$ holds; that is, \tilde{f}_{G^*} minimizes $R_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}, J}$. This proves statement (ii) of the lemma.

We will next show that $\tilde{\mathcal{F}}_{\mathcal{G}, J}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$ (statement (i) of the lemma). To obtain a contradiction, suppose that a classifier \tilde{f} with the form of (38) minimizes $R_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}, J}$ but does not minimize $R(\cdot)$ over $\mathcal{F}_{\mathcal{G}}$. As \tilde{f} does not minimize $R(\cdot)$ over $\mathcal{F}_{\mathcal{G}}$, $G_{\tilde{f}} \notin \mathcal{G}^*$ holds. Let m be the smallest integer in $\{1, \dots, J\}$ such that $\sum_{j=1}^m c_j \geq 1/2$ in (58). Then $G_m = G_{\tilde{f}}$. It then follows that

$$\begin{aligned} R_{\phi_h}(\tilde{f}) &= 2 \sum_{j=1}^J c_j \mathcal{R}(G_j) = 2c_m \mathcal{R}(G_m) + \sum_{j \in \{1, \dots, m-1, m+1, \dots, J\}} c_j \mathcal{R}(G_j) \\ &\geq 2c_m \mathcal{R}(\tilde{f}) + 2(1 - c_m) \mathcal{R}^* \\ &> 2\mathcal{R}^*, \end{aligned}$$

where the last line follows from $c_m > 0$ and $G_{\tilde{f}} \notin \mathcal{G}^*$. Hence $R_{\phi_h}(\tilde{f})$ does not take the minimum value of $R_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}, J}$ (i.e., $R_{\phi_h}(\tilde{f}) > 2\mathcal{R}^*$), which contradicts the assumption that \tilde{f} minimizes $R_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}, J}$. Therefore, \tilde{f} minimizes $R(\cdot)$ over $\mathcal{F}_{\mathcal{G}}$. Since this discussion is valid for any distribution P on $\{-1, 1\} \times \mathcal{X}$, we conclude $\tilde{\mathcal{F}}_{\mathcal{G}, J}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$. \square

Proof of Theorem 4.4. Fix a distribution P on $\{-1, 1\} \times \mathcal{X}$, and let $\tilde{f}^* \in \arg \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}(f)$. Define a class of step functions with at most $J(\geq 1)$ jumps

$$\tilde{\mathcal{F}}_J^* \equiv \left\{ f(x) = 2 \sum_{j=1}^J c_j 1\{x \in G_j\} - 1 : G_j \in \mathcal{G} \text{ and } c_j \geq 0 \text{ for } j = 1, \dots, J; \right. \\ \left. G_J \subseteq \dots \subseteq G_1; G_f = G_{\tilde{f}^*}; \sum_{j=1}^J c_j = 1 \right\},$$

which is a class of step functions whose prediction sets correspond to $G_{\tilde{f}^*}$, i.e., $G_f = G_{\tilde{f}^*}$ for any $f \in \tilde{\mathcal{F}}_J^*$. We look to find a sequence of functions $\{\tilde{f}_J^*\}_{J=1}^\infty$ such that $\tilde{f}_J^* \in \tilde{\mathcal{F}}_J^*$ for any J and $\tilde{f}_J^*(x) \rightarrow \tilde{f}^*(x)$, as $J \rightarrow \infty$, P_X -almost everywhere. Setting $G_j = \{x : \tilde{f}^*(x) \geq 2(j/J) - 1\}$ and $c_j = 1/J$ for $j = 1, \dots, J$ in the definition of $\tilde{\mathcal{F}}_J^*$, we define

$$\tilde{f}_J^*(\cdot) \equiv \frac{2}{J} \sum_{j=1}^J 1\{\tilde{f}^*(\cdot) \geq 2(j/J) - 1\} - 1.$$

For all $x \in \mathcal{X}$,

$$\left| \tilde{f}_J^*(x) - \tilde{f}^*(x) \right| = \left| \frac{2}{J} \sum_{j=1}^J 1\{\tilde{f}^*(x) \geq 2(j/J) - 1\} - 1 - \tilde{f}^*(x) \right| \\ \leq \frac{1}{J} \rightarrow 0 \text{ as } J \rightarrow \infty.$$

Thus, $\tilde{f}_J^*(x) \rightarrow \tilde{f}^*(x)$ holds P_X -almost everywhere.

Then it follows that

$$R_{\phi_h}(\tilde{f}^*) = \int_{\mathcal{X}} (1 - 2\eta(x)) \tilde{f}^*(x) dP_X(x) + 1 \\ = \lim_{J \rightarrow \infty} \int_{\mathcal{X}} (1 - 2\eta(x)) \tilde{f}_J^*(x) dP_X(x) + 1 \\ = \lim_{J \rightarrow \infty} R_{\phi_h}(\tilde{f}_J^*) \geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_J^*} R_{\phi_h}(\tilde{f}) \quad (41)$$

$$\geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_{G,J}} R_{\phi_h}(\tilde{f}), \quad (42)$$

where the first and third equalities follow from equation (37); the second equality follows from the dominated convergence theorem, which holds because both $(1 - 2\eta) \tilde{f}_J^*(x) \rightarrow (1 - 2\eta(x)) \tilde{f}^*(x)$ and $\left| (1 - 2\eta(x)) \tilde{f}_J^*(x) \right| < 1$ hold P_X -almost everywhere; the first inequality follows from $\tilde{f}_J^* \in \tilde{\mathcal{F}}_J^*$; the last inequality follows from $\tilde{\mathcal{F}}_J^* \subseteq \tilde{\mathcal{F}}_{G,J}$.

Lemma 4.3 shows that $\inf_{\tilde{f} \in \tilde{\mathcal{F}}_{G,J}} R_{\phi_h}(\tilde{f}) = 2\mathcal{R}^*$ for any J . Hence, from equation (42),

we have

$$R_{\phi_h}(\tilde{f}^*) \geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_{\mathcal{G}, J}} R_{\phi_h}(\tilde{f}) \geq 2\mathcal{R}^*,$$

On the other hand, Lemma 4.3 also shows that $R_{\phi_h}(\tilde{f}_{G^*}) = 2\mathcal{R}^*$. Therefore, \tilde{f}_{G^*} minimizes $R_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}}$, which proves statement (ii) of the theorem.

Next we will show that $\tilde{\mathcal{F}}_{\mathcal{G}}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$ (statement (i) of the theorem). To obtain a contradiction, suppose that \tilde{f}^* does not minimize $R(\cdot)$ over $\mathcal{F}_{\mathcal{G}}$ (i.e., $G_{\tilde{f}^*} \notin \mathcal{G}^*$). Then, from the proof of Lemma 4.3, for any J and $\tilde{f} \in \tilde{\mathcal{F}}_J^*$, $R_{\phi_h}(\tilde{f}) > 2\mathcal{R}^*$. Therefore, from equation (41),

$$R_{\phi_h}(\tilde{f}^*) \geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_J^*} R_{\phi_h}(\tilde{f}) > 2\mathcal{R}^*.$$

This contradicts the assumption that \tilde{f}^* minimizes $R_{\phi}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}}$ because, as we have seen, $\tilde{f}_{G^*} \in \tilde{\mathcal{F}}_{\mathcal{G}}$ achieves $R_{\phi_h}(\tilde{f}_{G^*}) = 2\mathcal{R}^*$. Therefore, $G_{\tilde{f}^*} \in \mathcal{G}^*$ must hold, i.e., $\tilde{f}^* \in \arg \inf_{f \in \mathcal{F}_{\mathcal{G}}} R(f)$. Since this discussion holds for any distribution P on $\{-1, 1\} \times \mathcal{X}$, $\tilde{\mathcal{F}}_{\mathcal{G}}$ is a classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$. \square

Proof of Corollary 4.5. Note that, for $G \in \mathcal{G}$,

$$\begin{aligned} \mathcal{R}(G) &= \int_{\mathcal{X}} [\eta(x)1\{x \notin G\} + (1 - \eta(x))1\{x \in G\}] dP_X(x) \\ &= \int_{\mathcal{X}} [\eta(x)1\{x \in G^c\} + (1 - \eta(x))(1 - 1\{x \in G^c\})] dP_X(x) \\ &= - \int_{G^c} (1 - 2\eta(x)) dP_X(x) + P(Y = -1). \end{aligned} \quad (43)$$

By equations (9) and (43), $cR(f)$ can be written as

$$cR(f) = \frac{1}{2} \left\{ \int_{\mathcal{X}} c(1 - 2\eta(x)) (1\{x \in G_f\} - 1\{x \notin G_f\}) dP_X(x) + c \right\}.$$

By equation (37), the term inside the braces equals $R_{\phi}(\tilde{f}_{G_f})$. Combining this result with Theorem 4.4 (i) leads to equation (16).

When $\tilde{\mathcal{F}}_{\mathcal{G}} = \mathcal{F}_{\mathcal{G}}$, by equation (16) and Corollary 3.4,

$$R_{\phi}(\tilde{f}_{G_f}) - R_{\phi}(f) \leq R_{\phi}(f) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R_{\phi}(f)$$

holds. Combining this with equation (16) leads to the second result. \square

B Proof of Theorems 5.1 and 5.2

This appendix provides proof of the results in Section 5 with some auxiliary results. The results below are related to the theory of empirical processes. We refer to Alexander (1984), Mammen and Tsybakov (1999), Tsybakov (2004), and Mbakop and Tabord-Meehan (2021) for the general strategy of the proof.

Given $G \in \mathcal{G}$, let $\mathbf{1}_G$ be an indicator function on \mathcal{X} such that $\mathbf{1}_G(x) = 1\{x \in G\}$. We first give the definition of bracketing entropy for a class of functions and a class of sets.

Definition B.1 (Bracketing entropy). (i) Let \mathcal{F} be a class of functions on \mathcal{X} . For $f \in \mathcal{F}$, let $\|f\|_{p,Q} \equiv (\int_{\mathcal{X}} |f(x)|^p dQ(x))^{1/p}$. $\|\cdot\|_{p,Q}$ is the $L_p(Q)$ -metric on \mathcal{X} , where Q is a measure on \mathcal{X} . Given a pair of functions (f_1, f_2) with $f_1 \leq f_2$, let $[f_1, f_2] \equiv \{f \in \mathcal{F} : f_1 \leq f \leq f_2\}$ be the bracket. Given $\epsilon > 0$, let $N_p^B(\epsilon, \mathcal{F}, Q)$ be the smallest k such that there exist pairs of functions (f_j^L, f_j^U) , $j = 1, \dots, k$, with $f_j^L \leq f_j^U$ that satisfy $\|f_j^U - f_j^L\|_{p,Q} < \epsilon$ and

$$\mathcal{F} \subseteq \cup_{j=1}^k [f_j^L, f_j^U].$$

We term $N_p^B(\epsilon, \mathcal{F}, Q)$ the $L_p(Q)$ -bracketing number of \mathcal{F} , and $H_p^B(\epsilon, \mathcal{F}, Q) \equiv \log N_p^B(\epsilon, \mathcal{F}, Q)$ the $L_p(Q)$ -bracketing entropy of \mathcal{F} . We also refer to $[f_j^L, f_j^U]$ as the ϵ -bracket with respect to $L_p(Q)$ if and only if $\|f_j^U - f_j^L\|_{p,Q} < \epsilon$ holds.

(ii) Given a class of measurable subsets $\mathcal{G} \subseteq 2^{\mathcal{X}}$, let $\mathcal{H}_{\mathcal{G}} \equiv \{\mathbf{1}_G : G \in \mathcal{G}\}$. We define $H_p^B(\epsilon, \mathcal{G}, Q) \equiv H_p^B(\epsilon, \mathcal{H}_{\mathcal{G}}, Q)$ and term this the $L_p(Q)$ -bracketing entropy of \mathcal{G} .

Note that in the definition of $N_p^B(\epsilon, \mathcal{F}, Q)$, the functions f_j^L and f_j^U do not have to belong to \mathcal{F} . Note also that if $\mathcal{F} \subseteq \tilde{\mathcal{F}}$, $H_p^B(\epsilon, \mathcal{F}, Q) \leq H_p^B(\epsilon, \tilde{\mathcal{F}}, Q)$ holds. When $\mathbf{1}_G \in \mathcal{F}$ for all $G \in \mathcal{G}$, $H_p^B(\epsilon, \mathcal{G}, Q) \leq H_p^B(\epsilon, \mathcal{F}, Q)$ holds.

The following theorem gives a finite-sample upper bound on the mean of the estimation error in Section 5, auxiliary results of which are provided below.

Theorem B.2. Let $\tilde{\mathcal{F}}$ be a class of classifiers whose members satisfy $-1 \leq f \leq 1$. Suppose that \mathcal{P} is a class of distributions on $\{-1, 1\} \times \mathcal{X}$ such that there exist positive constants C and r for which

$$H_1^B(\epsilon, \tilde{\mathcal{F}}, P_X) \leq C\epsilon^{-r}$$

holds for any $P \in \mathcal{P}$ and $\epsilon > 0$. Let q_n and τ_n be as in Theorem 5.1. Let \hat{f} minimize $\hat{R}_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}$. Then the following holds:

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[R_{\phi_h}(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}(f) \right] \leq \begin{cases} 4D_1\tau_n + 8D_2 \exp(-D_1^2 q_n^2) & \text{for } r \geq 1 \\ 4D_3\tau_n + 4n^{-1}D_4 & \text{for } 0 < r < 1 \end{cases}$$

for some positive constants D_1, D_2, D_3, D_4 , which depend only on C and r .

Proof. Fix $P \in \mathcal{P}$. Let \check{f}^* minimize $R_{\phi_h}(\cdot)$ over $\check{\mathcal{F}}$. Define a class of functions $\check{\mathcal{F}} \equiv \{(f+1)/2 : f \in \check{\mathcal{F}}\}$, which normalizes $\check{\mathcal{F}}$ so that $0 \leq f \leq 1$ for all $f \in \check{\mathcal{F}}$.

A standard argument gives

$$\begin{aligned}
E_{P^n} \left[R_{\phi_h}(\hat{f}) - \inf_{f \in \check{\mathcal{F}}} R_{\phi_h}(f) \right] &\leq E_{P^n} \left[R_{\phi_h}(\hat{f}) - \hat{R}_{\phi_h}(\hat{f}) + \hat{R}_{\phi_h}(\check{f}^*) - R_{\phi_h}(\check{f}^*) \right] \\
&\quad \left(\because \hat{R}_{\phi_h}(\hat{f}) \leq \hat{R}_{\phi_h}(\check{f}^*) \right) \\
&= 2E_{P^n} \left[R_{\phi_h} \left(\frac{\hat{f}+1}{2} \right) - \hat{R}_{\phi_h} \left(\frac{\hat{f}+1}{2} \right) \right] \\
&\quad + 2E_{P^n} \left[\hat{R}_{\phi_h} \left(\frac{\check{f}^*+1}{2} \right) - R_{\phi_h} \left(\frac{\check{f}^*+1}{2} \right) \right] \\
&\leq 4 \sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| R_{\phi_h}(f) - \hat{R}_{\phi_h}(f) \right| \right]. \tag{44}
\end{aligned}$$

Since $R_{\phi_h}(f) - \hat{R}_{\phi_h}(f)$ can be seen as the centered empirical process indexed by $f \in \check{\mathcal{F}}$, we can apply results in empirical process theory to (44) to obtain a finite-sample upper bound on the mean of the excess hinge risk.

