# Optimal Decision Rules Under Partial Identification

Kohei Yata*

August 31, 2023

**Abstract**

I consider a class of statistical decision problems in which the policy maker must decide between two alternative policies to maximize social welfare based on a finite sample. The central assumption is that the underlying, possibly infinite-dimensional parameter, lies in a known convex set, potentially leading to partial identification of the welfare effect. An example of such restrictions is the smoothness of counterfactual outcome functions. As the main theoretical result, I derive a finite-sample, exact minimax regret decision rule within the class of all decision rules under normal errors with known variance. When the error distribution is unknown, I obtain a feasible decision rule that is asymptotically minimax regret. I apply my results to the problem of whether to change a policy eligibility cutoff in a regression discontinuity setup, and illustrate them in an empirical application to a school construction program in Burkina Faso.

*Keywords:* Statistical decision theory, finite-sample minimax regret, partial identification, nonparametric regression models, regression discontinuity.

## 1 Introduction

A fundamental goal of empirical research in economics is to inform policy decisions. Evaluation of counterfactual policies often requires extrapolating from observables to unobservables. Without strong model restrictions such as functional form assumptions or exogeneity of an intervention, the performance of each counterfactual

*Department of Economics, University of Wisconsin–Madison. email: yata@wisc.edu. I am grateful to my thesis advisors Yuichi Kitamura, Timothy Armstrong, and Yusuke Narita for their invaluable advice and encouragement. I also thank Donald Andrews, Jack Porter, and seminar and conference participants at Bravo/JEA/SNSF Workshop, Bristol, CMStatistics 2022, Cornell, Happy Hour Seminar, IAAE 2023, Notre Dame, Penn State, TSE, UC Irvine, UCL, UW–Madison, Western, and Yale for helpful comments and suggestions.

policy may be only partially determined by observed data. In such situations, policy decision-making is challenging since we have no clear understanding of which policy is the best.

For example, a regression discontinuity (RD) design only credibly estimates the impact of treatment on the individuals at the eligibility cutoff. Therefore, without restrictive assumptions such as constant treatment effects, whether or not to offer the treatment to those away from the cutoff is ambiguous. Even randomized controlled trials may provide only partial knowledge of the impact of a new intervention, as can happen if participants do not perfectly comply with their assigned treatment or if the experimental sample is an unrepresentative subset of the target population.

This paper develops an optimal way of using data to make policy decisions in settings where social welfare under each counterfactual policy may be only partially identified. Specifically, I solve a class of statistical decision problems. The setup is as follows. The policy maker must decide between two alternative policies, policy 1 and policy 0, to maximize social welfare (e.g., the population mean of an outcome). The difference in welfare between the two policies is given by $L(\theta) \in \mathbb{R}$. $L(\cdot)$ is a known linear function of an unknown, possibly infinite-dimensional parameter $\theta$, such as nonparametric counterfactual outcome functions. By construction, it is optimal to choose policy 1 if $L(\theta) \geq 0$ and to choose policy 0 if $L(\theta) < 0$. The policy maker makes a decision after observing a finite sample $\boldsymbol{Y} = (Y_1, ..., Y_n)' \in \mathbb{R}^n$ whose expected value is given by $\boldsymbol{m}(\theta) \in \mathbb{R}^n$, where $\boldsymbol{m}(\cdot)$ is a known linear function of $\theta$. The main assumption is that $\theta$ belongs to a known convex set $\Theta$, potentially leading to partial identification of the welfare difference $L(\theta)$. An example of such restrictions is the smoothness of counterfactual outcome functions.

This setup is a general class of binary decision problems allowing for both point and partial identification. This class contains as special cases existing problems analyzed by, for example, Stoye (2012) and Ishihara and Kitagawa (2021), as well as novel problems under nonparametric regression models (including RD models) with a broad class of restrictions on the regression function. Section 2.3 presents several examples.

The main contribution of this paper is to derive a finite-sample, exact decision rule (i.e., a function that maps the sample $\boldsymbol{Y}$ to a probability of choosing policy 1) that is optimal under the minimax regret criterion, a standard criterion used in the literature on statistical treatment choice (e.g., Manski, 2004; Stoye, 2009; Kitagawa and Tetenov, 2018). That is, I obtain a decision rule $\delta^* : \mathbb{R}^n \to [0,1]$ that solves

$$\inf_{\delta:\mathbb{R}^n\to[0,1]} \sup_{\theta\in\Theta} R(\delta,\theta), \quad R(\delta,\theta) = \begin{cases} L(\theta)(1-\mathbb{E}_\theta[\delta(\boldsymbol{Y})]) & \text{if } L(\theta) \geq 0, \\ -L(\theta)\mathbb{E}_\theta[\delta(\boldsymbol{Y})] & \text{if } L(\theta) < 0. \end{cases}$$

Here, $R(\delta,\theta)$ is the expected welfare loss (i.e., regret) of decision rule $\delta$ under $\theta$. To derive the finite-sample optimality result, I assume that the sample $\boldsymbol{Y}$ is normally distributed with known variance and that the parameter space $\Theta$ is convex and centrosymmetric (i.e., $\theta \in \Theta$ implies $-\theta \in \Theta$), as well as mild regularity conditions. When the distribution of the error $\boldsymbol{Y} - \boldsymbol{m}(\theta)$ is unknown, I obtain a feasible decision rule that is asymptotically minimax regret under suitable regularity conditions. Appendix C.1 provides the asymptotic analysis.

Importantly, when I derive a finite-sample minimax regret rule, I do not impose any restrictions on the class of decision rules, thus allowing for, for example, nonrandomized rules based on nonlinear functions of the sample and randomized rules. This is in contrast to the approach by Ishihara and Kitagawa (2021), who focus on the class of nonrandomized rules that make a decision based on the sign of a linear function of the sample to characterize the maximum regret and reduce the dimension of the outer minimization problem. The main tool that I use to solve the minimax regret problem is what is called the *modulus of continuity*, defined in Section 3. It was originally introduced by Donoho (1994) to characterize minimax affine nonrandomized estimation and inference procedures, such as a minimax affine mean squared error (MSE) estimator, for a linear functional of a nonparametric regression function. I show that the minimax *regret* problem over the class of *all* decision rules can be simplified into an optimization problem with respect to the modulus of continuity. The optimization problem is analytically and computationally tractable.

The resulting decision rule is simple and thus easy to compute. It makes a decision based on a linear function of $\boldsymbol{Y}$. The minimax regret rule may be randomized or nonrandomized, depending on the restrictions imposed on the parameter space. More specifically, the rule is nonrandomized if the identified set of the welfare difference $L(\theta)$ is, in a certain formal sense, small relative to the variance of the sample $\boldsymbol{Y}$, including the case where $L(\theta)$ is point identified. Otherwise, it is a randomized rule, assigning a positive probability both to policies 1 and 0.

When the minimax regret rule is nonrandomized, it can be viewed as a rule that plugs a particular linear estimator of the welfare difference $L(\theta)$ into the optimal decision $\mathbf{1}\{L(\theta) \geq 0\}$. I compare this linear estimator with a linear minimax MSE estimator of $L(\theta)$, which minimizes the maximum of the MSE over the parameter space within the class of all linear estimators. The two estimators are shown to be generally different, which suggests that the plug-in rule based on the linear minimax MSE estimator is not optimal under the minimax regret criterion. More precisely, the linear estimator used by the minimax regret rule places more importance on the bias than on the variance compared to the linear minimax MSE estimator.

The main application considered by this paper is the problem of eligibility cutoff choice in an RD setup. In many policy domains, the eligibility for treatment is determined based on an individual's observable characteristics. One crucial policy question is whether we should change the eligibility criterion to achieve better outcomes (Dong and Lewbel, 2015). Specifically, I consider an RD setup and study the problem of whether or not to change the eligibility cutoff from a current value to a new value.[1] For an illustration of the results, I focus on the case where the welfare is the sample average outcome and the conditional mean counterfactual outcome function belongs to the class of Lipschitz functions with a known Lipschitz constant. Under the Lipschitz constraint, the welfare effect of the cutoff change is partially identified. Applying my general results, I obtain a minimax regret rule for this problem. It makes a decision based on the difference between a weighted average of outcomes for the treated units and that for the untreated units, with the weight vector determined by the choice of the new cutoff and Lipschitz constant. The rule can be easily computed by solving finite-dimensional convex optimization problems.

Finally, my approach is illustrated through an empirical application to the Burkinabé Response to Improve Girls' Chances to Succeed (BRIGHT) program, a school construction program in Burkina Faso (Kazianga, Levy, Linden and Sloan, 2013). Aiming to improve educational outcomes in rural villages, the program constructed

[1]This problem is related to the growing literature on extrapolation away from the cutoff in RD designs, including Rokkanen (2015), Angrist and Rokkanen (2015), Dong and Lewbel (2015), Bertanha and Imbens (2020), Bertanha (2020), Bennett (2020), and Cattaneo, Keele, Titiunik and Vazquez-Bare (2021). Unlike these papers, I explicitly consider the decision problem of whether or not to change the cutoff and derive an optimal decision rule.

primary schools in 132 villages from 2005 to 2008. To allocate schools, the Ministry of Education first computed a score summarizing village characteristics for each of the nominated 293 villages and then selected the highest-ranking villages to receive a school. This situation fits into an RD setup.

I ask whether we should expand this program or not. The more specific question considered in this analysis is whether or not to construct schools in the top 20% of previously ineligible villages. The analysis uses the enrollment rate as the welfare measure and assumes that the conditional mean counterfactual outcome function belongs to the class of Lipschitz functions with a known Lipschitz constant. For a plausible range of the Lipschitz constant, the minimax regret rule implies that building schools in the top 20% of previously ineligible villages is not cost-effective.

**Related Literature.** This paper adds new results to the existing body of research on finite-sample optimal decision rules for treatment choice under point identification (e.g., Stoye, 2009, 2012; Tetenov, 2012; Kitagawa, Lee and Qiu, 2023) and partial identification (e.g., Manski, 2007; Stoye, 2012; Ishihara and Kitagawa, 2021). The most closely related to this paper are Stoye (2012) and Ishihara and Kitagawa (2021). Proposition 7(iii) of Stoye (2012) provides a special case of my result in a setting where the experiment has imperfect internal or external validity. My result generalizes Stoye (2012)'s in three ways: (1) the sample can be multidimensional; (2) the parameter can be three or higher dimensional, even infinite dimensional; and (3) flexible forms of the parameter space are allowed. Ishihara and Kitagawa (2021) consider a special case of my setup, where the policy maker must decide whether or not to introduce a new policy based on results from multiple external studies. While they derive a minimax regret rule within the class of nonrandomized decision rules based on the sign of a linear function of the sample, I derive a minimax regret rule within the class of all decision rules.

More broadly, this paper contributes to the literature on statistical treatment choice and policy learning under partial identification, which has been recently growing in econometrics and statistics (e.g., Manski, 2000, 2007, 2009, 2010, 2011a,b, 2021; Chamberlain, 2011; Stoye, 2012; Kasy, 2018; Russell, 2020; Mo, Qi and Liu, 2021; Kallus and Zhou, 2021; Ishihara and Kitagawa, 2021; D'Adamo, 2022; Ben-Michael, Greiner, Imai and Jiang, 2022; Zhang, Ben-Michael and Imai, 2022; Christensen, Moon and Schorfheide, 2023). An approach to decisions under partial identification consists of the following two steps. The first step solves the population minimax regret problem with the perfect knowledge of the identified set. This leads to a point-identified policy that minimizes the maximum of regret over the identified set, which is generally different from the true optimal policy that is only partially identified. The second step uses data to estimate the point-identified policy. The papers taking this approach typically derive finite-sample risk bounds or asymptotic risk optimality as a statistical property, where the risk is defined as the expected value of the maximum of regret over the identified set.[2] See, for example, D'Adamo (2022) and Christensen *et al.* (2023). In contrast, this paper primarily focuses on finite-sample exact optimality and defines the risk as the expected value of regret, which corresponds to the standard notion of risk in Wald's statistical decision theory when the loss function is given by regret (Wald, 1950).

Extensive literature exists on the problem of learning treatment allocation policies that map an individual's covariates to a treatment. See Manski (2004), Dehejia (2005), Hirano and Porter (2009), Stoye (2009, 2012), Qian and Murphy (2011), Bhattacharya and Dupas (2012), Kitagawa and Tetenov (2018, 2021), Athey and Wager (2021), and Mbakop and Tabord-Meehan (2021), among others. When the covariates are discrete and no constraints are imposed on the class of treatment allocation policies, we can construct a decision rule by applying this paper's minimax regret rule to each of the subsamples with different covariate values separately. While this rule does not pool information across covariate values at all, it can be minimax regret if there are no cross-covariate parameter restrictions (Stoye, 2009).