We follow the general strategy of Theorem 1 in Mammen and Tsybakov (1999) and Proposition B.1 in Mbakop and Tabord-Meehan (2021). Note that

$$\sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| R_{\phi_h}(f) - \hat{R}_{\phi_h}(f) \right| \right] = \sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| E_P(Yf(x)) - \frac{1}{n} \sum_{i=1}^n Y_i f(X_i) \right| \right] \tag{45}$$

and that

$$\sup_{f \in \check{\mathcal{F}}} \left| E_P(Yf(x)) - \frac{1}{n} \sum_{i=1}^n Y_i f(X_i) \right| \leq 2$$

with probability one.

We first prove the result for the case of $r \geq 1$. For any $f \in \check{\mathcal{F}}$ and $D > 0$,

$$\begin{aligned}
&\frac{\sqrt{n}}{q_n} \sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| E(Yf(x)) - \frac{1}{n} \sum_{i=1}^n Y_i f(X_i) \right| \right] \\
&\leq D + \frac{2\sqrt{n}}{q_n} P^n \left(\sup_{f \in \check{\mathcal{F}}} \frac{\sqrt{n}}{q_n} \left| E(Yf(x)) - \frac{1}{n} \sum_{i=1}^n Y_i f(X_i) \right| > D \right).
\end{aligned}$$

We apply Corollary B.3 by setting $Z = (Y, X)$, $g(z_1) = z_1$ and $\mathcal{H} = \check{\mathcal{F}}$ as appear in the statement of the corollary. Note that, since $\check{\mathcal{F}}$ is an affine transformation multiplying $1/2$ to $f \in \check{\mathcal{F}}$, $H_1^B(\epsilon, \check{\mathcal{F}}, P_2) = H_1^B(2\epsilon, \mathcal{H}, P_2)$ holds. Then, by Corollary B.3 shown below with

$K = 2^{-r}C$, there exist $D_1, D_2, D_3 > 0$, depending only on C and r , such that

$$P^n \left(\sup_{f \in \tilde{\mathcal{F}}} \frac{\sqrt{n}}{q_n} \left| E(Yf(x)) - \frac{1}{n} \sum_{i=1}^n Y_i f(X_i) \right| > D \right) \leq D_2 \exp(-D^2 q_n^2),$$

for $D_1 \leq D \leq D_3 \sqrt{n}/q_n$. Thus when $r \geq 1$, we have

$$\begin{aligned} \tau_n^{-1} E_{P^n} \left[R_{\phi_h}(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}(f) \right] &\leq 4\tau_n^{-1} \sup_{f \in \tilde{\mathcal{F}}} E_{P^n} \left[\left| E(Yf(x)) - \frac{1}{n} \sum_{i=1}^n Y_i f(X_i) \right| \right] \\ &\leq 4D_1 + 8\tau_n^{-1} D_2 \exp(-D_1^2 q_n^2), \end{aligned}$$

which leads to the result for the case of $r \geq 1$.

The result for the case of $0 < r < 1$ follows immediately by applying Lemma B.4 to equation (45), where we set $Z = (Y, X)$, $g(z_1) = z_1$, and $\mathcal{H} = \tilde{\mathcal{F}}$. \square

We now give the proofs of Theorems 5.1 and 5.2.

Proof of Theorem 5.2. The result in Theorem 5.2 follows by combining equation (22) and Theorem B.2. \square

Proof of Theorem 5.1. Let $P \in \mathcal{P}$ be fixed. Define $\hat{f}^\dagger \equiv \tilde{f}_{G_{\hat{f}}}$. Since $R(\hat{f}^\dagger) = R(\hat{f})$, it follows from equation (18) that

$$R(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_G} R(f) = R(\hat{f}^\dagger) - \inf_{f \in \tilde{\mathcal{F}}_G} R(f) = \frac{1}{2} \left(R_{\phi_h}(\hat{f}^\dagger) - \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}(f) \right). \quad (46)$$

When (19) holds for all $\epsilon > 0$, the result follows by applying Theorem B.2 to (46).

We consider the case when (20) holds for all $\epsilon > 0$. Define a class of step functions $\mathcal{I}_G \equiv \{\tilde{f}_G : G \in \mathcal{G}\}$. We now show that (A) \hat{f}^\dagger minimizes $\hat{R}_{\phi_h}(\cdot)$ over \mathcal{I}_G , and that (B) $\inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}(f) = \inf_{f \in \mathcal{I}_G} R_{\phi_h}(f)$. If both hold, we can apply Theorem B.2 to the excess hinge risk in (46) with $\tilde{\mathcal{F}}$ replaced by \mathcal{I}_G .

We first prove (B). Since

$$\inf_{f \in \mathcal{I}_G} R_{\phi_h}(f) = 2 \inf_{G \in \mathcal{G}} \mathcal{R}(G) = 2\mathcal{R}^*,$$

Theorem 4.4 (i) shows that $\inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}(f) = \inf_{f \in \mathcal{I}_G} R_{\phi_h}(f)$.

We next prove (A). Note first that $\hat{f}^\dagger \in \mathcal{I}_G$ holds. Let P_n be the empirical distribution of the sample $\{(Y_i, X_i) : i = 1, \dots, n\}$. Replacing P with P_n in Theorem 4.4 shows that \hat{f} minimizes $\hat{R}(\cdot)$ over $\tilde{\mathcal{F}}_G$. Hence $\hat{\mathcal{R}}(G) \equiv \inf_{f \in \tilde{\mathcal{F}}_G} \hat{R}(f)$ is minimized by $G_{\hat{f}}$ over \mathcal{G} . Then Theorem 4.4 (ii) with P replaced by P_n shows that the new classifier \hat{f}^\dagger also minimizes

$\hat{R}_{\phi_h}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}}$. Then replacing R_{ϕ_h} with \hat{R}_{ϕ_h} in the statement of (B), \hat{f}^\dagger minimizes $\hat{R}_{\phi_h}(\cdot)$ over $\mathcal{I}_{\mathcal{G}}$.

From the definitions of $H_1^B(\epsilon, \mathcal{G}, P_X)$ and $\mathcal{I}_{\mathcal{G}}$, we have $H_1^B(\epsilon, \mathcal{I}_{\mathcal{G}}, P_X) = H_1^B(2\epsilon, \mathcal{G}, P_X)$. Therefore, we can apply Theorem B.2, with $\tilde{\mathcal{F}}$ replaced by $\mathcal{I}_{\mathcal{G}}$, to (46) and then obtain the inequality in (21). \square

The following corollary is similar to Corollary D.1 in Mbakop and Tabord-Meehan (2021). The difference is that the class of functions \mathcal{H} in the following corollary does not need to be a class of binary functions.

Corollary B.3. *Let $Z = (Z_1, Z_2) \sim P$, and $\{Z_i\}_{i=1}^n$ be a sequence of random variables that are i.i.d. as Z . Denote by P_2 the marginal distribution of Z_2 . Suppose P_2 is absolutely continuous with respect to Lebesgue measure and its density is bounded from above by a finite constant $A > 0$. Let \mathcal{F} be a class of real-valued functions of the form $f(z) = f(z_1, z_2) = g(z_1) \cdot h(z_2)$, where $h \in \mathcal{H}$, \mathcal{H} is a class of functions with values in $[0, 1]$, and g takes values in $[-1, 1]$. Suppose \mathcal{H} satisfies*

$$H_1^B(\epsilon, \mathcal{H}, P_2) \leq K\epsilon^{-r}$$

for some constants $K > 0$ and $r \geq 1$ and for all $\epsilon > 0$. Then there exist positive constants D_1, D_2, D_3 , depending only on K and r , such that for $n \geq 3$:

$$P^n \left(\sup_{f \in \mathcal{F}} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n (f(Z_i) - E_P[f(Z_i)]) \right| > xq_n \right) \leq D_2 \exp(-x^2 q_n^2),$$

for $D_1 \leq x \leq D_3 \sqrt{n}/q_n$, where

$$q_n = \begin{cases} \log n & r = 1 \\ n^{(r-1)/2(r+2)} & r > 1 \end{cases}.$$

Proof. Let $[h_j^L, h_j^U]$, $j = 1, \dots, N_1^B(\epsilon, \mathcal{H}, P_2)$, be a set of ϵ -brackets of \mathcal{H} with respect to $L_1(P_2)$ such that $\|h_j^U - h_j^L\|_{1, P_2} < \epsilon$ and $\mathcal{H} \subseteq \cup_{j=1}^{N_1^B(\epsilon, \mathcal{H}, P_2)} [h_j^L, h_j^U]$. Since $|h_j^U - h_j^L| < 1$, $\|h_j^U - h_j^L\|_{2, P_2}^2 \leq \|h_j^U - h_j^L\|_{1, P_2} < \epsilon$ holds. We hence have

$$H_2^B(\epsilon, \mathcal{H}, P_2) \leq H_1^B(\epsilon^2, \mathcal{H}, P_2) \leq K\epsilon^{-2r}.$$

The result immediately follows by applying Proposition B.1. \square

Lemma B.4. *Maintain the same definitions and assumptions as in Corollary B.3 with $r \geq 1$ replaced by $0 < r < 1$. Then, there exist positive constants D_3 and D_4 , depending only on K and r , such that:*

$$\sup_{f \in \mathcal{F}} E_{P^n} \left[\left| \frac{1}{n} \sum_{i=1}^n f(Z_i) - E_P[f(Z)] \right| \right] \leq \frac{D_3}{\sqrt{n}} + \frac{D_4}{n}.$$

Proof. We look to apply Proposition 3.5.15 in Giné and Nickl (2016). Note first that $|f| \leq 1$ and $\|f\|_{2,P} \leq 1$ for all $f \in \mathcal{F}$. Then we can apply Proposition 3.5.15 in Giné and Nickl (2016), with $F = 1$ and $\delta = 1$, and obtain

$$\begin{aligned} \sup_{f \in \mathcal{F}} E_{P^n} \left[\left| \frac{1}{n} \sum_{i=1}^n f(Z_i) - E_P[f(Z)] \right| \right] &\leq \left(\frac{58}{\sqrt{n}} + \frac{1}{3n} \int_0^2 \sqrt{\log(2N_2^B(\epsilon, \mathcal{F}, P))} d\epsilon \right) \\ &\quad \times \int_0^2 \sqrt{\log(2N_2^B(\epsilon, \mathcal{F}, P))} d\epsilon \\ &\leq \left(\frac{58}{\sqrt{n}} + \frac{2}{3n} + \frac{1}{3n} \int_0^2 \sqrt{H_2^B(\epsilon, \mathcal{F}, P)} d\epsilon \right) \\ &\quad \times \left(\frac{2}{3} + \frac{1}{3} \int_0^2 \sqrt{H_2^B(\epsilon, \mathcal{F}, P)} d\epsilon \right). \end{aligned} \quad (47)$$

By combining the arguments from the proofs of Corollary B.3 and Proposition B.1 below, we have

$$H_2^B(\epsilon, \mathcal{F}, P) \leq K\epsilon^{-2r}.$$

Substituting this upper bound into (47) yields

$$\begin{aligned} \sup_{f \in \mathcal{F}} E_{P^n} \left[\left| \frac{1}{n} \sum_{i=1}^n f(Z_i) - E_P[f(Z)] \right| \right] &\leq \left(\frac{58}{\sqrt{n}} + \frac{2}{3n} + \frac{1}{3n} \int_0^2 K\epsilon^{-r} d\epsilon \right) \\ &\quad \times \left(\frac{2}{3} + \frac{1}{3} \int_0^2 K\epsilon^{-r} d\epsilon \right) \\ &= \left(\frac{58}{\sqrt{n}} + \frac{2}{3n} + \frac{2^{1-r}K}{3n(1-r)} \right) \left(\frac{2}{3} + \frac{2^{1-r}K}{3(1-r)} \right). \end{aligned}$$

Therefore, setting

$$\begin{aligned} D_3 &\equiv \left(\frac{116}{3} + \frac{29 \cdot 2^{2-r}K}{3(1-r)} \right), \\ D_4 &\equiv \left(\frac{2}{3} + \frac{2^{1-r}K}{3(1-r)} \right)^2, \end{aligned}$$

leads to the result. \square

Proposition B.1. *Let $Z = (Z_1, Z_2) \sim P$, and $\{Z_i\}_{i=1}^n$ be a sequence of random variables that are i.i.d. as Z . Denote by P_2 the marginal distribution of Z_2 . Let \mathcal{F} be a class of real-valued functions of the form $f(z) = f(z_1, z_2) = g(z_1) \cdot h(z_2)$, where $h \in \mathcal{H}$, \mathcal{H} is a class of functions with values in $[0, 1]$, and g takes values in $[-1, 1]$. Suppose \mathcal{H} satisfies*

$$H_2^B(\epsilon, \mathcal{H}, P_2) \leq K\epsilon^{-r} \quad (48)$$

for some constants $K > 0$ and $r \geq 2$ and for all $\epsilon > 0$. Then there exist positive constants C_1, C_2, C_3 , depending only on K and r , such that if

$$\xi \leq \frac{\sqrt{n}}{128} \quad (49)$$

and

$$\xi \geq \begin{cases} C_1 n^{(r-2)/2(r+2)} & r \geq 2 \\ C_2 \log \max(n, e) & r = 2 \end{cases}, \quad (50)$$

then

$$P^n \left(\sup_{f \in \mathcal{F}} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n (f(Z_i) - E_P[f(Z_i)]) \right| > \xi \right) \leq C_3 \exp(-\xi^2).$$

Proof. We follow the general strategy of Theorem 2.3 and Corollary 2.4 in Alexander (1984) and Proposition D.1 in Mbakop and Tabord-Meehan (2021). Define

$$v_n(f) \equiv \frac{1}{\sqrt{n}} \sum_{i=1}^n [f(Z_i) - E(f(Z_i))].$$

We start by giving some definitions. Let $\delta_0 > \delta_1 > \dots > \delta_N > 0$ be a sequence of real numbers where $\{\delta_j\}_{j=0}^N$ and N are specified later. For each δ_j , there exists a set of δ_j -brackets \mathcal{H}_j^B of \mathcal{H} with respect to $L_2(P_2)$ such that $|\mathcal{H}_j^B| = N_2^B(\delta_j, \mathcal{H}, P_2)$. Define the function $H(\cdot) : (0, \infty) \rightarrow [0, \infty)$ as follows:

$$H(u) = \begin{cases} Ku^{-r} & \text{if } u < 1 \\ 0 & \text{if } u \geq 1 \end{cases}.$$

Note that by Assumption (48) and the fact that \mathcal{H} has unit diameter by definition, $N_2^B(\delta_j, \mathcal{H}, P_2) \leq \exp(H(\delta_j))$ for all $\delta_j > 0$. For each $0 \leq j \leq N$ and any $f = g \cdot h \in \mathcal{F}$, define $f_j^L \equiv g \cdot h_j^L 1\{g \geq 0\} + g \cdot h_j^U 1\{g < 0\}$ and $f_j^U \equiv g \cdot h_j^U 1\{g \geq 0\} + g \cdot h_j^L 1\{g < 0\}$ for some (h_j^L, h_j^U) that forms a δ_j -bracket for h with respect to $L_2(P_2)$ such that $h \in [h_j^L, h_j^U]$

and $[h_j^L, h_j^U] \in \mathcal{H}_j^B$. From the construction, $[f_j^U, f_j^L]$ is a δ_j -bracket for f with respect to $L_2(P)$. Let $f_j = f_j^L$, and let $\mathcal{F}_j = \{f_j : f \in \mathcal{F}\}$. We have $|\mathcal{F}_j| \leq \exp(H(\delta_j))$ and $\|f - f_j\|_{2,P} < \delta_j$ for every $f \in \mathcal{F}$.