The Gaussian model used in this paper has been studied for the problem of minimax estimation and inference in nonparametric regression models. Donoho (1994) uses the modulus of continuity to characterize minimax affine nonrandomized estimation and inference procedures. I show that the modulus of continuity can be used for deriving a minimax regret decision rule within the class of all decision rules. While the derivation of my result and that of Donoho (1994)'s consist of similar steps, the proof of each step is significantly different. Recently, Donoho (1994)'s framework

[2] An exception is Manski (2007), who considers a setup where an unbiased estimator for the population minimax regret rule exists and shows that any unbiased estimator is finite-sample minimax regret.

has been applied to estimation and inference on treatment effects in a variety of settings, including RD and difference-in-differences designs (Armstrong and Kolesár, 2018; Imbens and Wager, 2019; Kwon and Kwon, 2020; Armstrong and Kolesár, 2021; de Chaisemartin, 2021; Rambachan and Roth, 2023). My framework can be applied to the problem of treatment choice under these settings.

# 2 Setup, Optimality Criterion, and Examples

## 2.1 Setup

**Data-generating Model.** Suppose that the policy maker observes a sample $\boldsymbol{Y} = (Y_1, ..., Y_n)' \in \mathbb{R}^n$ of the form

$$\boldsymbol{Y} \sim \mathcal{N}(\boldsymbol{m}(\theta), \boldsymbol{\Sigma}), \tag{1}$$

where $\theta$ is an unknown parameter that lies in a known subset $\Theta$ of a vector space $\mathbb{V}$, $\boldsymbol{m} : \mathbb{V} \to \mathbb{R}^n$ is a known linear function, and $\boldsymbol{\Sigma}$ is a known, positive-definite $n \times n$ matrix.[3] I allow $\theta$ to be an infinite-dimensional parameter such as a function.

The linearity of $\boldsymbol{m}$ is not necessarily restrictive. If we specify $\theta$ so that it contains each of the expected values of $Y_1, ..., Y_n$ as its element, $\boldsymbol{m}$ is a function that extracts those expected values from $\theta$, which is linear in $\theta$.

This model allows the expected value of $\boldsymbol{Y}$ to depend on other observed variables such as covariates and treatment by treating them as fixed and subsuming them into $\boldsymbol{m}$ and $\boldsymbol{\Sigma}$. For example, a regression model with fixed regressors

$$Y_i = f(x_i) + u_i, \quad u_i \sim \mathcal{N}(0, \sigma^2(x_i)) \text{ independent across } i$$

is a special case where $\boldsymbol{Y} = (Y_1, ..., Y_n)'$, $\theta = f$, $\Theta$ is a class of functions, $\boldsymbol{m}(f) = (f(x_1), ..., f(x_n))'$, and $\boldsymbol{\Sigma} = \mathrm{diag}(\sigma^2(x_1), ..., \sigma^2(x_n))$.

The normality of $\boldsymbol{Y}$ and the assumption of known variance are restrictive, but are often imposed to deliver finite-sample optimality results for problems of estimation,

[3] Donoho (1994), Low (1995), and Armstrong and Kolesár (2018) investigate optimal estimation and inference of a linear functional of $\theta$ in a slightly more general version of this model that allows $\boldsymbol{Y}$ to be infinite dimensional.

inference, and treatment choice. In some cases, it is plausible to assume the normality of $\boldsymbol{Y}$. For example, suppose that unit $i$ represents a group of individuals defined by place, time, and individual characteristics among others and that $Y_1, ..., Y_n$ are group-level mean outcomes. If the number of groups is fixed at $n$, the distribution of $\boldsymbol{Y}$ approaches a normal distribution as the size of each group grows to infinity by the central limit theorem. The normal model (1) can be viewed as an asymptotic approximation if each group is large enough.[4]

In Appendix C.1, I consider an asymptotic framework where the distribution of the error $\boldsymbol{Y} - \boldsymbol{m}(\theta)$ is unknown and the sample size $n$ goes to infinity. I derive asymptotic optimality results under high-level regularity conditions.

I assume that the parameter space $\Theta$ is convex and centrosymmetric (i.e., $\theta \in \Theta$ implies $-\theta \in \Theta$) throughout the paper. Typical parameter spaces considered in empirical analyses are convex. For example, in the regression model above, classes of functions with bounded derivatives (e.g., the class of Lipschitz functions with a known Lipschitz constant) are convex. The centrosymmetry simplifies the analysis, but rules out some shape restrictions. In the regression model above, the class of convex functions and the class of concave functions fail to be centrosymmetric.[5]

**Policy Choice Problem.** Now, suppose that the policy maker is interested in choosing between two alternative policies, policy 1 and policy 0, to maximize social welfare. Examples of a binary policy decision include whether to introduce a program to a target population and whether to change a policy from the status quo to a new

[4] More generally, suppose that the policy maker observes an $n$-dimensional vector of statistics of the original data and that it is an asymptotically normal estimator of its population counterpart. For example, the mean outcome difference between the treatment and control groups in a randomized experiment is a statistic that is asymptotically normal for the population mean difference. If we regard the $n$-dimensional vector of statistics as $\boldsymbol{Y}$, the normal model (1) can again be viewed as an asymptotic approximation (Stoye, 2012; Tetenov, 2012; Ishihara and Kitagawa, 2021; Rambachan and Roth, 2023; Andrews, Kitagawa and McCloskey, 2023).

[5] On the other hand, in some cases, it is possible to impose the monotonicity of the regression function by normalizing the sample $\boldsymbol{Y}$ so that the new parameter space is centrosymmetric. Suppose, for example, that $\Theta = \{f \in \mathcal{F}_{\mathrm{Lip}}(C) : f(x) \text{ is nondecreasing in } x\}$, where $\mathcal{F}_{\mathrm{Lip}}(C) = \{f : |f(x) - f(\tilde{x})| \leq C|x - \tilde{x}| \text{ for every } x, \tilde{x} \in \mathbb{R}\}$. $\mathcal{F}_{\mathrm{Lip}}(C)$ is centrosymmetric while $\Theta$ is not. It is easy to show that $\Theta = \{\tilde{f} + f_0 : \tilde{f} \in \mathcal{F}_{\mathrm{Lip}}(C/2)\}$, where $f_0(x) = \frac{C}{2}x$ for all $x \in \mathbb{R}$. Therefore, the model $\boldsymbol{Y} \sim \mathcal{N}(\boldsymbol{m}(f), \boldsymbol{\Sigma})$, $f \in \Theta$, is equivalent to the model $\tilde{\boldsymbol{Y}} \sim \mathcal{N}(\boldsymbol{m}(\tilde{f}), \boldsymbol{\Sigma})$, $\tilde{f} \in \mathcal{F}_{\mathrm{Lip}}(C/2)$, where $\tilde{\boldsymbol{Y}} = \boldsymbol{Y} - \boldsymbol{m}(f_0) = (Y_1 - f_0(x_1), ..., Y_n - f_0(x_n))'$; the set of distributions of $\boldsymbol{Y}$ over $f \in \Theta$ is identical to the set of distributions of $\tilde{\boldsymbol{Y}} + \boldsymbol{m}(f_0)$ over $\tilde{f} \in \mathcal{F}_{\mathrm{Lip}}(C/2)$.

one. Suppose that the welfare resulting from implementing policy $a \in \{0, 1\}$ under $\theta$ is $W_a(\theta)$, where $W_a : \mathbb{V} \to \mathbb{R}$ is a known function specified by the policy maker. The welfare difference between policy 1 and policy 0 is given by

$$L(\theta) \coloneqq W_1(\theta) - W_0(\theta).$$

I assume that $L : \mathbb{V} \to \mathbb{R}$ is a linear function. The optimal policy under $\theta$ is policy 1 if $L(\theta) > 0$, policy 0 if $L(\theta) < 0$, and either of the two if $L(\theta) = 0$.

One example of a welfare criterion is a weighted average of an outcome across individuals. For example, suppose that a policy could change the outcome of each individual in the population. Suppose also that we specify $\theta = (f_1(\cdot), f_0(\cdot))$, where $f_a(x)$ represents the counterfactual mean outcome under policy $a$ across individuals whose observed covariates are $x$. The welfare under policy $a$ can be defined, for example, by the population mean outcome $W_a(\theta) = \int f_a(x) dP_X$, where $P_X$ is the probability measure of covariates and is assumed to be known. In this case, the welfare difference $L(\theta) = \int [f_1(x) - f_0(x)] dP_X$ is linear in $\theta = (f_1(\cdot), f_0(\cdot))$. If we are required to take the policy cost into account, we can incorporate it into the welfare by redefining the outcome to be the raw outcome minus the cost. On the other hand, the linearity of $L$ may rule out welfare criteria that depend on the distribution of the counterfactual outcome. See Kitagawa and Tetenov (2021) for such welfare criteria.

Importantly, this framework allows for cases in which $L(\theta)$ is not point identified in the sense that the identified set of $L(\theta)$ when $\boldsymbol{m}(\theta) = \boldsymbol{\mu}$, namely

$$\{L(\theta) : \boldsymbol{m}(\theta) = \boldsymbol{\mu}, \theta \in \Theta\},$$

is nonsingleton for some or all $\boldsymbol{\mu} \in \mathbb{R}^n$. This is the set of possible values of $L(\theta)$ consistent with the observed value of $\boldsymbol{Y} \in \mathbb{R}^n$ when there is no sampling uncertainty. If the identified set contains both positive and negative values, which policy we should choose is ambiguous even without sampling uncertainty. Whether $L(\theta)$ is point identified or not depends on the parameter space $\Theta$.

This framework nests some existing setups of treatment choice, such as limit experiments under parametric models by Hirano and Porter (2009), Gaussian experiments with limited validity by Stoye (2012), and a setup of policy choice based on multiple studies by Ishihara and Kitagawa (2021).[6] One of the essential departures from these setups is that the parameter $\theta$ can be infinite dimensional, accommodating nonparametric regression models, for example.[7]

## 2.2 Optimality Criterion

This paper considers the minimax regret criterion as an optimality criterion, following existing treatment choice studies (e.g., Manski, 2004, 2007; Stoye, 2009, 2012; Kitagawa and Tetenov, 2018).[8] I define a few concepts to introduce the minimax regret criterion. The *welfare regret loss* for policy choice $a \in \{0, 1\}$ is

$$l(a,\theta) \coloneqq \max_{a' \in \{0,1\}} W_{a'}(\theta) - W_a(\theta) = \begin{cases} L(\theta) \cdot (1-a) & \text{if } L(\theta) \geq 0, \\ -L(\theta) \cdot a & \text{if } L(\theta) < 0. \end{cases}$$

The welfare regret loss $l(a, \theta)$ is the difference between the welfare under the optimal policy and the welfare under policy $a$ under $\theta$. If the policy maker chooses the superior policy, they do not incur any loss; otherwise, they incur a loss of the absolute value of the welfare difference $L(\theta)$. A (randomized) *decision rule* is a measurable function $\delta : \mathbb{R}^n \to [0, 1]$, where $\delta(\boldsymbol{y})$ represents the probability of choosing policy 1 when the realization of the sample $\boldsymbol{Y}$ is $\boldsymbol{y}$. The risk or *regret* of decision rule $\delta$ under $\theta$ is the expected welfare regret loss

$$R(\delta,\theta) \coloneqq \begin{cases} L(\theta)(1-\mathbb{E}_\theta[\delta(\boldsymbol{Y})]) & \text{if } L(\theta) \geq 0, \\ -L(\theta)\mathbb{E}_\theta[\delta(\boldsymbol{Y})] & \text{if } L(\theta) < 0, \end{cases}$$

[6] Ishihara and Kitagawa (2021) do not impose convexity of the parameter space to derive their results. However, the specific examples of the parameter space that they consider satisfy convexity.

[7] Hirano and Porter (2009) consider a model with an infinite Gaussian sequence as a limit experiment under semiparametric models. They do not allow for a partially identified welfare difference—one of the crucial aspects of this paper's setup.