By a standard chaining argument,

$$\begin{aligned} P\left(\sup_{f \in \mathcal{F}} |v_n(f)| > \xi\right) &\leq |\mathcal{F}_0| \sup_{f \in \mathcal{F}} P\left(|v_n(f)| > \frac{7}{8}\xi\right) \\ &\quad + \sum_{j=0}^{N-1} |\mathcal{F}_j| |\mathcal{F}_{j+1}| \sup_{f \in \mathcal{F}} P(|v_n(f_j - f_{j+1})| > \eta_j) \\ &\quad + P\left(\sup_{f \in \mathcal{F}} |v_n(f_N - f)| > \frac{\xi}{16} + \eta_N\right), \end{aligned}$$

where $\{\eta_j\}_{j=0}^N$ are to be chosen so as to satisfy $\sum_{j=0}^N \eta_j \leq \xi/16$. Define

$$\begin{aligned} R_1 &= |\mathcal{F}_0| \sup_{f \in \mathcal{F}} P\left(|v_n(f)| > \frac{7}{8}\xi\right), \\ R_2 &= \sum_{j=0}^{N-1} |\mathcal{F}_j| |\mathcal{F}_{j+1}| \sup_{f \in \mathcal{F}} P(|v_n(f_j - f_{j+1})| > \eta_j), \\ R_3 &= P\left(\sup_{f \in \mathcal{F}} |v_n(f_N - f)| > \frac{\xi}{16} + \eta_N\right). \end{aligned}$$

We now choose $\{\delta_j\}_{j=0}^N$, $\{\eta_j\}_{j=0}^N$ and N to make the three terms sufficiently small.

First we study R_1 . Set δ_0 such that $H(\delta_0) = \xi^2/4$. Then, applying Hoeffding's inequality, we have

$$R_1 \leq 2 |\mathcal{F}_0| \exp\left(-2 \left(\frac{7}{8}\xi\right)^2\right) \leq 2 \exp(-\xi^2),$$

where we use the fact that $|\mathcal{F}_0| \leq \exp(H(\delta_0)) = \exp(\xi^2/4)$ in the second inequality.

Next, we study R_2 . Since $\|f_j - f_{j+1}\|_{2,P} \leq 2\delta_j$ by construction, applying Bennet's inequality (Lemma B.5) to each $\sup_{f \in \mathcal{F}} P(|v_n(f_j - f_{j+1})| > \eta_j)$ in R_2 leads to

$$R_2 \leq \sum_{j=0}^{N-1} 2 \exp(2H(\delta_{j+1})) \exp(-\psi_1(\eta_j, n, 4\delta_j^2)),$$

where ψ_1 satisfies the properties described in Lemma B.5.

Next, we study R_3 . Given the construction of \mathcal{F}_N ,

$$\begin{aligned} |v_n(f_N - f)| &\leq |v_n(f_N^U - f_N^L)| + 2\sqrt{n} \|f_N^U - f_N^L\|_{1,P} \\ &\leq |v_n(f_N^U - f_N^L)| + 2\sqrt{n}\delta_N. \end{aligned}$$

The last inequality holds because $\|f_N^U - f_N^L\|_{1,P} \leq \|h_N^U - h_N^L\|_{1,P_2}$ and

$$\|h_N^U - h_N^L\|_{1,P_2} \leq \|h_N^U - h_N^L\|_{2,P_2} \leq \delta_N,$$

which holds from Hölder's inequality. Set $\delta_N \leq s \equiv \xi / (32\sqrt{n})$. Then, by the above derivation and Bennet's inequality,

$$\begin{aligned} R_3 &\leq P \left(\sup_{f \in \mathcal{F}} |v_n(f_N^U - f_N^L)| > \eta_N \right) \\ &\leq 2 |\mathcal{F}_N| \exp(-\psi_1(\eta_N, n, \delta_N^2)). \end{aligned}$$

To develop upper bounds on R_2 and R_3 , we consider two distinct cases. First we consider the case $\delta_0 \leq s$. Set $N = 0$ and $\eta_0 = \xi/16$. Then we have that $R_2 = 0$ and

$$R_3 \leq 2 |\mathcal{F}_0| \exp(-\psi_1(\eta_0, n, \delta_0^2)).$$

Since Assumption (49) and $\delta_0 \leq s$ hold, we have

$$2\eta_0 = \frac{\xi}{8} \geq 4\sqrt{n} \left(\frac{\xi}{32\sqrt{n}} \right)^2 \geq 4\sqrt{n}\delta_0^2.$$

Hence by the properties of ψ_1 specified in Lemma B.5,

$$\psi_1(\eta_0, n, \delta_0^2) \geq \frac{1}{4} \psi_1(2\eta_0, n, \delta_0^2) \geq \frac{1}{4} \eta_0 \sqrt{n}.$$

Using $\eta_0 = \xi/16$ and Assumption (49), we obtain

$$\psi_1(\eta_0, n, \delta_0^2) \geq \frac{1}{4} \eta_0 \sqrt{n} = \frac{\xi}{64} \sqrt{n} \geq 2\xi^2.$$

By the definition of δ_0 , we also have $|\mathcal{F}_0| \leq \exp(\xi^2/4)$. Therefore, combining these results gives

$$R_2 + R_3 \leq 2 \exp(-\xi^2).$$

Next we consider the case $\delta_0 > s$. We here apply Lemma B.6, where we let N and $\{\delta_j\}_{j=0}^N$ be as in Lemma B.6 and $t = \delta_0$ and s be as defined above. Let $\eta_j = 8\sqrt{2}\delta_j H(\delta_{j+1})^{1/2}$ for $0 \leq j < N$ and $\eta_N = 8\sqrt{2}\delta_N H(\delta_N)^{1/2}$. Then Lemma B.6 leads to

$$\sum_{j=0}^N \eta_j = 8\sqrt{2} \sum_{j=0}^N H(\delta_{j+1})^{1/2} \leq 64\sqrt{2} \int_{s/4}^{\delta_0} H(u)^{1/2} du.$$

We have that for $0 < s < t$,

$$\int_s^t H(u)^{1/2} du \leq \begin{cases} K^{1/2} \log(1/s) & r = 2 \\ 2K^{1/2} (r-2)^{-1} s^{(2-r)/2} & r > 2. \end{cases}$$

Combining this with Assumption (50), where C_1 and C_1 are set to be sufficiently large, we have

$$\sum_{j=0}^N \eta_j \leq \frac{\xi}{16},$$

which is consistent with our choice of $\{\eta_j\}_j$. Setting C_1 and C_2 sufficiently large, it follows from Assumption (50) that

$$H(s) \leq \frac{\xi\sqrt{n}}{16}.$$

Hence we have

$$\left(\frac{\eta_j}{4\delta_j^2\sqrt{n}} \right)^2 < \frac{8H(s)}{ns^2} \leq 16.$$

Then from the properties of ψ_1 ,

$$\psi_1(\eta_j, n, 4\delta_j^2) \geq \frac{\eta_j^2}{16\delta_j^2}.$$

Using our bound on R_2 , we obtain

$$R_2 \leq \sum_{j=0}^{N-1} 2 \exp \left(2H(\delta_{j+1}) - \frac{\eta_j^2}{16\delta_j^2} \right) \leq \sum_{j=0}^{N-1} 2 \exp(-4^{j+1}H(\delta_0)).$$

Similarly, we obtain

$$R_3 \leq 2 \exp(-4^{N+1}H(\delta_0)).$$

Putting these results together and using Assumption (50), we have

$$R_2 + R_3 \leq \sum_{j=0}^{\infty} 2 \exp(-4^{j+1}H(\delta_0)) \leq C \exp(-\xi^2),$$

where C is a constant that depends only on K and r . □

Lemma B.5 (Bennet's inequality: see Theorem 2.9 in Boucheron et al. (2013)). *Let $\{Z_i\}_{i=1}^n$ be a sequence of independent random vectors with distribution P . Let f be some function taking values in $[0, 1]$ and define*

$$v_n(f) \equiv \frac{1}{\sqrt{n}} \sum_{i=1}^n [f(Z_i) - E_P(f(Z_i))].$$

Then, for any $\xi \geq 0$, the following holds:

$$P^n(|v_n(f)| > \xi) \leq 2 \exp(-\psi_1(\xi, n, a)),$$

where $a = \text{var}(v_n(f))$ and

$$\psi_1(\xi, n, a) = \xi \sqrt{nh} \left(\frac{\xi}{\sqrt{na}} \right),$$

with $h(x) = (1 + x^{-1}) \log(1 + x) - 1$.

Furthermore, ψ_1 has the following two properties:

$$\psi_1(\xi, n, \alpha) \geq \psi_1(C\xi, n, \rho\alpha) \geq C^2 \rho^{-1} \psi_1(\xi, n, \alpha)$$

for $C \leq 1$ and $\rho \geq 1$, and

$$\psi_1(\xi, n, \alpha) \geq \begin{cases} \frac{\xi^2}{4\alpha} & \text{if } \xi < 4\sqrt{n\alpha} \\ \frac{\xi}{2}\sqrt{n} & \text{if } \xi \geq 4\sqrt{n\alpha} \end{cases}.$$

Lemma B.6 (Lemma 3.1 in Alexander (1984)). *Let $H : (0, t] \rightarrow \mathbb{R}^+$ be a decreasing function, and let $0 < s < t$. Set $\delta_0 \equiv t$, $\delta_{j+1} \equiv s \vee \sup\{x \leq \delta_j/2 : H(x) \geq 4H(\delta_j)\}$ for $j \geq 0$, and $N \equiv \min\{j : \delta_j = s\}$. Then*

$$\sum_{j=0}^N \delta_j H(\delta_{j+1})^{1/2} \leq 8 \int_{s/4}^t H(x)^{1/2} dx.$$

C Proof of the results in Section 6

This appendix provides proof of the results in Section 6. Throughout this appendix, we suppose $\mathcal{X} = [0, 1]^{d_x}$ as in Section 6.

We first provide the proof of Lemma 6.1.

Proof of Lemma 6.1. Let μ_X be the Lebesgue measure on \mathcal{X} . From Theorem 8.3.2 in Dudley (1999), $H_1^B(\epsilon, \mathcal{G}_M, \mu_X) \leq K\epsilon^{d_x-1}$ holds for some positive constant K and for all $\epsilon > 0$. Since P_X is absolutely continuous with respect to μ_X and has a density that is bounded from above by A , we have $H_1^B(A^{-1}\epsilon, \mathcal{G}_M, P_X) \leq H_1^B(\epsilon, \mathcal{G}_M, \mu_X)$. Thus result (i) follows by setting $C = A^{-d_x}K$. \square

The following lemma is used in the proof of Theorem 6.3.

Lemma C.1. *Suppose that P_X is absolutely continuous with respect to the Lebesgue measure on \mathcal{X} and has a density that is bounded from above by a finite constant $A > 0$. Then there exists a constant \tilde{C} , which depends only on A , such that*

$$H_1^B(\epsilon, \mathcal{F}_M, P_X) \leq \tilde{C}\epsilon^{-d_x}.$$

holds for all $\epsilon > 0$.

Proof. Transform \mathcal{F}_M into $\tilde{\mathcal{F}}_M = \{(f+1)/2 : f \in \mathcal{F}_M\}$, which is a class of monotonically increasing functions taking values in $[0, 1]$. Following this transformation, $N_1^B(\epsilon, \mathcal{F}_M, P_X) = N_1^B(\epsilon/2, \tilde{\mathcal{F}}_M, P_X)$ holds. Then the result follows by applying Corollary 1.3 in Gao and Wellner (2007) to $\tilde{\mathcal{F}}_M$, in which we set $\tilde{C} = 2^{-d_x}C_2$, where C_2 is the same constant that appears in Corollary 1.3 in Gao and Wellner (2007). Note that this corollary requires that P_X is absolutely continuous with respect to the Lebesgue measure on \mathcal{X} and has a bounded density. \square

The following lemma gives finite upper bounds for two approximation errors:

$$\inf_{f \in \mathcal{B}_k} R_{\phi_h}(f) - \inf_{f \in \mathcal{F}_M} R_{\phi_h}(f) \text{ and } R_{\phi_h}(1\{\cdot \in G_{\hat{f}_B^\dagger}\} - 1\{\cdot \notin G_{\hat{f}_B^\dagger}\}) - R_{\phi_h}(\hat{f}_B^\dagger)$$

in (24) with $(\check{\mathcal{F}}, \tilde{\mathcal{F}}, \hat{f}) = (\mathcal{B}_k, \mathcal{F}_M, \hat{f}_B^\dagger)$.

Lemma C.2. *Let $k_j \geq 1$, for $j = 1, \dots, d_x$, be fixed. Suppose that the density of P_X is bounded from above by some finite constant $A > 0$.*

(i) *The following holds for the approximation error to the best classifier:*

$$\inf_{f \in \mathcal{B}_k} R_{\phi_h}(f) - \inf_{f \in \mathcal{F}_M} R_{\phi_h}(f) \leq 2A \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{4}{\sqrt{k_j}}.$$

(ii) *For $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_k} \hat{R}_{\phi_h}(f)$ such that the associated coefficients of the Bernstein bases take values in $\{-1, 1\}$, the following holds for the approximation error to the step*

function:

$$R_{\phi_h} \left(1 \left\{ \cdot \in G_{\hat{f}_B} \right\} - 1 \left\{ \cdot \notin G_{\hat{f}_B} \right\} \right) - R_{\phi_h}(\hat{f}_B) \leq 2A \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{4}{\sqrt{k_j}}.$$

The two approximation errors have the same upper bound which converges to zero as k_j ($j = 1, \dots, d_x$) increases. The convergence rate is $\max_{j=1, \dots, d_x} \sqrt{(\log k_j)/k_j}$. Note also that the upper bound on the approximation error to the step function does not depend on the sample size n .

The following two lemmas will be used in the proof of Lemma C.2.

Lemma C.3. Let $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_k} \hat{R}_{\phi_h}(f)$, and $\hat{\theta} \equiv \left\{ \hat{\theta}_{j_1 \dots j_{d_x}} \right\}_{j_1=1, \dots, k_1; \dots; j_{d_x}=1, \dots, k_{d_x}}$ be the vector of the coefficients of the Bernstein bases in \hat{f}_B . Let r_1^+ and r_1^- be the smallest non-negative value and the largest negative value in $\hat{\theta}$, respectively.

(i) If all non-negative elements in $\hat{\theta}$ take the same value r_1^+ , let r_2^+ be 1; otherwise, let r_2^+ be the second smallest non-negative value in $\hat{\theta}$. Propose a $(k_1 + 1) \times \dots \times (k_{d_x} + 1)$ -dimensional vector $\tilde{\theta} \equiv \left\{ \tilde{\theta}_{j_1 \dots j_{d_x}} \right\}_{j_1=1, \dots, k_1; \dots; j_{d_x}=1, \dots, k_{d_x}}$ such that, for all j_1, \dots, j_{d_x} , if $\hat{\theta}_{j_1 \dots j_{d_x}} = r_1^+$, $\tilde{\theta}_{j_1 \dots j_{d_x}} = r_2^+$; otherwise, $\tilde{\theta}_{j_1 \dots j_{d_x}} = \hat{\theta}_{j_1 \dots j_{d_x}}$. Then a new classifier

$$\tilde{f}_B(x) \equiv \sum_{j_1=1}^{k_1} \dots \sum_{j_{d_x}=1}^{k_{d_x}} \tilde{\theta}_{j_1 \dots j_{d_x}} (b_{k_1 j_1}(x_1) \times \dots \times b_{k_{d_x} j_{d_x}}(x_{d_x}))$$

minimizes $\hat{R}_{\phi_h}(\cdot)$ over \mathcal{B}_k .

(ii) Similarly, if all negative elements in $\hat{\theta}$ take the same value r_1^- , let r_2^- be -1 ; otherwise, let r_2^- be the second largest negative value in $\hat{\theta}$. Propose a $(k_1 + 1) \times \dots \times (k_{d_x} + 1)$ -dimensional vector $\check{\theta} \equiv \left\{ \check{\theta}_{j_1 \dots j_{d_x}} \right\}_{j_1=1, \dots, k_1; \dots; j_{d_x}=1, \dots, k_{d_x}}$ such that, for all j_1, \dots, j_{d_x} , if $\hat{\theta}_{j_1 \dots j_{d_x}} = r_1^-$, $\check{\theta}_{j_1 \dots j_{d_x}} = r_2^-$; otherwise, $\check{\theta}_{j_1 \dots j_{d_x}} = \hat{\theta}_{j_1 \dots j_{d_x}}$. Then a new classifier

$$\check{f}_B(x) \equiv \sum_{j_1=1}^{k_1} \dots \sum_{j_{d_x}=1}^{k_{d_x}} \check{\theta}_{j_1 \dots j_{d_x}} (b_{k_1 j_1}(x_1) \times \dots \times b_{k_{d_x} j_{d_x}}(x_{d_x}))$$

minimizes $\hat{R}_{\phi_h}(\cdot)$ over \mathcal{B}_k .

(iii) A classifier

$$\hat{f}_B^\dagger(x) \equiv \sum_{j_1=1}^{k_1} \dots \sum_{j_{d_x}=1}^{k_{d_x}} \text{sign} \left(\hat{\theta}_{j_1 \dots j_{d_x}} \right) \cdot (b_{k_1 j_1}(x_1) \times \dots \times b_{k_{d_x} j_{d_x}}(x_{d_x}))$$

minimizes $\hat{R}_{\phi_h}(\cdot)$ over \mathcal{B}_k .

Proof. First, note that $\tilde{\theta}, \check{\theta} \in \tilde{\Theta}$ holds by construction. We now prove (i). The proof of (ii) follows using a similar argument. Define

$$L_n(\theta) \equiv \sum_{i=1}^n \left\{ Y_i \cdot \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \theta_{j_1 \dots j_{d_x}} \sum_{i=1}^n (b_{k_1 j_1}(X_{1i}) \times \cdots \times b_{k_{d_x} j_{d_x}}(X_{d_x i})) \right\}.$$

Minimization of $\hat{R}_{\phi_h}(\cdot)$ over $\mathcal{B}_{\mathbf{k}}$ is equivalent to the maximization of $L_n(\cdot)$ over $\tilde{\Theta}$. Thus, $\hat{\theta}$ maximizes $L_n(\cdot)$ over $\tilde{\Theta}$.

We prove the result by contradiction. Suppose $\tilde{\theta} \notin \arg \max_{\theta \in \tilde{\Theta}} L_n(\theta)$. Let

$$J_1 \equiv \left\{ (j_1, \dots, j_{d_x}) : \hat{\theta}_{j_1 \dots j_{d_x}} = r_1^+ \right\}.$$

Then,

$$L_n(\tilde{\theta}) - L_n(\hat{\theta}) = \sum_{(j_1, \dots, j_{d_x}) \in J_1} \left\{ (r_2^+ - r_1^+) \sum_{i=1}^n Y_i (b_{k_1 j_1}(X_{1i}) \times \cdots \times b_{k_{d_x} j_{d_x}}(X_{d_x i})) \right\} < 0.$$

Since $r_2^+ - r_1^+ \geq 0$, the above inequality implies that there exists some $(j_1, \dots, j_{d_x}) \in J_1$ such that $\sum_{i=1}^n Y_i (b_{k_1 j_1}(X_{1i}) \times \cdots \times b_{k_{d_x} j_{d_x}}(X_{d_x i})) < 0$. For such (j_1, \dots, j_{d_x}) , setting $\hat{\theta}_{j_1 \dots j_{d_x}}$ to r_1^- can increase the value of $L_n(\hat{\theta})$ without violating the constraints in $\tilde{\Theta}$. But this contradicts the requirement that $\hat{\theta}_{j_1 \dots j_{d_x}}$ is non-negative. Therefore, $\tilde{\theta}$ maximizes $L_n(\cdot)$ over $\tilde{\Theta}$, or equivalently \tilde{f}_B minimizes $\hat{R}_{\phi_h}(\cdot)$ over $\mathcal{B}_{\mathbf{k}}$.

Result (iii) is shown by applying Lemma C.3 (i) and (ii) repeatedly to \hat{f}_B . \square

Lemma C.4. Fix $G \in \mathcal{G}$ and $k_j \geq 1$ for $j = 1, \dots, d_x$. Define a classifier

$$f_G(x) \equiv \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \theta_{j_1 \dots j_{d_x}} (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x})),$$

such that, for all j_1, \dots, j_{d_x} ,

$$\theta_{j_1 \dots j_{d_x}} = \begin{cases} 1 & \text{if } (j_1/k_1, \dots, j_{d_x}/k_{d_x}) \in G \\ -1 & \text{if } (j_1/k_1, \dots, j_{d_x}/k_{d_x}) \notin G \end{cases}.$$

Then the following holds:

$$|R_{\phi_h}(f_G) - R_{\phi_h}(1\{\cdot \in G\} - 1\{\cdot \notin G\})| \leq 2A \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{4}{\sqrt{k_j}}.$$

Proof. Define

$$J_{\mathbf{k}} \equiv \{(j_1, \dots, j_{d_x}) : (j_1/k_1, \dots, j_{d_x}/k_{d_x}) \in G\},$$

which is a set of grid points on G . It follows that

$$\begin{aligned} & R_{\phi_h}(f_G) - R_{\phi_h}(1\{\cdot \in G\} - 1\{\cdot \notin G\}) \\ &= \int_{[0,1]^{d_x}} (2\eta(x) - 1) (1\{x \in G\} - 1\{x \notin G\} - B_{\mathbf{k}}(\theta, x)) dP_X(x) \\ &= \int_{[0,1]^{d_x}} (2\eta(x) - 1) 1\{x \in G\} dP_X(x) - \int_{[0,1]^{d_x}} (2\eta(x) - 1) 1\{x \notin G\} dP_X(x) \\ &\quad - \underbrace{\int_{[0,1]^{d_x}} (2\eta(x) - 1) B_{\mathbf{k}}(\theta, x) dP_X(x)}_{(I)}. \end{aligned} \tag{51}$$

(I) can be written as

$$\begin{aligned} (I) &= \int_{[0,1]^{d_x}} (2\eta(x) - 1) \sum_{(j_1, \dots, j_{d_x}) \in J_{\mathbf{k}}} \left(\prod_{v=1}^{d_x} b_{k_v j_v}(x_v) \right) dP_X(x) \\ &\quad - \int_{[0,1]^{d_x}} (2\eta(x) - 1) \sum_{(j_1, \dots, j_{d_x}) \notin J_{\mathbf{k}}} \left(\prod_{v=1}^{d_x} b_{k_v j_v}(x_v) \right) dP_X(x). \end{aligned}$$

Thus,

$$\begin{aligned} (51) &= \int_{[0,1]^{d_x}} (2\eta(x) - 1) \underbrace{\left(1\{x \in G\} - \sum_{(j_1, \dots, j_{d_x}) \in J_{\mathbf{k}}} \left[\prod_{v=1}^{d_x} b_{k_v j_v}(x_v) \right] \right)}_{(II)} dP_X(x) \\ &\quad + \int_{[0,1]^{d_x}} (2\eta(x) - 1) \underbrace{\left(\sum_{(j_1, \dots, j_{d_x}) \notin J_{\mathbf{k}}} \left[\prod_{v=1}^{d_x} b_{k_v j_v}(x_v) \right] - 1\{x \in G^c\} \right)}_{(III)} dP_X(x). \end{aligned}$$

Let $Bin(k_j, x_j)$, $j = 1, \dots, d_x$, be independent binomial variables with parameters (k_j, x_j) . Then, both (II) and (III) are equivalent to

$$\begin{aligned} & \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) 1\{x \in G\} \\ & - \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) 1\{x \in G^c\}. \end{aligned}$$

Hence,

$$(51) = 2 \int_G (2\eta(x) - 1) \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) dP_X(x) \\ - 2 \int_{G^c} \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) dP_X(x),$$

and therefore

$$|R_{\phi_h}(f_G) - R_{\phi_h}(1\{\cdot \in G\} - 1\{\cdot \notin G\})| \\ \leq 2 \underbrace{\int_G \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) dP_X(x)}_{(IV)} \\ + 2 \underbrace{\int_{G^c} \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) dP_X(x)}_{(V)}.$$

We first evaluate (V). Let ϵ be a small positive value which converges to zero as $k_v \rightarrow \infty$. For small $\Delta_v \leq \epsilon/\sqrt{d_x}$, $v = 1, \dots, d_x$, which converges to zero as $k_v \rightarrow \infty$, define

$$\tilde{G}^c \equiv \{x \in G^c : (x_1 + \Delta_1, \dots, x_{d_x} + \Delta_{d_x}) \in G^c\}.$$

This set is either nonempty or empty. We consider these cases separately. First, suppose that \tilde{G}^c is nonempty. For each $x \in \tilde{G}^c$, let

$$(j_1(x), \dots, j_{d_x}(x)) \in \arg \min_{(j_1, \dots, j_{d_x}) \in J_{\mathbf{k}}: j_1/k_1 \geq x_1 + \Delta_1, \dots, j_{d_x}/k_{d_x} \geq x_{d_x} + \Delta_{d_x}} \|x - (j_1/k_1, \dots, j_{d_x}/k_{d_x})\|.$$

Then

$$(V) \leq \int_{G^c \setminus \tilde{G}^c} \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) dP_X(x) \\ + \int_{\tilde{G}^c} \left(\sum_{v=1}^{d_x} \Pr(Bin(k_{d_x}, x_{d_x}) \geq j_v(x)) \right) dP_X(x) \\ \leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}) \\ + \int_{\tilde{G}^c} \sum_{v=1}^{d_x} \exp \left\{ -2k_v \left(\frac{j_v(x)}{k_v} - x_v \right)^2 \right\} dP_X(x) \\ \leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}) + \sum_{v=1}^{d_x} \int_{\tilde{G}^c} \exp(-2k_v \Delta_v^2) dP_X(x).$$

To obtain the second inequality, we apply Hoeffding's inequality to $\Pr(Bin(k_{d_x}, x_{d_x}) \geq j_v(x))$,

which is applicable since $k_v x_v \leq j_v(x)$ for each $x \in \tilde{G}^c$, and use the following:

$$\begin{aligned} & \int_{G^c \setminus \tilde{G}^c} \Pr((\text{Bin}(k_1, x_1), \dots, \text{Bin}(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) dP_X(x) \\ & \leq A \int_{G^c \setminus \tilde{G}^c} dx \leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}), \end{aligned}$$

where the second inequality holds because $\int_{G^c \setminus \tilde{G}^c} dx$ is bounded from above by $\sum_{v=1}^{d_x} \Delta_v - (d_x - 1) \prod_{v=1}^{d_x} \Delta_v$, which is derived when $G^c = \mathcal{X}$. The last inequality follows from the fact that $j_v(x)/k_v - x_v \geq \Delta_v$ for all $v = 1, \dots, d_x$ and $x \in \tilde{G}$.

Next, we consider the case that \tilde{G}^c is empty. In this case,

$$\begin{aligned} (V) &= \int_{G^c} \Pr((\text{Bin}(k_1, x_1), \dots, \text{Bin}(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) dP_X(x) \\ &\leq \int_{G^c} dP_X(x) \leq A \cdot \max_{v=1, \dots, d_x} \Delta_v. \end{aligned}$$

The inequality follows because $\int_{G^c} dx$ is bounded from above by $\max_{v=1, \dots, d_x} \Delta_v$ when \tilde{G}^c is empty. Therefore, regardless of whether \tilde{G}^c is empty or not, we have

$$(V) \leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}) + \sum_{v=1}^{d_x} \int_{\tilde{G}^c} \exp(-2k_v \Delta_v^2) dP_X(x).$$

Set $\Delta_v = \sqrt{\log k_v} / (2\sqrt{k_v})$ for each $v = 1, \dots, d_x$. Then we have

$$\begin{aligned} (V) &\leq \frac{A}{2} \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \exp\left(-\frac{1}{2} \log k_v\right) \\ &= \frac{A}{2} \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \frac{1}{\sqrt{k_v}}. \end{aligned}$$

Next, we evaluate (IV). For small $\Delta_v \leq \epsilon/\sqrt{d_x}$, $v = 1, \dots, d_x$, which converges to zero as $k_v \rightarrow \infty$, define

$$\tilde{G} \equiv \{x \in G : (x_1 - \Delta_1, \dots, x_{d_x} - \Delta_{d_x}) \in G\}.$$

We again separately consider the two cases: \tilde{G} is nonempty or empty. First, suppose that \tilde{G} is nonempty. For each $x \in \tilde{G}$, let

$$(\tilde{j}_1(x), \dots, \tilde{j}_{d_x}(x)) \in \arg \min_{(j_1, \dots, j_{d_x}) \in (J_{\mathbf{k}})^c: j_1/k_1 \leq x_1 - \Delta_1, \dots, j_{d_x}/k_{d_x} \leq x_{d_x} - \Delta_{d_x}} \|x - (j_1/k_1, \dots, j_{d_x}/k_{d_x})\|.$$

Then,

$$\begin{aligned}
(IV) &\leq \int_{G \setminus \tilde{G}} \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) dP_X(x) \\
&\quad + \int_{\tilde{G}} \left(\sum_{v=1}^{d_x} \Pr(Bin(k_{d_x}, x_{d_x}) \leq \tilde{j}_v(x)) \right) dP_X(x) \\
&\leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}) + \int_{\tilde{G}} \left(\sum_{v=1}^{d_x} \exp \left\{ -2k_v \left(x_v - \frac{\tilde{j}_v(x)}{k_v} \right)^2 \right\} \right) dP_X(x) \\
&\leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}) + \sum_{v=1}^{d_x} \int_{\tilde{G}} \exp(-2k_v \Delta_v^2) dP_X(x).
\end{aligned}$$

The second inequality follows from Hoeffding's inequality and the fact that

$$\begin{aligned}
&\int_{G \setminus \tilde{G}} \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) dP_X(x) \\
&\leq A \int_{G \setminus \tilde{G}} dx \leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}),
\end{aligned}$$

where the inequality holds because $\int_{G \setminus \tilde{G}} dx$ attains its largest value, $\sum_{v=1}^{d_x} \Delta_v - \prod_{v=1}^{d_x} \Delta_v$, when $G = \mathcal{X}$. The last inequality follows from the fact that $j_v(x)/k_v \leq x_v - \Delta_v$ for all $v = 1, \dots, d_x$ and $x \in \tilde{G}$.

Next, we consider the case that \tilde{G} is empty. In this case,

$$\begin{aligned}
(IV) &= \int_G \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) dP_X(x) \\
&\leq \int_G dP_X(x) \leq A \cdot \max_{v=1, \dots, d_x} \Delta_v,
\end{aligned}$$

where the inequality follows because $\int_G dx$ is bounded from above by $\max_{v=1, \dots, d_x} \Delta_v$ when \tilde{G} is empty. Therefore, regardless of whether \tilde{G} is empty or not, we have

$$(IV) \leq A \cdot (\Delta_1 + \dots + \Delta_{d_x}) + \sum_{v=1}^{d_x} \int_{\tilde{G}} \exp(-2k_v \Delta_v^2) dP_X(x).$$

Set $\Delta_v = \sqrt{\log k_v} / (2\sqrt{k_v})$ for each $v = 1, \dots, d_x$. Then, we have

$$\begin{aligned}
(IV) &\leq \frac{A}{2} \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \exp\left(-\frac{1}{2} \log k_v\right) \\
&= \frac{A}{2} \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \frac{1}{\sqrt{k_v}}.
\end{aligned}$$

Consequently, combining above the results, we obtain

$$|R_{\phi_h}(f_G) - R_{\phi_h}(1\{\cdot \in G\} - 1\{\cdot \notin G\})| \leq 2A \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \frac{4}{\sqrt{k_v}}.$$

□

Finally, the following provides the proof of Lemma C.2.