[8] Alternative criteria include the maximin criterion, which solves $\sup_\delta \inf_{\theta\in\Theta} U(\delta,\theta)$, where $U(\delta,\theta) = W_1(\theta)\mathbb{E}_\theta[\delta(\boldsymbol{Y})] + W_0(\theta)(1-\mathbb{E}_\theta[\delta(\boldsymbol{Y})])$ is the expected welfare under decision rule $\delta$ under $\theta$. It has been pointed out that the maximin criterion is unreasonably pessimistic and can lead to pathological decision rules (Manski, 2004). Another approach is the Bayesian one, which solves $\sup_\delta \int U(\delta,\theta)d\pi(\theta)$, where $\pi$ is a prior on the vector space $\mathbb{V}$ that $\theta$ belongs to. In practice, when it is difficult to make a prior, the minimax regret criterion is a reasonable choice. See Manski (2021) for a detailed discussion.

where $\mathbb{E}_\theta$ denotes the expectation taken with respect to $\boldsymbol{Y}$ under $\theta$.[9]

Given a particular choice of $\Theta$, I evaluate decision rules based on the maximum regret over $\Theta$, $\sup_{\theta\in\Theta} R(\delta,\theta)$. My goal is to derive a *minimax regret* decision rule, which achieves

$$\inf_{\delta}\sup_{\theta\in\Theta} R(\delta,\theta),$$

where the infimum is taken over the set of all possible decision rules. I do not impose any restrictions on the class of decision rules.

To sum up, the minimax regret criterion deals with the sampling uncertainty given $\theta$ by taking the expectation of the welfare regret loss with respect to the distribution of $\boldsymbol{Y}$. It then deals with the parameter $\theta$ by considering the worst-case expected welfare regret loss. It does not distinguish between the case where the welfare difference $L(\theta)$ is point identified and the case where it is not. Nevertheless, as I show in Section 3, the minimax regret rule behaves differently in each case.

## 2.3 Examples

### 2.3.1 Motivating Example: Eligibility Cutoff Choice in RD Designs

In many policy domains, ranging from health to education to social programs, the eligibility for treatment is determined based on an individual's observable characteristics. A critical policy question is whether we should change the eligibility criterion to achieve better welfare (Dong and Lewbel, 2015).[10] My framework can be of use for policy makers interested in utilizing data to make such decisions.

Consider the following RD setup. For each unit $i = 1, .., n$, we observe a fixed running variable $x_i \in \mathbb{R}$, a binary treatment status $d_i \in \{0, 1\}$, and an outcome $Y_i \in \mathbb{R}$. The eligibility for treatment is determined based on whether the running variable exceeds a specific cutoff $c_0 \in \mathbb{R}$, so that $d_i = \mathbf{1}\{x_i \geq c_0\}$. Suppose that the

[9] More formally, the loss function of a randomized decision rule $\delta$ is defined as $\tilde{l}(\delta(\boldsymbol{y}),\theta) \coloneqq E[l(A,\theta)]$, where $A$ follows a Bernoulli distribution with success probability $\delta(\boldsymbol{y})$, for each possible realization $\boldsymbol{y}$. The risk of $\delta$ is then defined as $R(\delta,\theta) \coloneqq E_\theta[\tilde{l}(\delta(\boldsymbol{Y}),\theta)]$ (Berger, 1985, Chapter 1).

[10] For example, there is a heated debate about whether to extend Medicare eligibility in the United States (Song, 2020).

outcome $Y_i$ is of the form

$$Y_i = f(x_i, d_i) + u_i, \quad u_i \sim \mathcal{N}(0, \sigma^2(x_i, d_i)) \text{ independent across } i, \tag{2}$$

where $f : \mathbb{R} \times \{0,1\} \to \mathbb{R}$ is an unknown function and plays the role of the parameter $\theta$. We interpret $f(x, d)$ as the counterfactual mean outcome across individuals with running variable $x$ if their treatment status is set to $d \in \{0, 1\}$.[11] We can write the model in a vector form $\boldsymbol{Y} \sim \mathcal{N}(\boldsymbol{m}(f), \boldsymbol{\Sigma})$, where $\boldsymbol{Y} = (Y_1, ..., Y_n)'$, $\boldsymbol{m}(f) = (f(x_1, d_1), ..., f(x_n, d_n))'$, and $\boldsymbol{\Sigma} = \mathrm{diag}(\sigma^2(x_1, d_1), ..., \sigma^2(x_n, d_n))$.

Now, suppose that we are interested in changing the eligibility cutoff from $c_0$ to a specific value $c_1$. For illustration purposes, I assume $c_1 < c_0$. Suppose that the welfare under the cutoff $c_a$, with $a \in \{0, 1\}$, is an average of the counterfactual mean outcome across different values of the running variable

$$W_a(f) = \int [f(x, 1)\mathbf{1}\{x \geq c_a\} + f(x, 0)\mathbf{1}\{x < c_a\}]d\nu(x)$$

for some known measure $\nu$.[12] One choice of $\nu$ is an empirical measure, for which the welfare is the unweighted sample average: $W_a(f) = \frac{1}{n}\sum_{i=1}^n [f(x_i, 1)\mathbf{1}\{x_i \geq c_a\} + f(x_i, 0)\mathbf{1}\{x_i < c_a\}]$. The welfare difference between the two cutoffs is

$$L(f) = W_1(f) - W_0(f) = \int \mathbf{1}\{c_1 \leq x < c_0\}[f(x, 1) - f(x, 0)]d\nu(x),$$

which is a linear function of $f$. $L(f)$ is a weighted sum of the conditional average treatment effect $f(x, 1) - f(x, 0)$ across different values of the running variable between the two cutoffs $c_1$ and $c_0$.

To conclude the problem's setup, suppose that $f \in \mathcal{F}$, where $\mathcal{F}$ is a known set

[11] This interpretation is established in a potential outcomes model as follows (Armstrong and Kolesár, 2021). Suppose we observe a triple of the outcome, treatment status, and running variable $(Y_i, D_i, X_i)$. The observed outcome is $Y_i = Y_i(1)D_i + Y_i(0)(1 - D_i)$, where $Y_i(1)$ and $Y_i(0)$ are potential outcomes under treatment and no treatment, respectively. Let $f(x, d) = \mathbb{E}[Y_i(d)|X_i = x]$, which is equal to $\mathbb{E}[Y_i|X_i = x, D_i = d]$ if $D_i$ is a deterministic function of $X_i$. We obtain model (2) by conditioning on the realized values $\{(x_i, d_i)\}_{i=1}^n$ and assuming normal conditional errors.

[12] An implicit assumption made here is that $f$ is invariant to the cutoff change, which is in a similar spirit to (but stronger than) the local policy invariance condition introduced by Dong and Lewbel (2015).

of functions and plays the role of the parameter space $\Theta$. For an illustration of the results and empirical application, I focus on the *Lipschitz class* with a known Lipschitz constant $C \geq 0$:[13]

$$\mathcal{F}_{\text{Lip}}(C) = \{f : |f(x,d) - f(\tilde{x},d)| \leq C|x - \tilde{x}| \text{ for every } x, \tilde{x} \in \mathbb{R} \text{ and } d \in \{0,1\}\}.$$

The Lipschitz constraint bounds the maximum possible change in $f(x,d)$ in response to a shift in $x$ by one unit. In other words, the absolute value of the derivative of $f(x,d)$ with respect to $x$ must be at most $C$ if $f$ is differentiable. The Lipschitz class $\mathcal{F}_{\text{Lip}}(C)$ is both convex and centrosymmetric.

Imposing $f \in \mathcal{F}_{\text{Lip}}(C)$ is not strong enough to uniquely determine $L(f)$ from a given value of $\boldsymbol{m}(f) = (f(x_1,d_1),...,f(x_n,d_n))'$. Nevertheless, it produces an informative identified set of $L(f)$ since it gives finite upper and lower bounds on $f(x,d)$ for every $(x,d) \in \mathbb{R} \times \{0,1\}$ from the knowledge of $(f(x_1,d_1),...,f(x_n,d_n))$.[14]

In Section 4, I derive a minimax regret rule for this example when the welfare is the sample average outcome and $\mathcal{F} = \mathcal{F}_{\text{Lip}}(C)$.

**Remark 1.** In this example, $\delta(\boldsymbol{y}) \in [0,1]$ represents the probability of changing the cutoff from $c_0$ to $c_1$ when we apply decision rule $\delta$ to the realized sample $\boldsymbol{y}$. Alternatively, we can interpret $\delta(\boldsymbol{y})$ as the fraction of individuals to whom we would assign treatment within the subpopulation whose running variable lies between $c_1$ and $c_0$. To see this, suppose action $a \in [0,1]$ represents the fraction, instead of the binary cutoff choice, and define the welfare function as follows: for $a \in [0,1]$,

$$\tilde{W}_a(f) = W_0(f) + \int a \cdot \mathbf{1}\{c_1 \leq x < c_0\}[f(x,1) - f(x,0)]d\nu(x) = W_0(f) + a \cdot L(f).$$

Consider the following welfare regret loss: for $a \in [0,1]$,

$$\tilde{l}(a,f) = \max_{a' \in [0,1]} W_{a'}(f) - W_a(f) = \begin{cases} L(f) \cdot (1-a) & \text{if } L(f) \geq 0, \\ -L(f) \cdot a & \text{if } L(f) < 0. \end{cases}$$

[13] The Lipschitz class is used by, for example, Armstrong and Kolesár (2021) and Kwon and Kwon (2020) for inference on the average treatment effect under unconfoundedness and for inference on the RD parameter in RD designs, respectively.

[14] Given a value of $(f(x_1,d_1),...,f(x_n,d_n))$, the upper bound on $f(x,d)$ is $\min_{i:d_i=d}(f(x_i,d) + C|x_i - x|)$. The lower bound on $f(x,d)$ is $\max_{i:d_i=d}(f(x_i,d) - C|x_i - x|)$.

The risk of decision rule $\delta : \mathbb{R}^n \to [0,1]$ for this setup, $\mathbb{E}_f[\tilde{l}(\delta(\boldsymbol{Y}), f)]$, is equal to the risk for the setup with a binary action in Section 2.2. Therefore, these two setups lead to the mathematically equivalent minimax regret problem.

### 2.3.2 Existing Examples

One of the simplest examples is where the policy maker must decide whether to treat members of a population based on an estimator for the average treatment effect from a randomized experiment. In the notation of my framework, we observe a scalar estimator $Y \sim \mathcal{N}(m(\theta), \sigma^2)$, $\theta$ represents the average treatment effect, $\Theta = \mathbb{R}$, and $m(\theta) = L(\theta) = \theta$. The results from Hirano and Porter (2009) and Tetenov (2012) show that $\delta^*(Y) = \mathbf{1}\{Y \geq 0\}$ is a minimax regret rule.

Stoye (2012) considers an extended setup where the experiment may have limited validity because of selective noncompliance or because the treatment population is different from the sampling population. He formalizes such situations as a decision problem with partial identification. In the notation of my framework, we observe a scalar sample $Y \sim \mathcal{N}(m(\theta), \sigma^2)$, $\theta = (\theta_1, \theta_2)' \in \mathbb{R}^2$, $m(\theta) = \theta_1$, $\Theta = \{(\theta_1, \theta_2)' \in [-1,1]^2 : \theta_2 \in [a\theta_1 - b, a\theta_1 + b]\}$ for some known constants $a \in (0,1]$ and $b > 0$, and $L(\theta) = \theta_2$. Here, $\theta_1$ and $\theta_2$ represent the average treatment effects for the sampling and treatment populations, respectively, $Y$ is an estimator for $\theta_1$ from an experiment, and $[a\theta_1 - b, a\theta_1 + b] \cap [-1,1]$ is the identified set of $\theta_2$ given $\theta_1$. Stoye (2012) derives a minimax regret rule within the class of all decision rules (see Section 3.3 for its expression). My framework generalizes this setup in three ways: (1) the sample can be multidimensional; (2) the parameter can be three or higher dimensional, even infinite dimensional; and (3) flexible forms of the parameter space are allowed.

Ishihara and Kitagawa (2021) consider the problem of deciding whether or not to introduce a new policy to a specific local population based on causal evidence of similar policies implemented in other populations. In the notation of my framework, we observe an $n$-dimensional sample $\boldsymbol{Y} \sim \mathcal{N}(\boldsymbol{m}(\theta), \boldsymbol{\Sigma})$, $\theta = (\theta_0, \theta_1, ..., \theta_n)' \in \Theta \subset \mathbb{R}^{n+1}$, $\boldsymbol{m}(\theta) = (\theta_1, ..., \theta_n)'$, and $L(\theta) = \theta_0$. Here, $\theta_0$ is the average welfare effect of a new policy on the target population, $\theta_1, ..., \theta_n$ are the average welfare effects on $n$ study populations that are different from the target population, and $\Theta$ imposes restrictions on the differences between $\theta_i$'s. Ishihara and Kitagawa (2021) obtain a characterization of a minimax regret rule within the class of nonrandomized decision rules that make a decision based on the sign of a linear function of $\boldsymbol{Y}$. This paper derives a minimax regret rule among all decision rules in a more general setup.

### 2.3.3 Other Examples

The example in Section 2.3.1 can be easily generalized to a setup where the observed treatment is independent of counterfactual outcomes conditional on multidimensional covariates (i.e., the unconfoundedness assumption holds) and there is no or limited overlap in the covariate distribution between the treatment and control groups. This general setup covers the problem of which one of two eligibility criteria (or treatment allocation policies) based on multiple covariates should be implemented. In this setup, the welfare difference between the two criteria may be partially identified under smoothness restrictions such as Lipschitz constraints.

Another example is the policy adoption decision using a difference-in-differences design. Consider a group of units that have experienced a policy change and another group that has not. Suppose that the policy maker needs to decide whether to implement the new policy for the latter group. The average policy effect on this group may only be partially identified if either the parallel trends assumption is violated or the policy effect varies between the two groups. My framework can be applied to this problem by imposing a set of restrictions on the degree of the violation of parallel trends and the amount of heterogeneity in policy effects. Rambachan and Roth (2023) consider robust inference in a related setup.

# 3 Results in General Setup

In this section, I derive a minimax regret rule and discuss its interpretations and implications in the general setup. For simplicity, I normalize $\boldsymbol{\Sigma} = \sigma^2 \boldsymbol{I}_n$ for some $\sigma > 0$, where $\boldsymbol{I}_n$ is the identity matrix. This normalization is without loss of generality since $\boldsymbol{\Sigma}$ is known.

Solving minimax problems over the class of all decision rules is generally a difficult task. The main tool that I use to solve the minimax regret problem is the *modulus* *of continuity*, defined as

$$\omega(\epsilon; L, \boldsymbol{m}, \Theta) := \sup\{L(\theta) : \|\boldsymbol{m}(\theta)\| \leq \epsilon, \theta \in \Theta\}, \quad \epsilon \geq 0,$$

where $\|\cdot\|$ is the Euclidean norm. The modulus of continuity and its variants have been used in constructing minimax optimal estimators and confidence intervals on linear functionals in Gaussian models (Donoho, 1994; Low, 1995; Cai and Low, 2004; Armstrong and Kolesár, 2018).[15] This paper is the first to use the modulus of continuity to derive minimax regret rules for the problem of treatment choice. Below, I suppress the arguments $L, \boldsymbol{m}$, and $\Theta$ if they are clear from the context.

By definition, $\omega(\epsilon)$ is nondecreasing in $\epsilon$. Furthermore, $\omega(\epsilon)$ is concave in $\epsilon$ if $\Theta$ is convex.[16]

In the context of this paper, the modulus of continuity at $\epsilon$ is the largest possible welfare difference under the constraint that the norm of $\boldsymbol{m}(\theta)$, namely the expected value of $\boldsymbol{Y}$, is less than or equal to $\epsilon$. When $\epsilon = 0$ and hence the expected value of $\boldsymbol{Y}$ must be a vector of zeros, the sample $\boldsymbol{Y}$ is uninformative. When the norm constraint $\|\boldsymbol{m}(\theta)\| \leq \epsilon$ is relaxed, the strength of $\boldsymbol{Y}$ as a signal for $L(\theta)$ may increase, which makes it easier for the policy maker to detect the optimal policy. At the same time, the largest potential welfare loss when choosing the inferior policy may increase since the flexibility of $\theta$ increases because of the weaker norm constraint. The modulus of continuity is used to trade off these two criteria and find parameter values that are least favorable for the policy maker.

## 3.1 Sketch of Derivation of Minimax Regret Rule

Here, I sketch the derivation of a minimax regret rule, deferring the statement of the assumptions and the complete proof to Section 3.2 and Appendix A.3, respectively.

I introduce some notation. I say that $\theta_\epsilon \in \Theta$ *attains the modulus of continuity at* $\epsilon$ if $\theta_\epsilon \in \arg\max_{\theta \in \Theta : \|\boldsymbol{m}(\theta)\| \leq \epsilon} L(\theta)$. I use $\mathcal{R}(\sigma; \Theta)$ to denote the *minimax risk* $\inf_\delta \sup_{\theta \in \Theta} R(\delta, \theta)$, which may depend on the standard deviation $\sigma$ and on the choice of the parameter space $\Theta$ among others. Given any two parameter values $\tilde{\theta}, \bar{\theta} \in \mathbb{V}$,

[15] Donoho (1994) defines the modulus of continuity as $\tilde{\omega}(\epsilon) = \sup\{|L(\theta) - L(\tilde{\theta})| : \|\boldsymbol{m}(\theta - \tilde{\theta})\| \leq \epsilon, \theta, \tilde{\theta} \in \Theta\}$. If $\Theta$ is convex and centrosymmetric, the relationship $\tilde{\omega}(\epsilon) = 2\omega(\epsilon/2)$ holds.

[16] See, for example, Donoho (1994, Lemma 3) and Armstrong and Kolesár (2018, Appendix A).

where $\mathbb{V}$ is the vector space that the parameter $\theta$ belongs to, I define a *one-dimensional subfamily* as the set of all convex combinations of $\tilde{\theta}$ and $\bar{\theta}$, denoted by $[\tilde{\theta}, \bar{\theta}] = \{(1-\lambda)\tilde{\theta} + \lambda\bar{\theta} : \lambda \in [0,1]\}$. Let $\mathcal{R}(\sigma; [\tilde{\theta}, \bar{\theta}])$ denote the minimax risk $\inf_{\delta} \sup_{\theta \in [\tilde{\theta}, \bar{\theta}]} R(\delta, \theta)$ for this subproblem. Additionally, let $\Phi$ and $\phi$ denote the cumulative distribution function and the probability density function, respectively, of a standard normal random variable. Lastly, let $a^* \in \arg\max_{a \geq 0} a\Phi(-a)$, which is unique ($a^* \approx 0.752$).

Below, I first derive minimax regret rules for one-dimensional subproblems of the form $[-\bar{\theta}, \bar{\theta}]$. I then use the modulus of continuity to characterize the hardest one-dimensional subproblem, namely the one that has the largest minimax risk $\mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}])$ among all one-dimensional subproblems of this form. Lastly, I explain that the set of minimax regret rules for the hardest one-dimensional subproblem contains a minimax regret rule for the original problem.

### 3.1.1 Minimax Regret Rules for One-dimensional Subproblems

The following result derives minimax regret rules for one-dimensional subproblems.

**Lemma 1** (Minimax Regret Rules for One-dimensional Subproblems)**.** *Suppose that* $\Theta = [-\bar{\theta}, \bar{\theta}]$*, where* $\bar{\theta} \in \mathbb{V}$ *and* $L(\bar{\theta}) \geq 0$*. If* $\boldsymbol{m}(\bar{\theta}) \neq \mathbf{0}$*, then the decision rule* $\delta^*(\boldsymbol{Y}) = \mathbf{1}\left\{\boldsymbol{m}(\bar{\theta})'\boldsymbol{Y} \geq 0\right\}$ *is minimax regret. If* $\boldsymbol{m}(\bar{\theta}) = \mathbf{0}$*, then any decision rule* $\delta^*$ *such that* $\mathbb{E}[\delta^*(\boldsymbol{Y}^*)] = \frac{1}{2}$*, where* $\boldsymbol{Y}^* \sim \mathcal{N}(\mathbf{0}, \sigma^2\boldsymbol{I}_n)$*, is minimax regret. The minimax risk is given by*

$$\mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}]) = \begin{cases} L(\bar{\theta})\Phi\left(-\frac{\|\boldsymbol{m}(\bar{\theta})\|}{\sigma}\right) & \text{if } \|\boldsymbol{m}(\bar{\theta})\| \leq a^*\sigma, \\ a^*\sigma\frac{L(\bar{\theta})}{\|\boldsymbol{m}(\bar{\theta})\|}\Phi\left(-a^*\right) & \text{if } \|\boldsymbol{m}(\bar{\theta})\| > a^*\sigma. \end{cases}$$

*Proof.* See Appendix A.2. □

If $\boldsymbol{m}(\bar{\theta}) \neq \mathbf{0}$, the sample $\boldsymbol{Y}$ is informative for $\theta$ within the subfamily $[-\bar{\theta}, \bar{\theta}]$. In this case, $\boldsymbol{m}(\bar{\theta})'\boldsymbol{Y}$ is shown to be a sufficient statistic of $\boldsymbol{Y}$ for $\theta$. The minimax regret rule $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\bar{\theta})'\boldsymbol{Y} \geq 0\}$ makes a decision based on the sign of the sufficient statistic. On the other hand, if $\boldsymbol{m}(\bar{\theta}) = \mathbf{0}$, $\boldsymbol{Y} \sim \mathcal{N}(\mathbf{0}, \sigma^2\boldsymbol{I}_n)$ under any $\theta \in [-\bar{\theta}, \bar{\theta}]$, and therefore $\boldsymbol{Y}$ is uninformative for $\theta$. In this case, any decision rule that chooses each policy with probability one half over the distribution of $\boldsymbol{Y}$ is minimax regret for the subproblem $[-\bar{\theta}, \bar{\theta}]$. The minimax risk is obtained by computing the maximum regret $\sup_{\theta \in [-\bar{\theta}, \bar{\theta}]} R(\delta^*, \theta)$.

### 3.1.2 Hardest One-dimensional Subproblem

For simplicity, I assume here that for each $\epsilon \geq 0$, there exists a value of $\theta_\epsilon$ that attains the modulus of continuity at $\epsilon$ with $\|\boldsymbol{m}(\theta_\epsilon)\| = \epsilon$, so that $\omega(\epsilon) = \sup_{\theta \in \Theta: \|\boldsymbol{m}(\theta)\| = \epsilon} L(\theta)$. The minimax risk for the hardest one-dimensional subproblem can then be expressed in terms of the modulus of continuity:

$$
\begin{aligned}
&\sup_{\bar{\theta} \in \Theta} \mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}]) = \sup_{\bar{\theta} \in \Theta: L(\bar{\theta}) \geq 0} \mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}]) = \sup_{\epsilon \geq 0} \sup_{\bar{\theta} \in \Theta: \|\boldsymbol{m}(\bar{\theta})\| = \epsilon, L(\bar{\theta}) \geq 0} \mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}]) \\
&= \sup \left\{ \sup_{\epsilon \in [0, a^* \sigma]} \sup_{\bar{\theta} \in \Theta: \|\boldsymbol{m}(\bar{\theta})\| = \epsilon} L(\bar{\theta}) \Phi\left(-\frac{\|\boldsymbol{m}(\bar{\theta})\|}{\sigma}\right), \sup_{\epsilon > a^* \sigma} \sup_{\bar{\theta} \in \Theta: \|\boldsymbol{m}(\bar{\theta})\| = \epsilon} \frac{a^* \sigma L(\bar{\theta})}{\|\boldsymbol{m}(\bar{\theta})\|} \Phi\left(-a^*\right) \right\} \\
&= \sup \left\{ \sup_{\epsilon \in [0, a^* \sigma]} \omega(\epsilon) \Phi(-\epsilon/\sigma), \sup_{\epsilon > a^* \sigma} \frac{a^* \sigma \omega(\epsilon)}{\epsilon} \Phi(-a^*) \right\},
\end{aligned}
$$

where the first equality holds since restricting attention to $\bar{\theta}$ with $L(\bar{\theta}) \geq 0$ does not change the supremum by the centrosymmetry of $\Theta$ and the third equality uses Lemma 1. Furthermore, since $\frac{\omega(\epsilon)}{\epsilon}$ is continuous and nonincreasing on $(0, \infty)$ by the concavity of $\omega(\epsilon)$, $\sup_{\epsilon > a^* \sigma} \frac{a^* \sigma \omega(\epsilon)}{\epsilon} \Phi(-a^*) = \omega(a^* \sigma) \Phi(-a^*)$. The above expression can then be simplified into:

$$
\sup_{\bar{\theta} \in \Theta} \mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}]) = \sup_{\epsilon \in [0, a^* \sigma]} \omega(\epsilon) \Phi(-\epsilon/\sigma).
$$

Now, let $\epsilon^*$ solve the maximization problem on the right-hand side. $\epsilon^*$ balances the potential welfare loss $(\omega(\epsilon))$ and the probability of incurring loss $(\Phi\left(-\epsilon/\sigma\right))$. The corresponding subfamily $[-\theta_{\epsilon^*}, \theta_{\epsilon^*}]$ has the largest minimax risk among all one-dimensional subfamilies, where $\theta_{\epsilon^*}$ attains the modulus of continuity at $\epsilon^*$.

### 3.1.3 Minimax Regret Rule for the Original Problem

A minimax regret rule for the original problem can be derived by finding a decision rule $\delta^*$ and a one-dimensional subfamily $[-\bar{\theta}, \bar{\theta}] \subset \Theta$ that satisfy the following property.