Proof of Lemma C.2. We first prove (i). Let G^* minimize $\mathcal{R}(\cdot)$ over \mathcal{G}_M . From Theorem 4.4, a classifier $\tilde{f}^*(x) \equiv 1\{x \in G^*\} - 1\{x \in (G^*)^c\}$ minimizes the hinge risk $R_{\phi_h}(\cdot)$ over \mathcal{F}_M . Define a vector $\theta^* = \{\theta_{j_1 \dots j_d}^*\}_{j_1=0, \dots, k_1; \dots; j_d=0, \dots, k_d}$ such that for each j_1, \dots, j_d ,

$$\theta_{j_1 \dots j_d}^* = \begin{cases} 1 & \text{if } (j_1/k_1, \dots, j_d/k_d) \in G^* \\ -1 & \text{otherwise.} \end{cases}$$

Note that θ^* is contained by $\tilde{\Theta}$. Thus, it follows that

$$\inf_{f \in \mathcal{B}_{\mathbf{k}}} R_{\phi_h}(f) - R_{\phi_h}(\tilde{f}^*) \leq R_{\phi_h}(B_{\mathbf{k}}(\theta^*, \cdot)) - R_{\phi_h}(\tilde{f}^*).$$

Then, applying Lemma C.4 to $R_{\phi_h}(B_{\mathbf{k}}(\theta^*, \cdot)) - R_{\phi_h}(\tilde{f}^*)$ yields result (i).

Result (ii) follows immediately from Lemmas C.3 (iii) and C.4. □

D Statistical properties of weighted classification and their application to weighted classification

D.1 Statistical properties of weighted classification with hinge loss

This section extends the analysis of Section 5 to weighted classification with hinge losses. Let $\{(\omega_i, Y_i, X_i) : i = 1, \dots, n\}$ be a sample of observations that are i.i.d. as (ω, Y, X) . Let P^n denote the joint distribution of the sample of n observations. Given the sample, the empirical weighted classification risk and hinge risk for a classifier f are defined as

$$\hat{R}^\omega(f) \equiv n^{-1} \sum_{i=1}^n \omega_i 1\{Y_i \cdot \text{sign}(f(X_i)) \leq 0\},$$

$$\hat{R}_{\phi_h}^\omega(f) \equiv n^{-1} \sum_{i=1}^n \omega_i \max\{0, 1 - Y_i f(X_i)\},$$

respectively. Let $\check{\mathcal{F}}$ be a subclass of $\mathcal{F}_{\mathcal{G}}$, on which we learn a best classifier, and $\tilde{\mathcal{F}}_{\mathcal{G}}$ be a constrained classification-preserving reduction of $\mathcal{F}_{\mathcal{G}}$ in the sense of Definition 4.2.

As an analogue of Theorems 5.1 and 5.2, the following theorem gives general upper bounds on the mean of the \mathcal{G} -constrained excess weighted classification risk.

Theorem D.1. *Suppose that $\tilde{\mathcal{F}}_{\mathcal{G}}$ is a subclass of $\mathcal{F}_{\mathcal{G}}$ satisfying conditions (A1) and (A2) in Theorem 4.4. Let $\hat{f} \in \arg \inf_{f \in \check{\mathcal{F}}} \hat{R}_{\phi_h}^{\omega}(f)$, and $(q_n, \tau_n, L_C(r, n))$ be as in Theorem 5.1.*

(i) *Let \mathcal{P} be a class of distributions on $\mathbb{R}_+ \times \{-1, 1\} \times \mathcal{X}$ such that, for any $P \in \mathcal{P}$, Condition 7.4 holds and there exist positive constants C and r for which condition (19) holds for all $\epsilon > 0$ or condition (20) holds for all $\epsilon > 0$. Then if $\check{\mathcal{F}} = \tilde{\mathcal{F}}_{\mathcal{G}}$,*

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[R^{\omega}(\hat{f}) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R^{\omega}(f) \right] \leq ML_C(r, n).$$

(ii) *Suppose that \mathcal{P} is a class of distributions on $\mathbb{R}_+ \times \{-1, 1\} \times \mathcal{X}$ such that, for any $P \in \mathcal{P}$, Condition 7.4 holds and there exist positive constants C' and r' for which condition (23) holds for all $\epsilon > 0$. Then the following holds:*

$$\begin{aligned} \sup_{P \in \mathcal{P}} E_{P^n} \left[R^{\omega}(\hat{f}) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R^{\omega}(f) \right] &\leq ML_{C'}(r', n) + \frac{1}{2} \left(\inf_{f \in \check{\mathcal{F}}} R_{\phi_h}^{\omega}(f) - \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}^{\omega}(f) \right) \\ &\quad + \frac{1}{2} \left(R_{\phi_h}^{\omega}(\tilde{f}_{\mathcal{G}_f}) - R_{\phi_h}^{\omega}(\hat{f}) \right). \end{aligned} \quad (52)$$

Proof. See Appendix E. □

Similar comments as in Remark 5.3 apply to Theorem D.1 (ii). The two approximation errors in (52) are small as $\check{\mathcal{F}}$ is a good approximation of $\tilde{\mathcal{F}}_{\mathcal{G}}$.

D.2 Monotone weighted classification

Finally, we extend the results for monotone classification in Section 6 to weighted classification. Let $\mathcal{F}_{\mathcal{G}_M}$, \mathcal{F}_M , and $\mathcal{B}_{\mathbf{k}}$ be as in Section 6, and suppose $\mathcal{X} = [0, 1]^{d_x}$. Our aim is to find a best classifier that minimizes $R^{\omega}(\cdot)$ over $\mathcal{F}_{\mathcal{G}_M}$. We again consider using the whole class of monotone classifiers \mathcal{F}_M and sieve of Bernstein polynomials $\mathcal{B}_{\mathbf{k}}$ in the empirical hinge risk minimization for weighted classification.

The following theorem shows the statistical properties of monotone weighted classification using \mathcal{F}_M .

Theorem D.2. *Let \mathcal{P} be a class of distributions of (ω, Y, X) such that Condition 7.4 holds for any $P \in \mathcal{P}$, and that for any $P \in \mathcal{P}$ the marginal distribution P_X is absolutely*

continuous with respect to the Lebesgue measure on \mathcal{X} and has a density that is bounded from above by some finite constant $A > 0$. Let q_n and τ_n be as in Theorem 6.2, and let $\hat{f}_M \in \arg \inf_{f \in \mathcal{F}_M} \hat{R}_{\phi_h}^\omega(f)$. Then the following holds:

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[R^\omega(\hat{f}_M) - \inf_{f \in \mathcal{F}_{\mathcal{G}_M}} R^\omega(f) \right] \leq \begin{cases} 2MD_1\tau_n + 4MD_2 \exp(-D_1^2 q_n^2) & \text{if } d_x \geq 2 \\ 2MD_3\tau_n + 2Mn^{-1}D_4 & \text{if } d_x = 1 \end{cases}$$

for some positive constants D_1 and D_2 , which depend only on d_x and A .

Proof. Since \mathcal{F}_M satisfies conditions (A1) and (A2) in Theorem 4.4 with \mathcal{G} being \mathcal{G}_M (Example 4.7), the result follows from Theorem D.1 (i) and Lemma 6.1. \square

The classifier \hat{f}_M that minimizes $R_{\phi_h}^\omega(\cdot)$ over \mathcal{F}_M can be obtained by solving the linear program in (25) with $\sum_{i=1}^n Y_i f_i$ replaced by $\sum_{i=1}^n \omega_i Y_i f_i$.

For monotone weighted classification using Bernstein polynomials $\mathcal{B}_{\mathbf{k}}$, similarly to Section 6.2, we propose using

$$\hat{f}_B^\dagger(x) \equiv \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \text{sign} \left(\hat{\theta}_{j_1 \dots j_{d_x}} \right) \cdot (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x})),$$

where $\{\hat{\theta}_{j_1 \dots j_{d_x}}\}_{j_1=0, \dots, k_1; \dots; j_{d_x}=0, \dots, k_{d_x}}$ is the vector of coefficients characterizing some $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_{\mathbf{k}}} \hat{R}_{\phi_h}^\omega(f)$. \hat{f}_B^\dagger converts each estimated coefficient $\hat{\theta}_{j_1 \dots j_{d_x}}$ in \hat{f}_B to either -1 or 1 depending on its sign. The following theorem states the statistical properties of monotone weighted classification using \hat{f}_B^\dagger .

Theorem D.3. *Let \mathcal{P} be a class of distributions of (ω, Y, X) that satisfies the same conditions as in Theorem D.2. Let \tilde{q}_n and $\tilde{\tau}_n$ be as in Theorem 6.3. Then the following holds:*

$$\begin{aligned} \sup_{P \in \mathcal{P}} E_{P^n} \left[R^\omega(\hat{f}_B^\dagger) - \inf_{f \in \mathcal{F}_{\mathcal{G}_M}} R^\omega(f) \right] &\leq 2MD_1\tilde{\tau}_n + 4MD_2 \exp(-D_1^2 \tilde{q}_n^2) \\ &+ 4MA \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{8M}{\sqrt{k_j}}, \end{aligned}$$

where D_1 and D_2 are the same constants as in Theorem D.2, which depend only on d_x and A .

Proof. The result follows by combining Theorem D.1 (ii), Lemma C.1 in Appendix C, and Lemma E.5 in Appendix E. \square

Similar comments as those in Section 6 apply to Theorems D.2 and D.3. Using \mathcal{F}_M leads to the faster convergence rate than using the Bernstein polynomials \mathcal{B}_k . When using the Bernstein polynomials \mathcal{B}_k , to achieve the convergence rate of $\tilde{\tau}_n$ for the excess risk, Theorem D.3 suggests setting the tuning parameters k_j , $j = 1, \dots, d_x$, sufficiently large so that $\sqrt{\log k_j/k_j} = O(\tilde{\tau}_n)$.

E Proof of the results in Section 7 and Appendix D

This section provides proof of the results in Section 7 and Appendix D for weighted classification. Most of the various proofs are natural extensions of the proofs of the results in Sections 3–6. For simplicity of notation, define a function

$$L(\omega_+, \omega_-, \eta) \equiv -\omega_+ \eta + \omega_-(1 - \eta),$$

the right hand side of which appears in (28). Similar to (9) and (10), the following expressions are useful for the proofs in this appendix:

$$\mathcal{R}^\omega(G) = \int_G L(\omega_+(x), \omega_-(x), \eta(x)) dP_X(x) + \int_{\mathcal{X}} \omega_+(x) \eta(x) dP_X(x), \quad (53)$$

$$\mathcal{R}_\phi^\omega(G) = \int_G \Delta C_\phi^\omega(\omega_+(x), \omega_-(x), \eta(x)) dP_X(x) + \int_{\mathcal{X}} C_\phi^{\omega^-}(\omega_+(x), \omega_-(x), \eta(x)) dP_X(x). \quad (54)$$

We first state the proofs of Theorem 7.1 and Corollary 7.2.

Proof of Theorem 7.1. We first prove the ‘if’ part of the first statement. Fix P , \mathcal{G} , and $G_1, G_2 \in \mathcal{G}$. By equation (54),

$$\begin{aligned} \mathcal{R}_\phi^\omega(G_1) - \mathcal{R}_\phi^\omega(G_2) &= \int_{G_1 \setminus G_2} \Delta C_\phi^\omega(\omega_+(x), \omega_-(x), \eta(x)) dP_X(x) \\ &\quad - \int_{G_2 \setminus G_1} \Delta C_\phi^\omega(\omega_+(x), \omega_-(x), \eta(x)) dP_X(x). \end{aligned}$$

Thus, $\mathcal{R}_{\phi_1}^\omega(G_1) \geq \mathcal{R}_{\phi_1}^\omega(G_2)$ is equivalent to

$$\int_{G_1 \setminus G_2} \Delta C_{\phi_1}^\omega(\omega_+(x), \omega_-(x), \eta(x)) dP_X(x) \geq \int_{G_2 \setminus G_1} \Delta C_{\phi_1}^\omega(\omega_+(x), \omega_-(x), \eta(x)) dP_X(x).$$

Replacing $\Delta C_{\phi_1}^\omega$ by $\Delta C_{\phi_2}^\omega = c \Delta C_{\phi_1}^\omega$ with $c > 0$ does not change the above inequality. Moreover, replacing $\Delta C_{\phi_1}^\omega$ with $\Delta C_{\phi_2}^\omega$ in the above inequality is equivalent to $\mathcal{R}_{\phi_2}^\omega(G_1) \geq \mathcal{R}_{\phi_2}^\omega(G_2)$. Therefore, since the above discussion holds for any P , \mathcal{G} , and $G_1, G_2 \in \mathcal{G}$, the

‘if’ part of the first statement holds.

The ‘only if’ part of the first statement of the theorem follows directly from Theorem 3.6 upon setting $\Delta C_\phi^\omega(\omega_+(x), \omega_-(x), \eta(x)) = \Delta C_\phi(\eta(x))$, or equivalently $\omega_+(x) = \omega_-(x) = 1$, for all $x \in \mathcal{X}$.