**Property 1.** $\delta^*$ is minimax regret for the subproblem $[-\bar{\theta},\bar{\theta}] \subset \Theta$, that is, $\max_{\theta\in[-\bar{\theta},\bar{\theta}]} R(\delta^*,\theta) = \mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}])$. Furthermore, the maximum regret of $\delta^*$ over $\Theta$ is attained at some $\theta \in [-\bar{\theta},\bar{\theta}]$, that is, $\max_{\theta\in\Theta} R(\delta^*,\theta) = \max_{\theta\in[-\bar{\theta},\bar{\theta}]} R(\delta^*,\theta)$.

If $\delta^*$ and $[-\bar{\theta},\bar{\theta}]$ satisfy Property 1, then $\delta^*$ is shown to be minimax regret for the original problem as follows. By Property 1, $\max_{\theta\in\Theta} R(\delta^*,\theta) = \mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}])$. On the other hand, by the definition of the minimax risk,

$$\max_{\theta\in\Theta} R(\delta^*,\theta) \geq \mathcal{R}(\sigma;\Theta) \geq \mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}]).$$

Thus, $\max_{\theta\in\Theta} R(\delta^*,\theta) = \mathcal{R}(\sigma;\Theta)$, so $\delta^*$ is minimax regret for the original problem.

Here, I briefly explain how I find $\delta^*$ and $[-\bar{\theta},\bar{\theta}]$ that satisfy Property 1. First, for the case where $\epsilon^* > 0$ so that $\boldsymbol{m}(\theta_{\epsilon^*}) \neq \boldsymbol{0}$, I take the following three steps. The first step proves the existence of $\bar{\theta} \in \Theta$ such that the maximum regret of the rule $\delta^*(\boldsymbol{Y}) = \mathbf{1}\left\{\boldsymbol{m}(\bar{\theta})'\boldsymbol{Y} \geq 0\right\}$ over a specific subset $\tilde{\Theta}$ of $\Theta$ is attained at $\bar{\theta}$, that is, $\bar{\theta} \in \arg\max_{\theta\in\tilde{\Theta}} R(\delta^*,\theta)$, where $\tilde{\Theta}$ contains $\theta_{\epsilon^*}$. The second step shows that $\boldsymbol{m}(\bar{\theta}) = \boldsymbol{m}(\theta_{\epsilon^*})$ and $R(\delta^*,\bar{\theta}) = R(\delta^*,\theta_{\epsilon^*})$. The former implies that $\delta^*$ is minimax regret for $[-\theta_{\epsilon^*},\theta_{\epsilon^*}]$ by Lemma 1, and the latter implies that $\theta_{\epsilon^*}$, which may or may not be equal to $\bar{\theta}$, also attains the maximum regret of $\delta^*$ over $\tilde{\Theta}$. The last step shows that $\theta_{\epsilon^*}$ attains the maximum regret of $\delta^*$ even over the original parameter space $\Theta$, proving that $\delta^*$ and $[-\theta_{\epsilon^*},\theta_{\epsilon^*}]$ satisfy Property 1.

Next, consider the case where $\epsilon^* = 0$ so that $\boldsymbol{m}(\theta_{\epsilon^*}) = \boldsymbol{0}$. Let $\boldsymbol{w}^* = \lim_{\epsilon\to 0} \frac{\boldsymbol{m}(\theta_\epsilon)}{\|\boldsymbol{m}(\theta_\epsilon)\|}$, where $\theta_\epsilon$ attains the modulus of continuity at $\epsilon$. I consider the class of rules $\delta_{\sigma^*}$ indexed by $\sigma^* \geq \sigma$ such that

$$\delta_{\sigma^*}(\boldsymbol{Y}) = \begin{cases} \mathbf{1}\{(\boldsymbol{w}^*)'\boldsymbol{Y} \geq 0\} & \text{if } \sigma^* = \sigma, \\ \Phi\left(\dfrac{(\boldsymbol{w}^*)'\boldsymbol{Y}}{((\sigma^*)^2 - \sigma^2)^{1/2}}\right) & \text{if } \sigma^* > \sigma. \end{cases}$$

For any $\sigma^* \geq \sigma$, $\mathbb{E}[\delta_{\sigma^*}(\boldsymbol{Y}^*)] = \frac{1}{2}$, where $\boldsymbol{Y}^* \sim \mathcal{N}(\boldsymbol{0},\sigma^2\boldsymbol{I}_n)$, and therefore $\delta_{\sigma^*}$ is minimax regret for $[-\theta_{\epsilon^*},\theta_{\epsilon^*}]$ by Lemma 1. I then find a value of $\sigma^*$ such that the maximum regret of $\delta_{\sigma^*}$ is attained at $\theta_{\epsilon^*}$. Property 1 is satisfied by $\delta_{\sigma^*}$ and $[-\theta_{\epsilon^*},\theta_{\epsilon^*}]$.

**Remark 2.** Donoho (1994) shows the existence of a procedure and a one-dimensional subfamily that satisfy Property 1 for minimax affine estimation and inference problems. However, his proof relies on the following property of the risks (e.g., the MSE) considered by his paper: the risk can be expressed in terms of bias and variance of an affine estimator, and the maximum risk and maximum bias are attained at the same parameter values because the variance of an affine estimator is invariant to the parameter under known variance. His technique does not apply to minimax regret problems since regret cannot be expressed in terms of bias and variance.

## 3.2 Main Result: Minimax Regret Rule

I now present a minimax regret rule. To derive the result, I impose the following restrictions on $L$, $\boldsymbol{m}$, and $\Theta$.

**Assumption 1** (Regularity)**.**

*(i) For some $\bar{\epsilon} > 0$, there exists $\{\theta_\epsilon\}_{\epsilon \in [0,\bar{\epsilon}]}$ with $\theta_\epsilon \in \Theta$ such that the following holds.*

- *(a) For all $\epsilon \in [0, \bar{\epsilon}]$, $\theta_\epsilon$ attains the modulus of continuity at $\epsilon$.*
- *(b) There exists $\boldsymbol{w}^* \in \mathbb{R}^n$ such that $\lim_{\epsilon \to 0} \epsilon^{-1} \left( \boldsymbol{w}^* - \frac{\boldsymbol{m}(\theta_\epsilon)}{\|\boldsymbol{m}(\theta_\epsilon)\|} \right) = \mathbf{0}$.*
- *(c) There exists $\iota \in \Theta$ with $L(\iota) \neq 0$ such that for all $\epsilon \in [0, \bar{\epsilon}]$, $\theta_\epsilon + c\iota \in \Theta$ for all $c$ in a neighborhood of zero.*
- *(d) $\rho(\cdot)$ is differentiable at any $\epsilon \in \{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\}$, where $\rho(\epsilon) = \sup\{L(\theta) : (\boldsymbol{w}^*)'\boldsymbol{m}(\theta) = \epsilon, \theta \in \Theta\}$.*

*(ii) $\omega(\cdot)$ is differentiable at any $\epsilon \in [0, a^*\sigma]$.*

Assumption 1(i)(a) says that the modulus of continuity is attained for all sufficiently small $\epsilon \geq 0$, which typically holds if $\Theta$ is closed.[17] Assumption 1(i)(b) requires that the unit vector $\frac{\boldsymbol{m}(\theta_\epsilon)}{\|\boldsymbol{m}(\theta_\epsilon)\|} \in \mathbb{R}^n$ converge to some constant $\boldsymbol{w}^*$ faster than $\epsilon$ as $\epsilon \to 0$. The limit $\boldsymbol{w}^*$ can be viewed as the direction at which the welfare difference $L(\theta)$ increases the most when we move $\boldsymbol{m}(\theta)$ away from $\mathbf{0}$. Assumption 1(i)(c) says that there exists $\iota \in \Theta$ with $L(\iota) \neq 0$ such that $\theta_\epsilon$ lies in $\Theta$ even after receiving a small perturbation in the direction of $\iota$. Assumption 1(i)(d) and (ii) assume the differentiability of $\rho(\epsilon) = \sup\{L(\theta) : (\boldsymbol{w}^*)'\boldsymbol{m}(\theta) = \epsilon, \theta \in \Theta\}$ and $\omega(\epsilon)$. I provide sufficient conditions for the differentiability in Lemmas A.3 and A.5 in Appendix A.1.

[17] See Donoho (1994, Lemma 2) for sufficient conditions.

I use Assumption 1(i)(c)–(d) and (ii) to obtain an expression for a minimax regret decision rule in terms of the modulus of continuity. In Appendix A.3, I present a minimax regret rule under relaxed conditions.

The following theorem derives a minimax regret rule.

**Theorem 1** (Minimax Regret Rule)**.** *Let $\Theta$ be convex and centrosymmetric, and suppose that Assumption 1 holds.*[18] *Then, the following holds.*

*(i) There exists a unique solution $\epsilon^*$ to the maximization problem $\max_{\epsilon\in[0,a^*\sigma]}\omega(\epsilon)\Phi(-\epsilon/\sigma)$. Moreover, $\epsilon^*>0$ if and only if $\sigma > 2\phi(0)\frac{\omega(0)}{\omega'(0)}$.*[19]

*(ii) Suppose that there exists $\theta_{\epsilon^*}\in\Theta$ that attains the modulus of continuity at $\epsilon^*$. Then, the following decision rule is minimax regret:*

$$\delta^*(\boldsymbol{Y}) = \begin{cases} \mathbf{1}\left\{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y}\geq 0\right\} & \text{if } \sigma > 2\phi(0)\frac{\omega(0)}{\omega'(0)}, \\ \mathbf{1}\left\{(\boldsymbol{w}^*)'\boldsymbol{Y}\geq 0\right\} & \text{if } \sigma = 2\phi(0)\frac{\omega(0)}{\omega'(0)}, \\ \Phi\left(\dfrac{(\boldsymbol{w}^*)'\boldsymbol{Y}}{((2\phi(0)\omega(0)/\omega'(0))^2-\sigma^2)^{1/2}}\right) & \text{if } \sigma < 2\phi(0)\frac{\omega(0)}{\omega'(0)}, \end{cases}$$

*where $\boldsymbol{m}(\theta_{\epsilon^*})$ does not depend on the choice of $\theta_{\epsilon^*}$ among those that attain the modulus of continuity at $\epsilon^*$, and $\|\boldsymbol{m}(\theta_{\epsilon^*})\|=\epsilon^*$. The minimax risk is given by $\mathcal{R}(\sigma;\Theta)=\omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$.*

*Proof.* See Appendix A.3. □

The minimax regret rule takes different forms for the case where $\sigma \geq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$ and for the case where $\sigma < 2\phi(0)\frac{\omega(0)}{\omega'(0)}$. If $\sigma \geq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$, the minimax regret rule is nonrandomized, making a choice according to the sign of a weighted sum of the sample $\boldsymbol{Y}$. If $\sigma < 2\phi(0)\frac{\omega(0)}{\omega'(0)}$, the minimax regret rule is randomized, assigning a positive probability both to policies 1 and 0.

### 3.2.1 When Randomize?

The condition $\sigma \geq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$ determines whether the minimax regret rule is randomized or not. This condition is related to the strength of the restrictions imposed on the

[18] Assumption 1(i) is used only for obtaining the minimax regret rule for the case where $\sigma \leq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$. Without this assumption, the other statements in Theorem 1 hold.

[19] Lemma A.13 in Appendix A.3.2 shows that $\omega'(0)>0$ under the conditions of Theorem 1.

parameter space $\Theta$. To see this, note first that $\omega(0) = \sup\{L(\theta) : \boldsymbol{m}(\theta) = \mathbf{0}, \theta \in \Theta\}$ by definition. Since $L$ and $\boldsymbol{m}$ are linear and $\Theta$ is convex and centrosymmetric, it is shown that the closure of the identified set of $L(\theta)$ when $\boldsymbol{m}(\theta) = \mathbf{0}$ is given by $\mathrm{cl}(\{L(\theta) : \boldsymbol{m}(\theta) = \mathbf{0}, \theta \in \Theta\}) = [-\omega(0), \omega(0)]$.[20] We can thus interpret $\omega(0)$ as half of the length of the identified set of $L(\theta)$ when $\boldsymbol{m}(\theta) = \mathbf{0}$.

If $L(\theta)$ is identified, the length of the identified set is zero, which means that $\omega(0) = 0$. Since $\sigma \geq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$, the minimax regret rule is always nonrandomized for cases where $L(\theta)$ is identified. In the example of Section 2.3.1, $L(\theta)$ is identified if, for example, we specify a polynomial function for $f$.

On the other hand, if $L(\theta)$ is not identified, the length of the identified set is nonzero, which means that $\omega(0) > 0$. If the identified set is small relative to $\sigma$ (holding $\omega'(0)$ fixed), the condition $\sigma \geq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$ holds, and the minimax regret rule is nonrandomized. If the identified set is large relative to $\sigma$, the minimax regret rule is randomized. In Section 4, I show how this condition translates into one regarding the Lipschitz constant $C$ in the example of Section 2.3.1.

Simple calculations show that the randomized minimax regret rule is equivalent to $\delta^*(\boldsymbol{Y}) = \mathbb{P}\left((\boldsymbol{w}^*)'\boldsymbol{Y} + \xi \geq 0 | \boldsymbol{Y}\right)$, where $\xi | \boldsymbol{Y} \sim \mathcal{N}(0, (2\phi(0)\omega(0)/\omega'(0))^2 - \sigma^2)$. This rule is obtained through the following two-step procedure. We first add a noise $\xi$ to $(\boldsymbol{w}^*)'\boldsymbol{Y}$. This addition artificially increases the standard deviation of $(\boldsymbol{w}^*)'\boldsymbol{Y}$ from $\sigma$ to $2\phi(0)\frac{\omega(0)}{\omega'(0)}$, which is the threshold at which we switch from a nonrandomized rule to a randomized rule. We then make a decision according to the sign of $(\boldsymbol{w}^*)'\boldsymbol{Y} + \xi$. The larger $\omega(0)$ is, the larger the variance of $\xi$ is and the more dependent the choice is on the noise. As a result, given any realization of $\boldsymbol{Y}$, the probabilities of choosing policy 1 and policy 0 approach 1/2 as $\omega(0)$ increases, which suggests that the decisions become more mixed if we impose weaker restrictions on $\Theta$.

For understanding why the policy maker should randomize their decisions when $\omega(0)$ is large relative to $\sigma$, it is useful to consider the problem of finding worst-case parameter values for a generic decision rule $\delta$. The worst-case regret is attained at the parameter values that optimally trade off the potential welfare loss and the probability of incurring loss (i.e., $L(\theta)$ and $1 - \mathbb{E}_\theta[\delta(\boldsymbol{Y})]$ when $L(\theta) \geq 0$). Suppose that $\omega(0)$ is

[20] Since $L$ and $\boldsymbol{m}$ are linear and $\Theta$ is centrosymmetric, $-\omega(0) = \inf\{L(\theta) : \boldsymbol{m}(\theta) = \mathbf{0}, \theta \in \Theta\}$. Moreover, for any $\alpha \in (-\omega(0), \omega(0))$, we can find $\theta \in \Theta$ such that $L(\theta) = \alpha$ and $\boldsymbol{m}(\theta) = \mathbf{0}$ by the linearity of $L$ and $\boldsymbol{m}$ and the convexity of $\Theta$.

large relative to $\sigma$ and that the policy maker uses a nonrandomized rule. Since $\sigma$ is small, $\boldsymbol{Y}$ does not vary much across repeated samples, which makes the policy maker's choice based on the nonrandomized rule too predictable. By exploiting it, it is easy to find a value of $\theta$ under which the policy maker chooses the inferior policy with a high probability. If $\omega(0)$ is large enough, such choice of $\theta$ is not likely associated with a small welfare loss, leading to a large expected welfare loss of the decision rule. The policy maker can avoid this by randomizing their decisions; randomization makes their choice less predictable and protects against the exploitation of predictable choices. A similar discussion is provided by Stoye (2012).

### 3.3 Relation to Existing Results

Theorem 1 contains Proposition 7(iii) of Stoye (2012) as a special case. He considers a simple setup with a specific form of partial identification, given in Section 2.3.2.[21] He shows that the following rule is minimax regret if $b < 1$:[22]

$$\delta^*(Y) = \begin{cases} \mathbf{1}\{Y \geq 0\} & \text{if } \sigma \geq 2\phi(0)\frac{b}{a}, \\ \Phi\left(\dfrac{Y}{((2\phi(0)b/a)^2 - \sigma^2)^{1/2}}\right) & \text{if } \sigma < 2\phi(0)\frac{b}{a}. \end{cases}$$

The condition $\sigma \geq 2\phi(0)\frac{b}{a}$ is equivalent to $\sigma \geq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$ since $\omega(\epsilon) = \sup\{\theta_2 : \theta_1 \in [-\epsilon, \epsilon], (\theta_1, \theta_2)' \in \Theta\} = \min\{a\epsilon + b, 1\}$. Note that the nonrandomized minimax regret rule $\delta^*(Y) = \mathbf{1}\{Y \geq 0\}$ is insensitive to any of $\sigma$, $a$, and $b$ as long as $\sigma \geq 2\phi(0)\frac{b}{a}$.

Theorem 1 confirms that a minimax regret rule can be both nonrandomized and randomized even in much more general setups. At the same time, Theorem 1 suggests that the nonrandomized minimax regret rule $\mathbf{1}\{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y} \geq 0\}$ may be sensitive to $\sigma$ and $\Theta$, since $\boldsymbol{m}(\theta_{\epsilon^*})$ depends on them. Therefore, the robustness of the nonrandomized minimax regret rule to the error variance and to the parameter space is not

[21] Assumption 1(i)(d) and (ii) (the differentiability of $\rho$ and $\omega$) do not hold for some choices of constants $a$ and $b$. For such cases, the minimax regret rule is derived by Theorems A.1 and A.2 in Appendix A.3, which do not require the differentiability.

[22] Stoye (2012) also covers the case where $b \geq 1$. In this case, Assumption 1(i)(c) does not hold; $\theta^* = (0, 1)'$ attains the modulus of continuity at $\epsilon = 0$, but there exists no $\theta \in \Theta$ such that $L(\theta) = \theta_2 \neq 0$ and $\theta^* + c\theta \in \Theta$ for any $c$ in a neighborhood of zero. Theorem A.2 in Appendix A.3.2 covers this case.

a general property.

In a special case of this paper's setup, Ishihara and Kitagawa (2021) characterize the minimax regret rule within the class of decision rules of the form $\delta(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{w}'\boldsymbol{Y} \geq 0\}$, where $\boldsymbol{w} \in \mathbb{R}^n$. Theorem 1 shows that this restricted class contains the minimax regret rule when $\sigma \geq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$ and may not when $\sigma < 2\phi(0)\frac{\omega(0)}{\omega'(0)}$.

## 3.4 Comparison with a Plug-in Rule Based on a Linear Minimax Mean Squared Error Estimator

Here, I compare the nonrandomized minimax regret rule with a plug-in rule based on a linear minimax mean squared error (MSE) estimator.[23] To define the alternative rule, let $\hat{L}_{\text{MSE}}(\boldsymbol{Y}) = \boldsymbol{w}'_{\text{MSE}}\boldsymbol{Y}$ be a *linear minimax MSE estimator* of $L(\theta)$, where

$$\boldsymbol{w}_{\text{MSE}} \in \arg\min_{\boldsymbol{w}\in\mathbb{R}^n} \sup_{\theta\in\Theta} \mathbb{E}_\theta[(\boldsymbol{w}'\boldsymbol{Y} - L(\theta))^2].$$

$\hat{L}_{\text{MSE}}(\boldsymbol{Y})$ has the smallest maximum MSE within the class of linear estimators. I define the *plug-in MSE rule* as $\delta_{\text{MSE}}(\boldsymbol{Y}) = \mathbf{1}\{\hat{L}_{\text{MSE}}(\boldsymbol{Y}) \geq 0\}$, which makes a choice according to the sign of the linear minimax MSE estimator of $L(\theta)$.

Donoho (1994) characterizes $\hat{L}_{\text{MSE}}(\boldsymbol{Y})$ using the modulus of continuity. Let $\epsilon_{\text{MSE}} > 0$ solve $\frac{\epsilon^2}{\epsilon^2+\sigma^2} = \frac{\omega'(\epsilon)\epsilon}{\omega(\epsilon)}$. The linear minimax MSE estimator is then given by $\hat{L}_{\text{MSE}}(\boldsymbol{Y}) = \frac{\omega'(\epsilon_{\text{MSE}})}{\epsilon_{\text{MSE}}}\boldsymbol{m}(\theta_{\epsilon_{\text{MSE}}})'\boldsymbol{Y}$, where $\theta_{\epsilon_{\text{MSE}}}$ attains the modulus of continuity at $\epsilon_{\text{MSE}}$ with $\|\boldsymbol{m}(\theta_{\epsilon_{\text{MSE}}})\| = \epsilon_{\text{MSE}}$. The plug-in MSE rule is $\delta_{\text{MSE}}(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\theta_{\epsilon_{\text{MSE}}})'\boldsymbol{Y} \geq 0\}$.

Recall that the minimax regret rule is $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y} \geq 0\}$ if $\sigma > 2\phi(0)\frac{\omega(0)}{\omega'(0)}$, where $\epsilon^*$ solves $\max_{\epsilon\in[0,a^*\sigma]} \omega(\epsilon)\Phi(-\epsilon/\sigma)$.

**Proposition 1** (Comparison with Plug-in MSE Rule)**.** *Suppose that $\omega(\cdot)$ is differentiable with $\omega'(0) > 0$, and let $\epsilon_{\text{MSE}} > 0$ solve $\frac{\epsilon^2}{\epsilon^2+\sigma^2} = \frac{\omega'(\epsilon)\epsilon}{\omega(\epsilon)}$ and $\epsilon^*$ solve $\max_{\epsilon\in[0,a^*\sigma]} \omega(\epsilon)\Phi(-\epsilon/\sigma)$. Then, $\epsilon^* < \epsilon_{\text{MSE}}$.*

*Proof.* See Appendix A.5. □

I provide an implication of Proposition 1 through the following result from Low (1995): the optimal bias-variance frontier in the estimation of $L(\theta)$ can be traced out

[23] In Appendix C.2, I also compare the minimax regret rule with a hypothesis testing rule that chooses policy 1 if a hypothesis that supports policy 0 is rejected.

by a class of linear estimators $\{\hat{L}_\epsilon(\boldsymbol{Y})\}_{\epsilon>0}$ of the form $\hat{L}_\epsilon(\boldsymbol{Y}) = \frac{\omega'(\epsilon)}{\epsilon}\boldsymbol{m}(\theta_\epsilon)'\boldsymbol{Y}$. Here, $\theta_\epsilon$ attains the modulus of continuity at $\epsilon$ with $\|\boldsymbol{m}(\theta_\epsilon)\| = \epsilon$. Specifically, for each $\epsilon > 0$, $\hat{L}_\epsilon(\boldsymbol{Y})$ minimizes the maximum bias among all linear estimators with variance bounded by $\mathrm{Var}(\hat{L}_\epsilon(\boldsymbol{Y})) = (\sigma\omega'(\epsilon))^2$:

$$\frac{\omega'(\epsilon)}{\epsilon}\boldsymbol{m}(\theta_\epsilon) \in \arg\min_{\boldsymbol{w}\in\mathbb{R}^n} \overline{\mathrm{Bias}}_\Theta(\boldsymbol{w}'\boldsymbol{Y}) \;\; s.t. \;\; \mathrm{Var}(\boldsymbol{w}'\boldsymbol{Y}) \le (\sigma\omega'(\epsilon))^2,$$

where $\overline{\mathrm{Bias}}_\Theta(\boldsymbol{w}'\boldsymbol{Y}) = \sup_{\theta\in\Theta}\mathbb{E}_\theta[\boldsymbol{w}'\boldsymbol{Y} - L(\theta)]$ is the maximum bias of $\boldsymbol{w}'\boldsymbol{Y}$ over $\Theta$. As $\epsilon$ increases, the maximum bias $\overline{\mathrm{Bias}}_\Theta(\hat{L}_\epsilon(\boldsymbol{Y}))$ increases and the variance $\mathrm{Var}(\hat{L}_\epsilon(\boldsymbol{Y})) = (\sigma\omega'(\epsilon))^2$ decreases. $\epsilon_{\mathrm{MSE}}$ minimizes the worst-case MSE $\sup_{\theta\in\Theta}\mathbb{E}_\theta[(\hat{L}_\epsilon(\boldsymbol{Y}) - L(\theta))^2] = \overline{\mathrm{Bias}}_\Theta(\hat{L}_\epsilon(\boldsymbol{Y}))^2 + \mathrm{Var}(\hat{L}_\epsilon(\boldsymbol{Y}))$.

Since $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y} \ge 0\} = \mathbf{1}\{\hat{L}_{\epsilon^*}(\boldsymbol{Y}) \ge 0\}$, the minimax regret rule $\delta^*(\boldsymbol{Y})$ can be viewed as a rule that makes a choice according to the sign of the linear estimator $\hat{L}_{\epsilon^*}(\boldsymbol{Y})$. Proposition 1 implies that the corresponding linear estimator $\hat{L}_{\epsilon^*}(\boldsymbol{Y})$ places more importance on the bias than on the variance compared with the linear minimax MSE estimator $\hat{L}_{\mathrm{MSE}}(\boldsymbol{Y})$. This result suggests that the plug-in MSE rule is not necessarily optimal under the minimax regret criterion.