We next prove the second statement, or equivalently that $\Delta C_{\phi_{01}}^\omega(\omega_+, \omega_-, \eta) = L(\omega_+, \omega_-, \eta)$ holds for all $(\omega_+, \omega_-, \eta) \in \mathbb{R}_+ \times \mathbb{R}_+ \times [0, 1]$. For $f \in \mathcal{F}_G$, it follows that

$$\begin{aligned} C_{\phi_{01}}^\omega(\omega_+, \omega_-, f, \eta) &= \omega_+ 1\{f \leq 0\} \eta + \omega_- 1\{f > 0\} (1 - \eta), \\ \Delta C_{\phi_{01}}^{w+}(\omega_+, \omega_-, f, \eta) &= \omega_-(1 - \eta), \\ \Delta C_{\phi_{01}}^{w-}(\omega_+, \omega_-, f, \eta) &= \omega_+ \eta. \end{aligned}$$

Thus, we have

$$\Delta C_{\phi_{01}}^\omega(\omega_+, \omega_-, f, \eta) = \Delta C_{\phi_{01}}^{w+}(\omega_+, \omega_-, f, \eta) - \Delta C_{\phi_{01}}^{w-}(\omega_+, \omega_-, f, \eta) = L(\omega_+, \omega_-, \eta).$$

Finally, we prove the last statement. For the hinge loss function $\phi_h(\alpha) = c \max\{0, 1 - \alpha\}$ and $f \in \mathcal{F}_G$, we have

$$C_{\phi_h}(\omega_+, \omega_-, f, \eta) = c(\omega_+(1 - f)\eta + \omega_-(1 + f)(1 - \eta)).$$

Hence, we obtain

$$\begin{aligned} \Delta C_{\phi_h}^{w+}(\omega_+, \omega_-, \eta) &= \begin{cases} 2c\omega_-(1 - \eta) & \text{for } L(\omega_+, \omega_-, \eta) < 0 \\ c(\omega_+\eta + \omega_-(1 - \eta)) & \text{for } L(\omega_+, \omega_-, \eta) \geq 0 \end{cases}, \\ \Delta C_{\phi_h}^\omega(\omega_+, \omega_-, \eta) &= \begin{cases} c(\omega_+\eta + \omega_-(1 - \eta)) & \text{for } L(\omega_+, \omega_-, \eta) < 0 \\ 2c\omega_+\eta & \text{for } L(\omega_+, \omega_-, \eta) \geq 0 \end{cases}. \end{aligned}$$

Therefore, $\Delta C_{\phi_h}^\omega(\omega_+, \omega_-, \eta) = cL(\omega_+, \omega_-, \eta)$ holds for all $(\omega_+, \omega_-, \eta) \in \mathbb{R}_+ \times \mathbb{R}_+ \times [0, 1]$. \square

Proof of Corollary 7.2. Equation (29) follows from

$$\begin{aligned} R^\omega(f) - \inf_{f \in \mathcal{F}_G} R^\omega(f) &= \mathcal{R}^\omega(G_f) - \mathcal{R}^\omega(G^*) \\ &= \int_{\mathcal{X}} L(\omega_+(x), \omega_-(x), \eta(x)) (1\{x \in G_f\} - 1\{x \in G^*\}) dP_X(x) \\ &= c^{-1} \int_{\mathcal{X}} \Delta C_\phi^\omega(\omega_+(x), \omega_-(x), \eta(x)) (1\{x \in G_f\} - 1\{x \in G^*\}) dP_X(x) \\ &= c^{-1} (\mathcal{R}_\phi^\omega(G_f) - \mathcal{R}_\phi^\omega(G^*)) = c^{-1} \left(\inf_{\tilde{f} \in \mathcal{F}_{G_f}} R_\phi^\omega(\tilde{f}) - \inf_{f \in \mathcal{F}_G} R_\phi^\omega(f) \right) \end{aligned}$$

$$\leq c^{-1} \left(R_{\phi}^{\omega}(f) - \inf_{f \in \mathcal{F}_{\mathcal{G}}} R_{\phi}^{\omega}(f) \right),$$

where the first equality follows from (53); the second equality follows from the assumption; the third equality follows from (54). \square

We next provide the proof of Theorem 7.5. Beforehand, note that the weighted hinge risk can be expressed as

$$\begin{aligned} R_{\phi_h}^{\omega}(f) &= \int_{\mathcal{X}} (\omega_+(x) (1 - f(x)) \eta(x) + \omega_-(x) (1 + f(x)) (1 - \eta(x))) dP_X(x) \\ &= \int_{\mathcal{X}} L(\omega_+(x), \omega_-(x), \eta(x)) f(x) dP_X(x) + E_P[\omega]. \end{aligned} \quad (55)$$

Moreover, for $G \in \mathcal{G}$, $\mathcal{R}(G)$ can be written as

$$\begin{aligned} \mathcal{R}^{\omega}(G) &= \int_{\mathcal{X}} (\omega_+(x) \eta(x) 1\{x \in G^c\} + \omega_-(x) (1 - \eta(x)) 1\{x \in G\}) dP_X(x) \\ &= - \int_{G^c} L(\omega_+(x), \omega_-(x), \eta(x)) (1 - \eta(x)) dP_X(x) \\ &\quad + \int_{\mathcal{X}} \omega_-(x) (1 - \eta(x)) dP_X(x). \end{aligned} \quad (56)$$

The following lemma, which is an analogue of Lemma 4.3, will be used in the proof of Theorem 7.5.

Lemma E.1. *Let $\mathcal{G} \subseteq 2^{\mathcal{X}}$ be a class of measurable subsets and $\tilde{\mathcal{F}}_{\mathcal{G}, J}$ be defined as in (15). (i) Let \tilde{f}^* be a minimizer of the weighted hinge risk $R_{\phi_h}^{\omega}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}, J}$. Then \tilde{f}^* minimizes the weighted classification risk $R^{\omega}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}, J}$, and leads to $R_{\phi_h}^{\omega}(\tilde{f}^*) = 2\mathcal{R}^{\omega}$. (ii) For $G^* \in \mathcal{G}^*$, \tilde{f}_{G^*} minimizes $R_{\phi_h}^{\omega}(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G}, J}$.*

Proof. Let $\tilde{f} \in \tilde{\mathcal{F}}_{\mathcal{G}, J}$. The classifier \tilde{f} has the form of

$$\tilde{f}(x) = 2 \sum_{j=1}^J c_j 1\{x \in G_j\} - 1$$

for some $G_1, \dots, G_J \in \mathcal{G}$ such that $G_J \subseteq \dots \subseteq G_1$ and some $c_j \geq 0$ for $j = 1, \dots, J$ with $\sum_{j=1}^J c_j = 1$. Substituting \tilde{f} into (55) yields

$$R_{\phi_h}^{\omega}(\tilde{f}) = 2 \sum_{j=1}^J c_j \int_{G_j} L(\omega_+(x), \omega_-(x), \eta(x)) (1 - \eta(x)) dP_X(x) + 2 \int_{\mathcal{X}} \omega_+(x) \eta(x) dP_X(x).$$

Comparing this expression with equation (53),

$$R_{\phi_h}^\omega(\tilde{f}) = 2 \sum_{j=1}^J c_j \mathcal{R}^\omega(G_j). \quad (57)$$

From this expression and the assumption that $\sum_{j=1}^J c_j = 1$, $R_{\phi_h}^\omega(f) \geq 2\mathcal{R}^{w*}$ for any $f \in \tilde{\mathcal{F}}_{\mathcal{G},J}$.

For $G^* \in \mathcal{G}^*$, \tilde{f} is equivalent to \tilde{f}_{G^*} when $G_1 = G^*$ and $c_j = 1$. Thus, from equation (57), $R_{\phi_h}^\omega(\tilde{f}_{G^*}) = 2\mathcal{R}^{w*}$; that is, \tilde{f}_{G^*} minimizes $R_{\phi_h}^\omega$ over $\tilde{\mathcal{F}}_{\mathcal{G},J}$. This proves statement (ii) of the lemma.

We will next prove that the minimizer \tilde{f}^* of $R_{\phi_h}^\omega(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G},J}$ also minimizes $R^\omega(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G},J}$. Let \tilde{f}^* be denoted by

$$\tilde{f}^*(x) = 2 \sum_{j=1}^J c_j 1\{x \in G_j\} - 1 \quad (58)$$

for some $G_1, \dots, G_J \subseteq \mathcal{G}$ such that $G_J \subseteq \dots \subseteq G_1$ and some $c_j \geq 0$ for $j = 1, \dots, J$ such that $\sum_{j=1}^J c_j = 1$. To obtain a contradiction, suppose \tilde{f}^* does not minimize $R^\omega(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G},J}$. As \tilde{f}^* does not minimize $R^\omega(\cdot)$ over $\tilde{\mathcal{F}}_{\mathcal{G},J}$, $G_{\tilde{f}^*} \notin \mathcal{G}^*$. Letting m be the smallest number in $\{1, \dots, J\}$ such that $\sum_{j=1}^m c_j \geq 1/2$, $\tilde{G}_m = G_{\tilde{f}^*}$. Then

$$\begin{aligned} R_{\phi_h}^\omega(\tilde{f}) &= 2 \sum_{j=1}^J c_j \mathcal{R}^\omega(G_j) \\ &\geq 2c_m \mathcal{R}^\omega(G_{\tilde{f}^*}) + 2(1 - c_m) \mathcal{R}^{w*} \\ &> 2\mathcal{R}^{w*} \end{aligned}$$

where the last line comes from $c_m^- > 0$ and $G_{\tilde{f}^*} \notin \mathcal{G}^*$. Since $\inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G},J}} R_{\phi_h}^\omega(f) = 2\mathcal{R}^{w*}$, this contradicts the assumption that \tilde{f}^* minimizes $R_{\phi_h}^\omega$ over $\tilde{\mathcal{F}}_{\mathcal{G},J}$. \square

We are now equipped to give proof of Theorem 7.5.

Proof of Theorem 7.5. Given $\tilde{f}^* \in \arg \inf_{f \in \tilde{\mathcal{F}}_{\mathcal{G}}} R_{\phi_h}^\omega(f)$, let $\tilde{\mathcal{F}}_J^*$ be a class of classifiers as in the proof of Theorem 4.4. Note that $G_f = G_{\tilde{f}^*}$ for any $f \in \tilde{\mathcal{F}}_J^*$. Similarly to the proof of Theorem 4.4, we define a sequence of classifiers $\{\tilde{f}_J^*\}_{J=1}^\infty$ such that

$$\tilde{f}_J^*(x) \equiv \frac{2}{J} \sum_{j=1}^J 1\{\tilde{f}^*(x) \geq 2(j/J) - 1\} - 1.$$

We show in the proof of Theorem 4.4 that $\tilde{f}_J^* \rightarrow \tilde{f}^*$ as $J \rightarrow \infty$, P_X -almost everywhere.

Then

$$\begin{aligned}
R_{\phi_h}^\omega(\tilde{f}^*) &= \int_{\mathcal{X}} L(\omega_+(x), \omega_-(x), \eta(x)) \tilde{f}^*(x) dP_X(x) + E_P[\omega] \\
&= \lim_{J \rightarrow \infty} \int_{\mathcal{X}} \tilde{f}_J^*(x) dP_X(x) + E_P[\omega] \\
&= \lim_{J \rightarrow \infty} R_{\phi_h}^\omega(\tilde{f}_J^*) \geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_J^*} R_{\phi_h}^\omega(\tilde{f}) \\
&\geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_{G,J}} R_{\phi_h}^\omega(\tilde{f}),
\end{aligned} \tag{59}$$

where the first and third equalities follow from (55). The second equality follows from the dominated convergence theorem, which holds as both

$$L(\omega_+(X), \omega_-(X), \eta(X)) \tilde{f}_J^*(X) \rightarrow L(\omega_+(X), \omega_-(X), \eta(X)) \tilde{f}^*(X)$$

and

$$\left| L(\omega_+(X), \omega_-(X), \eta(X)) \tilde{f}_J^*(X) \right| < \infty$$

hold P_X -almost everywhere, where the second condition is satisfied by Condition 7.4. The first inequality follows from $\tilde{f}_J^* \in \tilde{\mathcal{F}}_J^*$, and the last inequality follows from $\tilde{\mathcal{F}}_J^* \subseteq \tilde{\mathcal{F}}_{G,J}$.

Lemma E.1 shows that $\inf_{\tilde{f} \in \tilde{\mathcal{F}}_{G,J}} R_{\phi_h}^\omega(\tilde{f}) = 2\mathcal{R}^{w*}$ for any J . Thus, we have

$$R_{\phi_h}^\omega(\tilde{f}^*) \geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_{G,J}} R_{\phi_h}^\omega(\tilde{f}) = 2\mathcal{R}^{w*}.$$

This means that the minimal value of $R_{\phi_h}^\omega$ on $\tilde{\mathcal{F}}_G$ is at least $2\mathcal{R}^{w*}$. Lemma E.1 also shows that \tilde{f}_{G^*} leads to $R_{\phi_h}^\omega(\tilde{f}_{G^*}) = 2\mathcal{R}^{w*}$. Therefore, \tilde{f}_{G^*} minimizes $R_{\phi_h}^\omega$ over $\tilde{\mathcal{F}}_G$, which proves statement (ii) of the theorem.

We will next prove statement (i) of the theorem. Let $\tilde{f}^* \in \arg \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}^\omega(f)$. To obtain a contradiction, suppose that \tilde{f}^* does not minimize $R^\omega(\cdot)$ over $\tilde{\mathcal{F}}_G$, or equivalently $G_{\tilde{f}^*} \notin \mathcal{G}^*$. Then, from equation (57), for any J and $\tilde{f} \in \tilde{\mathcal{F}}_J^*$, $R_{\phi_h}^\omega(\tilde{f}) > 2\mathcal{R}^{w*}$ holds. Therefore, from equation (59),

$$R_{\phi_h}^\omega(\tilde{f}^*) \geq \lim_{J \rightarrow \infty} \inf_{\tilde{f} \in \tilde{\mathcal{F}}_J^*} R_{\phi_h}^\omega(\tilde{f}) > 2\mathcal{R}^{w*}.$$

This contradicts the assumption that \tilde{f}^* minimizes $R_{\phi_h}^\omega$ over $\tilde{\mathcal{F}}_G$ because $R_{\phi_h}^\omega(\tilde{f}_{G^*}) = 2\mathcal{R}^{w*}$. \square

The following corollary shows a similar relationship between the \mathcal{G} -constrained excess weighted-classification risk and \mathcal{F}_G -constrained excess weighted-hinge risk as is present in

Corollary 4.5.

Corollary E.2. Assume $\tilde{\mathcal{F}}_G$ is a subclass of \mathcal{F}_G satisfying conditions (A1) and (A2) in Theorem 4.4. If ϕ satisfies $\Delta C_\phi^\omega(\omega_+, \omega_-, \eta) = c(-\omega_+ \eta + \omega_-(1 - \eta))$ for some $c > 0$,

$$c(R^\omega(f) - \inf_{f \in \mathcal{F}_G} R^\omega(f)) = \frac{1}{2} \left(R_\phi^\omega(f) - \inf_{f \in \tilde{\mathcal{F}}_G} R_\phi^\omega(f) \right) + \frac{1}{2} \left(R_\phi^\omega(\tilde{f}_{G_f}) - R_\phi^\omega(f) \right)$$

for any $f \in \mathcal{F}_G$.

Proof. By equations (53) and (56), $R^\omega(f)$ can be written as

$$cR^\omega(f) = \frac{c}{2} \left\{ \int_{\mathcal{X}} L(\omega_+(x), \omega_-(x), \eta(x)) \tilde{f}_{G_f}(x) dP_X(x) + E_P[\omega] \right\}.$$

From equation (55), the right-hand side is equal to $2^{-1}R_\phi^\omega(\tilde{f}_{G_f})$. \square

We now give the proof of Theorem D.1, which is an extension of the proof of Theorem 5.1.

Proof of Theorem D.1 (ii). For convenience of notation, we prove the result with C' and r' replaced by C and r , respectively. Let $P \in \mathcal{P}$ be fixed. First of all, Corollary E.2 and decomposing $R_{\phi_h}^\omega(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}^\omega(f)$ gives

$$\begin{aligned} R^\omega(\hat{f}) - \inf_{f \in \mathcal{F}_G} R^\omega(f) &= \frac{1}{2} \left(R_{\phi_h}^\omega(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}^\omega(f) \right) + \frac{1}{2} \left(\inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}^\omega(f) - \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}^\omega(f) \right) \\ &\quad + \frac{1}{2} \left(R_{\phi_h}^\omega(\tilde{f}_{G_f}) - R_{\phi_h}^\omega(\hat{f}) \right). \end{aligned} \quad (60)$$

Hence, to obtain the inequality in (52), we need to prove that

$$R_{\phi_h}^\omega(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}^\omega(f) \leq L_C(r, n).$$

We follow the same strategy as the proof of Theorem B.2. Let \check{f}^* minimizes $R_{\phi_h}^\omega(\cdot)$ over $\tilde{\mathcal{F}}$. A standard argument gives

$$\begin{aligned} E_{P^n} \left[R_{\phi_h}^\omega(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}} R_{\phi_h}^\omega(f) \right] &\leq E_{P^n} \left[R_{\phi_h}^\omega(\hat{f}) - \hat{R}_{\phi_h}^\omega(\hat{f}) + \hat{R}_{\phi_h}^\omega(\check{f}^*) - R_{\phi_h}^\omega(\check{f}^*) \right] \\ &\quad \left(\because \hat{R}_{\phi_h}^\omega(\hat{f}) \leq \hat{R}_{\phi_h}^\omega(\check{f}^*) \right) \\ &= 2E_{P^n} \left[R_{\phi_h}^\omega \left(\frac{\hat{f} + 1}{2} \right) - \hat{R}_{\phi_h}^\omega \left(\frac{\hat{f} + 1}{2} \right) \right] \end{aligned}$$

$$\begin{aligned}
& + 2E_{P^n} \left[\hat{R}_{\phi_h}^\omega \left(\frac{\check{f}^* + 1}{2} \right) - R_{\phi_h}^\omega \left(\frac{\check{f}^* + 1}{2} \right) \right] \\
& \leq 4 \sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| R_{\phi_h}^\omega(f) - \hat{R}_{\phi_h}^\omega(f) \right| \right],
\end{aligned}$$

where $\check{\mathcal{F}} = \{(f+1)/2 : f \in \check{\mathcal{F}}\}$ as in the proof of Theorem B.2.