# 4 Application to Eligibility Cutoff Choice

Theorem 1 provides a procedure to compute a minimax regret rule for the example in Section 2.3.1. In this section, I provide the formula of the minimax regret rule and discuss how the rule depends on the Lipschitz constant $C$ and the new cutoff $c_1$.

I first normalize $\boldsymbol{Y}$ and $\boldsymbol{m}(\cdot)$ by left multiplying them by $\boldsymbol{\Sigma}^{-1/2}$ so that the variance-covariance matrix of the sample is the identity matrix: $\tilde{\boldsymbol{Y}} \sim \mathcal{N}(\tilde{\boldsymbol{m}}(f), \boldsymbol{I}_n)$, where $\tilde{\boldsymbol{Y}} = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{Y} = (Y_1/\sigma(x_1, d_1), ..., Y_n/\sigma(x_n, d_n))'$, and $\tilde{\boldsymbol{m}}(f) = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{m}(f) = (f(x_1, d_1)/\sigma(x_1, d_1), ..., f(x_n, d_n)/\sigma(x_n, d_n))'$. For illustration, I focus on the Lipschitz class $\mathcal{F} = \mathcal{F}_{\mathrm{Lip}}(C)$ and suppose that the welfare is the sample average of the expected outcome. The welfare difference is given by $L(f) = \frac{1}{n}\sum_{i=1}^n \mathbf{1}\{c_1 \le x_i < c_0\}[f(x_i, 1) - f(x_i, 0)]$.

Below, I first verify that Assumption 1 holds and then apply Theorem 1 to derive a minimax regret rule.

## 4.1 Verifying Assumption 1

To verify Assumption 1, I derive a closed-form expression for a value of $f$ that attains the modulus of continuity at $\epsilon$ when $\epsilon$ is sufficiently small. The modulus of continuity $\omega(\epsilon; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C))$ is computed by solving

$$\sup_{f \in \mathcal{F}_{\text{Lip}}(C)} \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}\{c_1 \leq x < c_0\}[f(x_i, 1) - f(x_i, 0)] \quad s.t. \quad \sum_{i=1}^{n} \frac{f(x_i, d_i)^2}{\sigma^2(x_i, d_i)} \leq \epsilon^2. \tag{3}$$

The unknown parameter $f$ is infinite dimensional, but the objective and the norm constraint $\sum_{i=1}^{n} \frac{f(x_i,d_i)^2}{\sigma^2(x_i,d_i)} \leq \epsilon^2$ depend on $f$ only through its values at $(x_1, 0), ..., (x_n, 0), (x_1, 1), ..., (x_n, 1)$. This optimization problem can be reduced to the following convex optimization problem with $2n$ unknowns and $1 + n(n-1)$ inequality constraints by a slight modification of Theorem 2.2 in Armstrong and Kolesár (2021):

$$\max_{(f(x_i,0), f(x_i,1))_{i=1,...,n} \in \mathbb{R}^{2n}} \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}\{c_1 \leq x_i < c_0\}[f(x_i, 1) - f(x_i, 0)] \tag{4}$$

$$s.t. \quad \sum_{i=1}^{n} \frac{f(x_i, d_i)^2}{\sigma^2(x_i, d_i)} \leq \epsilon^2, \;\; f(x_i, d) - f(x_j, d) \leq C|x_i - x_j|, \quad d \in \{0, 1\}, i, j \in \{1, ..., n\}.$$

A solution to (4) exists since the objective function is continuous and the set of the vectors of $2n$ unknowns that satisfy the constraints is closed and bounded. Once we find a solution $(f(x_i, 0), f(x_i, 1))_{i=1,...,n}$, we can always find a function $f \in \mathcal{F}_{\text{Lip}}(C)$ that interpolates the points $(x_i, f(x_i, 0)), (x_i, f(x_i, 1))$, $i = 1, ..., n$ (Beliakov, 2006, Theorem 4), which is a solution to the original problem (3).

Now, I show that the problem (4) has a closed-form solution for any sufficiently small $\epsilon \geq 0$. The derivation utilizes the specific treatment assignment rule in the RD, namely $d_i = \mathbf{1}\{x_i \geq c_0\}$. Let $\tilde{n} = \sum_{i=1}^{n} \mathbf{1}\{c_1 \leq x_i < c_0\}$ denote the number of units whose treatment status would be changed if the cutoff were changed. Additionally, let $x_{+,\min} = \min\{x_i : x_i \geq c_0\}$ be the value of $x$ of the treated unit closest to the original cutoff $c_0$, and let $\sigma^2_{+,\min} = \sigma^2(x_{+,\min}, 1)$. To simplify the exposition, I assume that $x_i \neq x_j$ for any $i \neq j$, $i, j = 1, ..., n$, in what follows.[24]

[24] It is possible to obtain a closed-form solution without this assumption at the cost of making the presentation more complex.

**Proposition 2** (Solution to Modulus Problem for Cutoff Choice)**.** *Suppose that* $d_i = \mathbf{1}\{x_i \geq c_0\}$ *for all* $i = 1, ..., n$ *and that* $x_i \neq x_j$ *for any* $i \neq j$, $i, j = 1, ..., n$. *Then, there exists* $\bar{\epsilon} > 0$ *such that for any* $\epsilon \in [0, \bar{\epsilon}]$, *one solution to (4) is given by*

$$f_\epsilon(x_i, 0) = \begin{cases} 0 & \text{if } x_i < c_1 \text{ or } x_i \geq c_0, \\ -\dfrac{\sigma^2(x_i, 0)\epsilon}{\bar{\sigma}} & \text{if } c_1 \leq x_i < c_0, \end{cases}$$

$$f_\epsilon(x_i, 1) = \begin{cases} 0 & \text{if } x_i > x_{+,\min}, \\ C(x_{+,\min} - x_i) + \dfrac{\tilde{n}\sigma^2_{+,\min}\epsilon}{\bar{\sigma}} & \text{if } x_i \leq x_{+,\min}, \end{cases}$$

*and the modulus of continuity is given by*

$$\omega(\epsilon; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\mathrm{Lip}}(C)) = C\frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{c_1 \leq x_i < c_0\}[x_{+,\min} - x_i] + \frac{\bar{\sigma}\epsilon}{n},$$

*where* $\bar{\sigma} = (\tilde{n}^2\sigma^2_{+,\min} + \sum_{i:c_1 \leq x_i < c_0} \sigma^2(x_i, 0))^{1/2}$.

*Proof.* See Appendix A.6. □

A brief explanation of this result is as follows. Since $d_i = \mathbf{1}\{x_i \geq c_0\}$, the norm constraint of (4) does not depend on $f(x_i, 1)$ for $i$ with $x_i < c_0$. The upper bound on $f(x_i, 1)$ for such unit $i$ is $C(x_{+,\min} - x_i) + f(x_{+,\min}, 1)$ under the Lipschitz constraint; $(x_i, f(x_i, 1))$ lies on the straight line with slope $-C$ that goes through $(x_{+,\min}, f(x_{+,\min}, 1))$. Given a value of $f(x_{+,\min}, 1)$, the objective of (4) then becomes $C\frac{1}{n}\sum_{i:c_1 \leq x_i < c_0}[x_{+,\min} - x_i] + \frac{\tilde{n}}{n}f(x_{+,\min}, 1) - \frac{1}{n}\sum_{i:c_1 \leq x_i < c_0} f(x_i, 0)$, which is a constant plus a weighted sum of $f(x_{+,\min}, 1)$ and $f(x_i, 0)$ for $i$ with $c_1 \leq x_i < c_0$. By maximizing this under the norm constraint, we obtain the display of $f_\epsilon$ in Proposition 2, which turns out to satisfy the Lipschitz constraint for any sufficiently small $\epsilon$.

Assumption 1(i)(a) immediately follows from Proposition 2. Moreover, for any

$\epsilon \in [0, \bar{\epsilon}]$, $\frac{\tilde{\boldsymbol{m}}(f_\epsilon)}{\|\tilde{\boldsymbol{m}}(f_\epsilon)\|} = \frac{\tilde{\boldsymbol{m}}(f_\epsilon)}{\epsilon}$ is constant and equal to $\boldsymbol{w}^* = (w_1^*, ..., w_n^*)'$, where

$$w_i^* = \begin{cases} 0 & \text{if } x_i < c_1 \text{ or } x_i > x_{+,\min}, \\ -\dfrac{\sigma(x_i, 0)}{\bar{\sigma}} & \text{if } c_1 \leq x_i < c_0, \\ \dfrac{\tilde{n}\sigma_{+,\min}}{\bar{\sigma}} & \text{if } x_i = x_{+,\min}. \end{cases} \tag{5}$$

Therefore, Assumption 1(i)(b) holds. Furthermore, Assumption 1(i)(c) straightforwardly holds with $\iota(x, d) = d$ for all $x \in \mathbb{R}$. Assumption 1(i)(d) and (ii) are shown to hold in Appendix B.1.

## 4.2 Minimax Regret Rule

Now, I apply Theorem 1 to derive a minimax regret rule. By Proposition 2, we obtain closed-form expressions for $\omega(0; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C))$ and $\omega'(0; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C))$:

$$\omega(0; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C)) = \frac{C}{n} \sum_{i=1}^{n} \mathbf{1}\{c_1 \leq x_i < c_0\}[x_{+,\min} - x_i], \ \omega'(0; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C)) = \frac{\bar{\sigma}}{n}.$$

Let

$$\sigma^* := 2\phi(0) \frac{\omega(0; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C))}{\omega'(0; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C))} = 2\phi(0) C \sum_{i=1}^{n} \mathbf{1}\{c_1 \leq x_i < c_0\}[x_{+,\min} - x_i]/\bar{\sigma}. \tag{6}$$

Recall that $\tilde{\boldsymbol{Y}} = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{Y} = (Y_1/\sigma(x_1, d_1), ..., Y_n/\sigma(x_n, d_n))'$, $\tilde{\boldsymbol{m}}(f) = \boldsymbol{\Sigma}^{-1/2}\boldsymbol{m}(f) = (f(x_1, d_1)/\sigma(x_1, d_1), ..., f(x_n, d_n)/\sigma(x_n, d_n))'$, and the variance of $\tilde{\boldsymbol{Y}}$ is $\boldsymbol{I}_n$. Let $\epsilon^* \in \arg\max_{\epsilon \in [0, a^*]} \omega(\epsilon; L, \tilde{\boldsymbol{m}}, \mathcal{F}_{\text{Lip}}(C))\Phi(-\epsilon)$ and $(f_{\epsilon^*}(x_i, 0), f_{\epsilon^*}(x_i, 1))_{i=1,...,n}$ solve the problem (4) for $\epsilon = \epsilon^*$. By Theorem 1, the following rule is minimax regret:

$$\delta^*(\boldsymbol{Y}) = \begin{cases} \mathbf{1}\left\{\sum_{i=1}^{n} f_{\epsilon^*}(x_i, d_i) Y_i/\sigma^2(x_i, d_i) \geq 0\right\} & \text{if } 1 > \sigma^*, \\ \mathbf{1}\left\{\sum_{i=1}^{n} w_i^* Y_i/\sigma(x_i, d_i) \geq 0\right\} & \text{if } 1 = \sigma^*, \\ \Phi\left(\dfrac{\sum_{i=1}^{n} w_i^* Y_i/\sigma(x_i, d_i)}{((\sigma^*)^2 - 1)^{1/2}}\right) & \text{if } 1 < \sigma^*. \end{cases} \tag{7}$$

The minimax regret rule makes a decision based on a weighted sum of $Y_1, ..., Y_n$.

To understand how the rule differs across different values of the Lipschitz constant $C$, suppose first that the magnitude of $C$ is moderate so that $\sigma^*$ is marginally smaller than 1. In this case, $\epsilon^*$ tends to be sufficiently small, which implies that $\frac{f_{\epsilon^*}(x_i,d_i)/\sigma(x_i,d_i)}{\epsilon^*} = w_i^*$ by Proposition 2. The minimax regret rule is given by

$$\begin{aligned}\delta^*(\boldsymbol{Y}) &= \mathbf{1}\left\{\sum_{i=1}^n f_{\epsilon^*}(x_i,d_i)Y_i/\sigma^2(x_i,d_i) \geq 0\right\}\\ &= \mathbf{1}\left\{\sum_{i=1}^n w_i^* Y_i/\sigma(x_i,d_i) \geq 0\right\} = \mathbf{1}\left\{Y_{+,\min} - \frac{1}{\tilde{n}}\sum_{i:c_1\leq x_i<c_0} Y_i \geq 0\right\},\end{aligned}$$

where $Y_{+,\min} = Y_i$ for $i$ with $x_i = x_{+,\min}$. $Y_{+,\min} - \frac{1}{\tilde{n}}\sum_{i:c_1\leq x_i<c_0} Y_i$ is the difference between the outcome of the treated unit closest to the status quo cutoff $c_0$ and the mean outcome across the untreated units between the two cutoffs $c_0$ and $c_1$. This difference can be interpreted as an estimator of the effect of the cutoff change. The outcomes of the other units are not used to construct the estimator. The minimax regret rule makes a decision according to its sign.