Note that

$$\begin{aligned}
& \sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| R_{\phi_h}^\omega(f) - \hat{R}_{\phi_h}^\omega(f) \right| \right] \\
& = M \sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| E_P \left[\left(\frac{\omega}{M} \right) Y f(x) \right] - \frac{1}{n} \sum_{i=1}^n \left(\frac{\omega_i}{M} \right) Y_i f(X_i) \right| \right] \tag{61}
\end{aligned}$$

and that

$$\sup_{f \in \check{\mathcal{F}}} \left| E_P \left[\left(\frac{\omega}{M} \right) Y f(x) \right] - \frac{1}{n} \sum_{i=1}^n \left(\frac{\omega_i}{M} \right) Y_i f(X_i) \right| \leq 2.$$

We first prove the result for the case of $r \geq 1$. For any $f \in \check{\mathcal{F}}$ and $D > 0$,

$$\begin{aligned}
& \frac{\sqrt{n}}{q_n} \sup_{f \in \check{\mathcal{F}}} E_{P^n} \left[\left| E_P \left[\left(\frac{\omega}{M} \right) Y f(x) \right] - \frac{1}{n} \sum_{i=1}^n \left(\frac{\omega_i}{M} \right) Y_i f(X_i) \right| \right] \\
& \leq D + \frac{2\sqrt{n}}{q_n} P^n \left(\sup_{f \in \check{\mathcal{F}}} \frac{\sqrt{n}}{q_n} \left| E_P \left[\left(\frac{\omega}{M} \right) Y f(x) \right] - \frac{1}{n} \sum_{i=1}^n \left(\frac{\omega_i}{M} \right) Y_i f(X_i) \right| > D \right).
\end{aligned}$$

Then, applying Corollary B.3, setting $Z_1 = (\omega, Y)$, $Z_2 = X$, $g(Z_1) = (\omega/M) \cdot Y$ and $\mathcal{H} = \check{\mathcal{F}}$, shows that there exist $D_1, D_2, D_3 > 0$, depending only on r and C , such that

$$P^n \left(\sup_{f \in \check{\mathcal{F}}} \frac{\sqrt{n}}{q_n} \left| E_P \left[\left(\frac{\omega}{M} \right) Y f(x) \right] - \frac{1}{n} \sum_{i=1}^n \left(\frac{\omega_i}{M} \right) Y_i f(X_i) \right| > D \right) \leq D_2 \exp(-D^2 q_n^2),$$

for $D_1 \leq D \leq D_3 \sqrt{n}/q_n$. Therefore, when $r \geq 1$, we have

$$\tau_n^{-1} E_{P^n} \left[R_{\phi_h}^\omega(\hat{f}) - \inf_{f \in \check{\mathcal{F}}} R_{\phi_h}^\omega(f) \right] \leq 4MD_1 + 8M\tau_n^{-1} D_2 \exp(-D_1^2 q_n^2).$$

Combining this result with (60) leads to the inequality in (52) for the case of $r \geq 1$.

The inequality in (52) for the case of $r < 1$ follows immediately by applying Lemma B.4 to equation (61). \square

Proof of Theorem D.1 (i). Let $P \in \mathcal{P}$ be fixed. We follow the same strategy as in the proof of Theorem 5.1. Define $\hat{f}^\dagger(x) = 1\{x \in G_{\hat{f}}\} - 1\{x \notin G_{\hat{f}}\}$. Then equation (60)

becomes

$$\begin{aligned} R^\omega(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_G} R^\omega(f) &= R^\omega(\hat{f}^\dagger) - \inf_{f \in \tilde{\mathcal{F}}_G} R^\omega(f) \\ &= \frac{1}{2} \left(R_{\phi_h}^\omega(\hat{f}^\dagger) - \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}^\omega(f) \right). \end{aligned}$$

It follows that

$$\begin{aligned} R_{\phi_h}^\omega(\hat{f}^\dagger) - \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}^\omega(f) &= E_P[\omega Y \hat{f}^\dagger(X)] - \inf_{f \in \tilde{\mathcal{F}}_G} E_P[\omega Y f(X)] \\ &\leq M \left| E_P[Y \hat{f}^\dagger(X)] - \inf_{f \in \tilde{\mathcal{F}}_G} E_P[Y f(X)] \right| \\ &= M \left(R_{\phi_h}(\hat{f}^\dagger) - \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}(f) \right), \end{aligned}$$

where the third line follows because $\tilde{\mathcal{F}}_G$ is a classification-preserving reduction of \mathcal{F}_G and, accordingly, $E_P[Y \hat{f}^\dagger(X)] \geq \inf_{f \in \tilde{\mathcal{F}}_G} E_P[Y f(X)]$ holds. Thus we have

$$R^\omega(\hat{f}) - \inf_{f \in \tilde{\mathcal{F}}_G} R^\omega(f) \leq M \left(R_{\phi_h}(\hat{f}^\dagger) - \inf_{f \in \tilde{\mathcal{F}}_G} R_{\phi_h}(f) \right).$$

Then the result follows by applying the same argument in the proof of Theorem 5.1 to the above equation. \square

The following are extensions of Lemmas C.3, C.4, and C.2.

Lemma E.3. *Let $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_k} \hat{R}_{\phi_h}^\omega(f)$, and $\hat{\theta} \equiv \left\{ \hat{\theta}_{j_1 \dots j_{d_x}} \right\}_{j_1=1, \dots, k_1; \dots; j_{d_x}=1, \dots, k_{d_x}}$ be the vector of the coefficients characterizing \hat{f}_B . Let r_1^+ and r_1^- be the smallest non-negative value and the largest negative value in $\hat{\theta}$, respectively.*

(i) If all non-negative elements in $\hat{\theta}$ take the same value r_1^+ , let r_2^+ be 1; otherwise, let r_2^+ be the second smallest non-negative value in $\hat{\theta}$. Propose a $(k_1 + 1) \times \dots \times (k_{d_x} + 1)$ -dimensional vector $\tilde{\theta} \equiv \left\{ \tilde{\theta}_{j_1 \dots j_{d_x}} \right\}_{j_1=1, \dots, k_1; \dots; j_{d_x}=1, \dots, k_{d_x}}$ such that for all j_1, \dots, j_{d_x} if $\hat{\theta}_{j_1 \dots j_{d_x}} = r_1^+$, $\tilde{\theta}_{j_1 \dots j_{d_x}} = r_2^+$; otherwise, $\tilde{\theta}_{j_1 \dots j_{d_x}} = \hat{\theta}_{j_1 \dots j_{d_x}}$. Then, a classifier

$$\tilde{f}_B(x) \equiv \sum_{j_1=1}^{k_1} \dots \sum_{j_{d_x}=1}^{k_{d_x}} \tilde{\theta}_{j_1 \dots j_{d_x}} (b_{k_1 j_1}(x_1) \times \dots \times b_{k_1 j_1}(x_{d_x}))$$

minimizes \hat{R}_{ϕ_h} over \mathcal{B}_k .

(ii) Similarly, if all negative elements in $\hat{\theta}$ take the same value r_1^- , let r_2^- be -1 ; otherwise, let r_2^- be the second largest negative value in $\hat{\theta}$. Propose a $(k_1 + 1) \times \dots \times (k_{d_x} + 1)$ -vector

$\check{\theta} \equiv \{\check{\theta}_{j_1 \dots j_{d_x}}\}_{j_1=1, \dots, k_1; \dots; j_{d_x}=1, \dots, k_{d_x}}$ such that for all j_1, \dots, j_{d_x} if $\check{\theta}_{j_1 \dots j_{d_x}} = r_1^-$, $\check{\theta}_{j_1 \dots j_{d_x}} = r_2^-$; otherwise, $\check{\theta}_{j_1 \dots j_{d_x}} = \hat{\theta}_{j_1 \dots j_{d_x}}$. Then, a classifier

$$\check{f}_B(x) \equiv \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \check{\theta}_{j_1 \dots j_{d_x}} (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x}))$$

minimizes $\hat{R}_{\phi_h}^\omega$ over \mathcal{B}_k .

(iii) A classifier

$$\hat{f}_B^\dagger(x) \equiv \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \text{sign}(\hat{\theta}_{j_1 \dots j_{d_x}}) \cdot (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x}))$$

minimizes $\hat{R}_{\phi_h}(\cdot)$ over \mathcal{B}_k .

Proof. First, note that $\tilde{\theta}, \check{\theta} \in \tilde{\Theta}$ holds by construction. We here prove (i). The proof of (ii) follows using a similar argument. Define

$$L_n(\theta) = \sum_{i=1}^n \left\{ \omega_i Y_i \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \theta_{j_1 \dots j_{d_x}} \sum_{i=1}^n (b_{k_1 j_1}(X_{1i}) \times \cdots \times b_{k_{d_x} j_{d_x}}(X_{d_x i})) \right\}.$$

Minimization of \hat{R}_{ϕ_h} over \mathcal{B}_k is equivalent to the maximization of $L_n(\theta)$ over $\tilde{\Theta}$. Thus, $\hat{\theta}$ maximizes $L_n(\theta)$ over $\tilde{\Theta}$.

We prove the result by contradiction. Suppose $\tilde{\theta} \notin \arg \max_{\theta \in \tilde{\Theta}} L_n(\theta)$. Let

$$J_1 \equiv \left\{ (j_1, \dots, j_{d_x}) : \hat{\theta}_{j_1 \dots j_{d_x}} = r_1^+ \right\}.$$

Then,

$$\begin{aligned} L_n(\tilde{\theta}) - L_n(\hat{\theta}) &= \sum_{(j_1, \dots, j_{d_x}) \in J_1} \left\{ (r_2^+ - r_1^+) \sum_{i=1}^n \omega_i Y_i (b_{k_1 j_1}(X_{1i}) \times \cdots \times b_{k_{d_x} j_{d_x}}(X_{d_x i})) \right\} \\ &< 0 \end{aligned}$$

holds. Since $r_2^+ - r_1^+ \geq 0$, the above equation implies that there exists some (j_1, \dots, j_{d_x}) in J_1 such that $\sum_{i=1}^n \omega_i Y_i (b_{k_1 j_1}(X_{1i}) \times \cdots \times b_{k_{d_x} j_{d_x}}(X_{d_x i})) < 0$. For such (j_1, \dots, j_{d_x}) , setting $\hat{\theta}_{j_1 \dots j_{d_x}}$ to r_1^- increases the value of $L_n(\hat{\theta})$ without violating the constraints in $\tilde{\Theta}$.

But this contradicts the requirement that $\hat{\theta}_{j_1 \dots j_{d_x}}$ is non-negative. Therefore, $\tilde{\theta}$ maximizes $L_n(\theta)$ over $\tilde{\Theta}$, or equivalently \check{f}_B minimizes \hat{R}_{ϕ_h} over \mathcal{B}_k .

Result (iii) follows by applying results (i) and (ii) repeatedly to \hat{f}_B . \square

Lemma E.4. Fix $G \in \mathcal{G}$ and k_1, \dots, k_{d_x} . Define a classifier

$$f_G(x) = \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \theta_{j_1 \dots j_{d_x}} (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x})),$$

such that, for all j_1, \dots, j_{d_x} , $\theta_{j_1 \dots j_{d_x}} = 1$ if $(j_1/k_1, \dots, j_{d_x}/k_{d_x}) \in G$, and $\theta_{j_1 \dots j_{d_x}} = -1$ otherwise. Then the following holds:

$$|R_{\phi_h}^\omega(f_G) - R_{\phi_h}^\omega(1\{\cdot \in G\} - 1\{\cdot \notin G\})| \leq 2MA \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{4M}{\sqrt{k_j}}.$$

Proof. Define

$$J_{\mathbf{k}} \equiv \{(j_1, \dots, j_{d_x}) : (j_1/k_1, \dots, j_{d_x}/k_{d_x}) \in G\},$$

which is a set of grid points on G , and

$$L(x) \equiv -\omega_+(x)\eta(x) + \omega_-(x)(1 - \eta(x)).$$

It follows that

$$\begin{aligned} & R_{\phi_h}^\omega(f_G) - R_{\phi_h}^\omega(1\{\cdot \in G\} - 1\{\cdot \notin G\}) \\ &= \int_{[0,1]^{d_x}} L(x) (1\{x \in G\} - 1\{x \notin G\} - B_{\mathbf{k}}(\theta, x)) dP_X(x) \\ &= \int_{[0,1]^{d_x}} L(x) 1\{x \in G\} dP_X(x) - \int_{[0,1]^{d_x}} L(x) 1\{x \notin G\} dP_X(x) \\ &\quad - \underbrace{\int_{[0,1]^{d_x}} L(x) B_{\mathbf{k}}(\theta, x) dP_X(x)}_{(I)}. \end{aligned} \tag{62}$$

(I) can be written as

$$\begin{aligned} (I) &= \int_{[0,1]^{d_x}} L(x) \sum_{(j_1, \dots, j_{d_x}) \in J_{\mathbf{k}}} \left(\prod_{v=1}^{d_x} b_{k_v j_v}(x_v) \right) dP_X(x) \\ &\quad - \int_{[0,1]^{d_x}} L(x) \sum_{(j_1, \dots, j_{d_x}) \notin J_{\mathbf{k}}} \left(\prod_{v=1}^{d_x} b_{k_v j_v}(x_v) \right) dP_X(x). \end{aligned}$$

Thus,

$$\begin{aligned}
(62) &= \int_{[0,1]^{d_x}} L(x) \underbrace{\left(1 \{x \in G\} - \sum_{(j_1, \dots, j_{d_x}) \in J_{\mathbf{k}}} b_{k_v j_v}(x_v) \right)}_{(II)} dP_X(x) \\
&+ \int_{[0,1]^{d_x}} L(x) \underbrace{\left(\sum_{(j_1, \dots, j_{d_x}) \notin J_{\mathbf{k}}} b_{k_v j_v}(x_v) - 1 \{x \in G^c\} \right)}_{(III)} dP_X(x).
\end{aligned}$$

Let $Bin(k_j, x_j)$, $j = 1, \dots, d_x$, be independent binomial variables with parameters (k_j, x_j) . Then (II) and (III) are equivalent to

$$\begin{aligned}
&\Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) 1 \{x \in G\} \\
&- \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) 1 \{x \in G^c\}.
\end{aligned}$$

Hence,

$$\begin{aligned}
(62) &= 2 \int_G L(x) \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) dP_X(x) \\
&- 2 \int_{G^c} L(x) \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) dP_X(x),
\end{aligned}$$

and therefore

$$\begin{aligned}
&|R_{\phi_h}^\omega(f_G) - R_{\phi_h}^\omega(1 \{\cdot \in G\} - 1 \{\cdot \notin G\})| \\
&\leq 2M \underbrace{\int_G \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}^c) dP_X(x)}_{(IV)} \\
&+ 2M \underbrace{\int_{G^c} \Pr((Bin(k_1, x_1), \dots, Bin(k_{d_x}, x_{d_x})) \in J_{\mathbf{k}}) dP_X(x)}_{(V)},
\end{aligned}$$

because $|L(x)| < M$ for all $x \in \mathcal{X}$. The proof of Lemma C.4 shows that

$$\begin{aligned}
(IV) &\leq \frac{A}{2} \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \frac{1}{\sqrt{k_v}}, \\
(V) &\leq \frac{A}{2} \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \frac{1}{\sqrt{k_v}}.
\end{aligned}$$

Therefore,

$$|R_{\phi_h}^\omega(f_G) - R_{\phi_h}^\omega(1\{\cdot \in G\} - 1\{\cdot \notin G\})| \leq 2MA \left(\sum_{v=1}^{d_x} \sqrt{\frac{\log k_v}{k_v}} \right) + \sum_{v=1}^{d_x} \frac{4M}{\sqrt{k_v}}.$$

□

Lemma E.5. *Let $k_j \geq 1$, for $j = 1, \dots, d_x$, be fixed. Suppose that the density of P_X is bounded from above by $A > 0$. Suppose further that $\mathbf{k} = (k_1, \dots, k_{d_x})$ satisfies $\sqrt{d_x \log k_j} / (2\sqrt{k_j}) \leq \epsilon$ for all $j = 1, \dots, d_x$ and some $\epsilon > 0$.*

(i) *The following holds for the approximation error to the best classifier:*

$$\inf_{f \in \mathcal{B}_{\mathbf{k}}} R_{\phi_h}^\omega(f) - \inf_{f \in \mathcal{F}_M} R_{\phi_h}^\omega(f) \leq 2AM \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{4M}{\sqrt{k_j}}.$$

(ii) *For $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_{\mathbf{k}}} \hat{R}_{\phi_h}^\omega(f)$ such that its coefficients take values in $\{-1, 1\}$, the following holds for the approximation error to the step function:*

$$R_{\phi_h}^\omega(1\{\cdot \in G_{\hat{f}_B}\} - 1\{\cdot \notin G_{\hat{f}_B}\}) - R_{\phi_h}^\omega(\hat{f}_B) \leq 2AM \sum_{j=1}^{d_x} \sqrt{\frac{\log k_j}{k_j}} + \sum_{j=1}^{d_x} \frac{4M}{\sqrt{k_j}}.$$

Proof. We first prove (i). Let G^* minimize $\mathcal{R}^\omega(\cdot)$ over \mathcal{G}_M . From Theorem 7.5, a classifier $\tilde{f}^*(x) \equiv 1\{x \in G^*\} - 1\{x \in (G^*)^c\}$ minimizes the hinge risk $R_{\phi_h}^\omega(\cdot)$ over \mathcal{F}_M . Define a vector $\theta^* = \{\theta_{j_1 \dots j_d}^*\}_{j_1=0, \dots, k_1; \dots; j_d=0, \dots, k_d}$ such that for each j_1, \dots, j_d ,

$$\theta_{j_1 \dots j_d}^* = \begin{cases} 1 & \text{if } (j_1/k_1, \dots, j_d/k_d) \in G^* \\ -1 & \text{otherwise.} \end{cases}$$

Note that θ^* is contained in $\tilde{\Theta}$. Thus, it follows that

$$\inf_{f \in \mathcal{B}_{\mathbf{k}}} R_{\phi_h}^\omega(f) - \inf_{f \in \mathcal{F}_M} R_{\phi_h}^\omega(f) \leq R_{\phi_h}^\omega(B_{\mathbf{k}}(\theta^*, \cdot)) - R_{\phi_h}^\omega(\tilde{f}^*).$$

Then, applying Lemma E.4 to $R_{\phi_h}^\omega(B_{\mathbf{k}}(\theta^*, \cdot)) - R_{\phi_h}^\omega(\tilde{f}^*)$ establishes result (i).

The inequality in Lemma E.5 (ii) follows immediately from Lemma E.4. Applying Lemma E.3 (iii) to any $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_{\mathbf{k}}} \hat{R}_{\phi_h}^\omega(f)$ shows that a classifier

$$\hat{f}_B^\dagger(x) = \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d_x}=1}^{k_{d_x}} \text{sign}(\hat{\theta}_{j_1 \dots j_{d_x}}) (b_{k_1 j_1}(x_1) \times \cdots \times b_{k_{d_x} j_{d_x}}(x_{d_x}))$$

minimizes $\hat{R}_{\phi_h}^\omega(\cdot)$ over \mathcal{B}_k , which proves the existence of $\hat{f}_B \in \arg \inf_{f \in \mathcal{B}_k} \hat{R}_{\phi_h}^\omega(f)$ such that its coefficients take values in $\{-1, 1\}$. \square

References

- AGARWAL, A., A. BEYGEZIMER, M. DUDÍK, J. LANGFORD, AND H. WALLACH (2018): “A reductions approach to fair classification,” in *International Conference on Machine Learning (ICML)*, PMLR, 60–69.
- ALEXANDER, K. S. (1984): “Probability inequalities for empirical processes and a law of the iterated logarithm,” *Annals of Probability*, 12, 1041–1067.
- ATHEY, S. AND S. WAGER (2021): “Policy learning with observational data,” *Econometrica*, 89, 133–161.
- BABII, A., X. CHEN, E. GHYSELS, AND R. KUMAR (2020): “Binary choice with asymmetric loss in a data-rich environment: Theory and an application to racial justice,” *arXiv preprint arXiv:2010.08463*.
- BARTLETT, P. L., M. I. JORDAN, AND J. D. MCAULIFFE (2006): “Convexity, classification, and risk bounds,” *Journal of the American Statistical Association*, 101, 138–156.
- BEYGEZIMER, A. AND J. LANGFORD (2009): “The offset tree for learning with partial labels,” in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, 129–137.
- BHATTACHARYA, D. AND P. DUPAS (2012): “Inferring welfare maximizing treatment assignment under budget constraints,” *Journal of Econometrics*, 167, 168–196.
- BOUCHERON, S., G. LUGOSI, AND P. MASSART (2013): *Concentration Inequalities: A Nonasymptotic Theory of Independence*, Oxford, UK: Oxford University Press.
- BREIMAN, L., J. FRIEDMAN, C. STONE, AND R. OLSHEN (1984): *Classification and Regression Trees*, The Wadsworth and Brooks-Cole statistics-probability series, Taylor & Francis.
- CALDERS, T. AND S. VERWER (2010): “Three naive Bayes approaches for discrimination-free classification,” *Data Mining and Knowledge Discovery*, 21, 277–292.
- CANO, J. R., P. A. GUTIÉRREZ, B. KRAWCZYK, M. WOŹNIAK, AND S. GARCÍA (2019): “Monotonic classification: An overview on algorithms, performance measures and data sets,” *Neurocomputing*, 341, 168–182.

- CHAMBERLAIN, G. (2011): “Bayesian aspects of treatment choice,” in *The Oxford Handbook of Bayesian Econometrics*, ed. by J. Geweke, G. Koop, and H. van Dijk, Oxford University Press, 11–39.
- CHEN, C. C. AND S. T. LI (2014): “Credit rating with a monotonicity-constrained support vector machine model,” *Expert Systems with Applications*, 41, 7235–7247.
- CHOULDECHOVA, A. AND A. ROTH (2018): “The frontiers of fairness in machine learning,” *arXiv preprint arXiv:1810.08810*.
- CORTES, C. AND V. VAPNIK (1995): “Support-vector networks,” *Machine Learning*, 20, 273–297.
- DEHEJIA (2005): “Program evaluation as a decision problem,” *Journal of Econometrics*, 125, 141–173.
- DEVROYE, L., L. GYÖRFI, AND G. LUGOSI (1996): *A Probabilistic Theory of Pattern Recognition*, Springer.
- DONINI, M., L. ONETO, S. BEN-DAVID, J. SHAWE-TAYLOR, AND M. PONTIL (2018): “Empirical risk minimization under fairness constraints,” in *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS 2018)*, 2796–2806.
- DUDLEY, R. M. (1999): *Uniform Central Limit Theorems*, Cambridge University Press.
- DWORK, C., M. HARDT, T. PITASSI, O. REINGOLD, AND R. ZEMEL (2012): “Fairness through awareness,” in *Proceedings of the 2012 Innovations in Theoretical Computer Science Conference (ITCS 2012)*, 214–226.
- FELDMAN, M., S. A. FRIEDLER, J. MOELLER, C. SCHEIDEGGER, AND S. VENKATASUBRAMANIAN (2015): “Certifying and removing disparate impact,” in *Proceedings of the 2015 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 259–268.
- FREUND, Y. AND R. SCHAPIRE (1997): “A decision-theoretic generalization of on-line learning and an application to boosting,” *Journal of Computer System Sciences*, 55, 119–139.
- GAO, F. AND J. A. WELLNER (2007): “Entropy estimate for high-dimensional monotonic functions,” *Journal of Multivariate Analysis*, 98, 1751–1764.
- GINÉ, E. AND R. NICKL (2016): *Mathematical Foundations of Infinite-Dimensional Statistical Models*, New York: Cambridge University Press.

- HIRANO, K. AND J. R. PORTER (2009): “Asymptotics for statistical treatment rules,” *Econometrica*, 77, 1683–1701.
- HOROWITZ, J. L. (1992): “A smoothed maximum score estimator for the binary response model,” *Econometrica*, 60, 505–531.
- JIANG, W. (2004): “Process consistency for adaboost,” *Annals of Statistics*, 32, 13–29.
- KALLUS, N. (2021): “More efficient policy learning via optimal retargeting,” *Journal of the American Statistical Association*, 116, 646–658.
- KAMISHIMA, T., S. AKAHO, AND J. SAKUMA (2011): “Fairness-aware learning through regularization approach,” in *Proceedings of the 2011 IEEE 11th International Conference on Data Mining (ICDM 2011)*, 643–650.
- KARLAN, D., S. MULLAINATHAN, AND B. N. ROTH (2019): “Debt traps? Market vendors and moneylender debt in India and the Philippines,” *American Economic Review: Insights*, 1, 27–42.
- KASY, M. (2018): “Optimal taxation and insurance using machine learning — Sufficient statistics and beyond,” *Journal of Public Economics*, 167, 205–219.
- KITAGAWA, T. AND Y.-L. LIN (2021): “Boosting for treatment choice,” *unpublished manuscript*.
- KITAGAWA, T. AND A. TETENOV (2018): “Who should be treated? Empirical welfare maximization methods for treatment choice,” *Econometrica*, 86, 591–616.
- (2021): “Equality-minded treatment choice,” *Journal of Business & Economic Statistics*, 39, 561–574.
- KITAGAWA, T. AND G. WANG (2020): “Who should get vaccinated? Individualized allocation of vaccines over SIR network,” *cemmap working paper*.
- KOLTCHINSKII, V. (2006): “Local Rademacher complexities and oracle inequalities in risk minimization,” *Annals of Statistics*, 34, 2593–2656.
- LIELI, R. P. AND H. WHITE (2010): “The construction of empirical credit scoring rules based on maximization principles,” *Journal of Econometrics*, 157, 110–119.
- LUGOSI, G. (2002): “Pattern classification and learning theory,” in *Principles of Non-parametric Learning*, ed. by L. Györfi, Vienna: Springer, 1–56.
- LUGOSI, G. AND N. VAYATIS (2004): “On the Bayes-risk consistency of regularized boosting methods,” *Annals of Statistics*, 32, 30–55.

- MAMMEN, E. AND A. B. TSYBAKOV (1999): “Smooth discrimination analysis,” *Annals of Statistics*, 27, 1808–1829.
- MANNOR, S., R. MEIR, AND T. ZHANG (2003): “Greedy algorithms for classification – consistency, convergence rates, and adaptivity,” *Journal of Machine Learning Research*, 4, 713–742.
- MANSKI, C. F. (1975): “Maximum score estimation of the stochastic utility model of choice,” *Journal of Econometrics*, 3, 205–228.
- (2004): “Statistical treatment rules for heterogeneous populations,” *Econometrica*, 72, 1221–1246.
- MANSKI, C. F. AND T. THOMPSON (1989): “Estimation of best predictors of binary response,” *Journal of Econometrics*, 40, 97–123.
- MBAKOP, E. AND M. TABORD-MEEHAN (2021): “Model selection for treatment choice: Penalized welfare maximization,” *Econometrica*, 89, 825–848.
- NGUYEN, X., M. J. WAINWRIGHT, AND M. I. JORDAN (2009): “On surrogate loss functions and f-divergences,” *Annals of Statistics*, 37, 876–904.
- QIAN, M. AND S. A. MURPHY (2011): “Performance guarantees for individualized treatment rules,” *Annals of Statistics*, 39, 1180–1210.
- RAMBACHAN, A., J. KLEINBERG, J. LUDWIG, AND S. MULLAINATHAN (2020): “An economic approach to regulating algorithms,” *NBER working paper*.
- SAKAGUCHI, S. (2021): “Estimation of optimal dynamic treatment assignment rules under policy constraint,” *arXiv preprint arXiv:2106.05031*.
- SCOTT, C. (2012): “Calibrated asymmetric surrogate losses,” *Electronic Journal of Statistics*, 6, 958–992.
- STEINWART, I. (2005): “Consistency of support vector machines and other regularized kernel classifiers,” *IEEE Transactions on Information Theory*, 51, 713–742.
- (2007): “How to compare different loss functions and their risks,” *Constructive Approximation*, 26, 225–287.
- STOYE, J. (2009): “Minimax regret treatment choice with finite samples,” *Journal of Econometrics*, 151, 70–81.
- (2012): “Minimax regret treatment choice with covariates or with limited validity of experiments,” *Journal of Econometrics*, 166, 138–156.

- SWAMINATHAN, A. AND T. JOACHIMS (2015): “Counterfactual risk minimization: Learning from logged bandit feedback,” *Journal of Machine Learning Research*, 16, 1731–1755.
- TETENOV, A. (2012): “Statistical treatment choice based on asymmetric minimax regret criteria,” *Journal of Econometrics*, 166, 157–165.
- TSYBAKOV, A. B. (2004): “Optimal aggregation of classifiers in statistical learning,” *Annals of Statistics*, 32, 135–166.
- VAPNIK, V. N. (1998): *Statistical Learning Theory*, John Wiley & Sons.
- VIVIANO, D. (2019): “Policy targeting under network interference,” *arXiv preprint arXiv:1906.10258*.
- WANG, J. AND S. K. GHOSH (2012): “Shape restricted nonparametric regression based on multivariate Bernstein polynomials,” *North Carolina State University Department of Statistics technical report*.
- ZADROZNY, B. (2003): “Policy mining: Learning decision policies from fixed sets of data,” *Ph.D Thesis, University of California, San Diego*.
- ZENG, J., B. USTUN, AND C. RUDIN (2017): “Interpretable classification models for recidivism prediction,” *Journal of the Royal Statistical Society, Series A: Statistics in Society*, 180, 689–722.
- ZHANG, T. (2004): “Statistical behavior and consistency of classification methods based on convex risk minimization,” *Annals of Statistics*, 32, 56–85.
- ZHANG, Y., E. B. LABER, M. DAVIDIAN, AND A. A. TSIATIS (2018): “Interpretable dynamic treatment regimes,” *Journal of the American Statistical Association*, 113, 1541–1549.
- ZHAO, Y., D. ZENG, E. B. LABER, AND M. R. KOSOROK (2015): “New statistical learning methods for estimating optimal dynamic treatment regimes,” *Journal of the American Statistical Association*, 110, 583–598.
- ZHAO, Y., D. ZENG, A. J. RUSH, AND M. R. KOSOROK (2012): “Estimating individualized treatment rules using outcome weighted learning,” *Journal of the American Statistical Association*, 107, 1106–1118.