On the other hand, if the Lipschitz constant $C$ is small enough so that $\sigma^*$ is substantially smaller than 1, nonzero weights may be assigned to some of the other units, that is, $f_{\epsilon^*}(x_i,d_i)$ may be nonzero for some of the units with $x_i < c_1$ or $x_i > x_{+,\min}$. If the Lipschitz constant $C$ is large enough so that $\sigma^* > 1$, the minimax regret rule is a randomized rule based on $Y_{+,\min} - \frac{1}{\tilde{n}}\sum_{i:c_1\leq x_i<c_0} Y_i$.

Whether the minimax regret rule is randomized or not depends not only on the Lipschitz constant $C$ but also on the cutoffs $c_0$ and $c_1$ and $\bar{\sigma} = (\tilde{n}^2\sigma_{+,\min}^2 + \sum_{i:c_1\leq x_i<c_0}\sigma^2(x_i,0))^{1/2}$. To investigate their relationships, suppose that $\sigma^2(x_i,d_i) = \sigma^2$ for all $i$ for some $\sigma^2 > 0$ for simplicity. In this situation, $\bar{\sigma} = (\tilde{n}^2+\tilde{n})^{1/2}\sigma$, and

$$\sigma^* = \frac{2\phi(0)C\frac{1}{\tilde{n}}\sum_{i=1}^n \mathbf{1}\{c_1 \leq x_i < c_0\}[x_{+,\min} - x_i]}{(1+1/\tilde{n})^{1/2}\,\sigma}.$$

$\sigma^*$ is nonincreasing in $c_1$ since $\frac{1}{\tilde{n}}\sum_{i=1}^n \mathbf{1}\{c_1 \leq x_i < c_0\}[x_{+,\min} - x_i]$ and $\tilde{n}$ are nonincreasing in $c_1$.[25] Furthermore, $\sigma^*$ is decreasing in $\sigma$. Therefore, the minimax regret rule is nonrandomized when $c_1$ is large (i.e., when the cutoff change $c_0 - c_1$ is small)

[25]Whether $\sigma^*$ is increasing in $c_0$ or not depends on the empirical distribution of $x_i$.

or $\sigma$ is large. The minimax regret rule is randomized otherwise.

### 4.3 Practical Implementation

Here, I summarize the procedure for computing the minimax regret rule and discuss practical issues. Given the conditional variances $\sigma^2(x_i, d_i)$, $i = 1, ..., n$ and the Lipschitz constant $C$, the minimax regret rule is computed as follows.

1. Compute $\sigma^*$ using the closed-form expression (6).
2. If $1 > \sigma^*$, find $\epsilon^* \in \arg\max_{\epsilon \in [0,a^*]} \omega(\epsilon)\Phi(-\epsilon)$ and compute $f_{\epsilon^*}$ that attains the modulus of continuity at $\epsilon^*$. For each $\epsilon \geq 0$, $\omega(\epsilon)$ is computed by solving the convex optimization problem (4).[26] An efficient method for computing $\epsilon^*$ is provided in Appendix B.2.
3. If $1 \leq \sigma^*$, compute $\boldsymbol{w}^*$ using the closed-form expression (5).
4. Construct the decision rule according to (7).

In practice, the conditional variance $\sigma^2(x_i, d_i)$ is unknown. I suggest using a consistent estimator in place of the true $\sigma^2(x_i, d_i)$. The conditional variance can be estimated, for example, by applying a local linear regression to the squared residuals (Fan and Yao, 1998) or by the nearest-neighbor variance estimator (Abadie and Imbens, 2006). In the case where unit $i$ represents a group of individuals and where $Y_i$ is the sample mean outcome within group $i$, it is natural to use the conventional standard error of the sample mean as $\sigma(x_i, d_i)$.

Implementation of the minimax regret rule requires choosing the Lipschitz constant $C$. In principle, it is not possible to choose the Lipschitz constant $C$ that applies to both sides of the cutoff $c_0$ in a data-driven way since we only observe outcomes either under treatment or under no treatment on each side. It is, however, possible to estimate a lower bound on $C$ by using the following observation: if $f \in \mathcal{F}_{\text{Lip}}(C)$ is differentiable, a lower bound on $C$ is given by $\max\left\{\max_{\tilde{x} \geq c_0}\left|\frac{\partial f(\tilde{x},1)}{\partial x}\right|, \max_{\tilde{x} < c_0}\left|\frac{\partial f(\tilde{x},0)}{\partial x}\right|\right\}$

[26] In the empirical application in Section 5, I solve the convex optimization problem using CVXPY, a Python-embedded modeling language for convex optimization problems (Diamond and Boyd, 2016; Agrawal, Verschueren, Diamond and Boyd, 2018).

since $\left|\frac{\partial f(\tilde{x},d)}{\partial x}\right| \leq C$ for all $\tilde{x}$ and $d$. To estimate a lower bound, we could estimate the derivatives $\frac{\partial f(\tilde{x},1)}{\partial x}$ for $\tilde{x} \geq c_0$ and $\frac{\partial f(\tilde{x},0)}{\partial x}$ for $\tilde{x} < c_0$ by a local polynomial regression and then take the maximum of their absolute values over a plausible, relevant interval.[27] In practice, I recommend considering a range of plausible choices of $C$, including the estimated lower bound, to conduct a sensitivity analysis. I implement this approach for my empirical application in Section 5.

# 5 Empirical Policy Application

I now illustrate my approach in an empirical application to the BRIGHT program in Burkina Faso. I consider the hypothetical problem of whether or not to expand the program and empirically compare the performance of the minimax regret rule with alternative decision rules.

## 5.1 Background and Data

The goal of the BRIGHT program was to improve children's, especially girls' educational outcomes in rural villages by constructing well-resourced village-based schools. The program was funded by the Millennium Challenge Corporation, a U.S. government agency, and implemented by a consortium of non-governmental organizations.

The program constructed primary schools with three classrooms for grades 1 to 3 in 132 villages from 47 departments during the period from 2005 to 2008. The Ministry of Education determined the villages where schools would be built through the following process. First, 293 villages were nominated based on low school enrollment rates. Second, the Ministry administered a survey in each village and assigned each village a score using a set formula. The formula attached a large weight to the estimated number of children to be served from the nominated and neighboring villages, giving additional weight to girls. The Ministry then ranked villages within each department and selected the top half of the villages to receive a school. For further details on

[27] Simply taking the maximum of the estimated derivatives could raise a concern of upward bias. One could use the method for intersection bounds developed by Chernozhukov, Lee and Rosen (2013) to address it.

Figure 1: Distribution of Relative Score

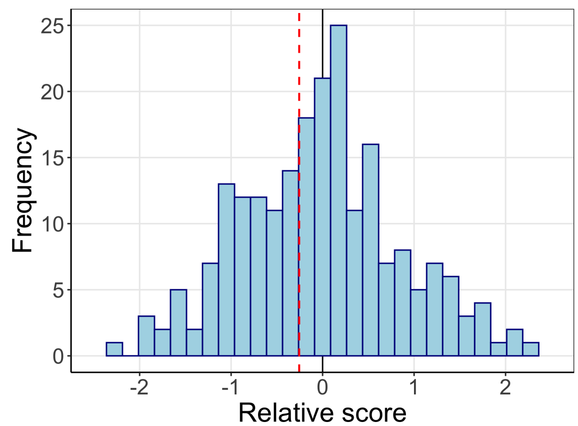


*Notes*: This figure shows the histogram of the relative score of villages on the interval $[-2.5, 2.5]$. The vertical dashed line indicates the new cutoff $-0.256$, which corresponds to the hypothetical policy of constructing schools in previously ineligible villages whose relative scores are in the top 20%. The villages with zero observed enrollment rates are excluded.

the BRIGHT program and allocation process, see Levy, Sloan, Linden and Kazianga (2009) and Kazianga *et al.* (2013).

Since the school allocation was determined at department level, the cutoff score for the program eligibility was different across departments. Following Kazianga *et al.* (2013), I define the *relative score* as the score for each village minus the cutoff score for the department that the village belongs to. As a result, a village is eligible for the program when the relative score is larger than zero. Kazianga *et al.* (2013) use the relative score as a running variable and evaluate the causal effect of the program on educational outcomes using an RD design. Figure 1 reports the distribution of the relative score.

I use the replication data for Kazianga *et al.* (2013)'s results (Kazianga, Levy, Linden and Sloan, 2019) and consider whether we should expand the program or not. The dataset contains survey results about 30 households from 287 nominated villages, yielding a total sample of 23,282 children between the ages of 5 and 12. The survey was conducted in 2008, namely 2.5 years after the start of the program. Table 1 reports summary statistics about child educational outcomes and characteristics. Children in eligible villages are more likely to attend school, achieve higher test scores,

Table 1: Child Educational Outcomes and Characteristics

| | All | Eligible villages | Ineligible villages |
|---|---|---|---|
| | (1) | (2) | (3) |
| Panel A. Educational outcomes (child-level means) | | | |
| Enrollment | 0.366 | 0.494 | 0.259 |
| Normalized total test scores | 0.000 | 0.248 | −0.209 |
| Highest grade child has achieved | 0.876 | 1.132 | 0.636 |
| Panel B. Child and household characteristics (child-level means) | | | |
| Child's age | 8.121 | 8.174 | 8.071 |
| Child is female | 0.503 | 0.476 | 0.525 |
| Head's age | 47.653 | 47.387 | 47.904 |
| Head years of schooling | 0.156 | 0.198 | 0.117 |
| Number of members | 10.812 | 10.815 | 10.808 |
| Number of children | 5.971 | 6.098 | 5.850 |
| Muslim | 0.587 | 0.576 | 0.597 |
| Basic roofing | 0.516 | 0.534 | 0.500 |
| Number of motorbikes | 0.299 | 0.319 | 0.279 |
| Number of phones | 0.185 | 0.199 | 0.172 |
| Total number of children | 23,282 | 10,645 | 12,637 |
| Total number of villages | 287 | 136 | 151 |

*Notes*: This table reports child-level averages of educational outcomes and characteristics by program eligibility in the year 2008, namely 2.5 years after the start of the BRIGHT program. Panel A reports the educational outcomes' means. Panel B reports the means of child and household characteristics. Column (1) shows the means for children in all villages. Columns (2) and (3) show the means for children in villages selected for BRIGHT school and in unselected villages, respectively.

and complete a higher grade. Household heads in eligible villages completed slightly more years of schooling. Furthermore, households in eligible villages tend to have more assets such as basic roofing and motorbikes.

I consider school enrollment as the target outcome. Since the score and program eligibility are determined at village level, I use the village-level mean outcome, namely the enrollment rate for each village. This setting fits into the setup in Section 2.3.1, where $i$ represents a village, $Y_i$ is the observed enrollment rate of village $i$, $d_i$ is the program eligibility, and $x_i$ is the relative score. The original cutoff is $c_0 = 0$, that is, $d_i = \mathbf{1}\{x_i \geq 0\}$. The parameter is a function $f : \mathbb{R} \times \{0, 1\} \to \mathbb{R}$, where $f(x, d)$ represents the counterfactual probability of enrollment conditional on the relative score if the eligibility status were set to $d \in \{0, 1\}$. Since $Y_i$ is a village-level sample mean, it is plausible to assume that $Y_i$ is approximately normally distributed. I use the conventional standard error of the sample mean as the standard deviation of $Y_i$.[28]

## 5.2 Hypothetical Policy Choice Problem

I ask whether we should scale up the program and build BRIGHT schools in other villages. Specifically, I consider the following decision problem. The counterfactual policy is to build BRIGHT schools in previously ineligible villages whose relative scores are in the top 20%, which corresponds to lowering the cutoff from 0 to $-0.256$.[29] I use the average enrollment rate across villages as the welfare criterion, so that the welfare effect of this policy relative to the status quo is

$$L(f) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{-0.256 \le x_i < 0\}[f(x_i, 1) - f(x_i, 0)].$$

When deciding whether to implement the policy, it is important to consider the benefit relative to the cost. Kazianga *et al.* (2013) provide an estimate of the cost of constructing a BRIGHT school, which is \$4,758 per village.[30] To incorporate the cost into the decision problem, I suppose that the policy maker cares about the cost-effectiveness of this new policy relative to similar programs. Cost-effectiveness is defined as the ratio of the policy cost to the increase in the target outcome, namely the enrollment in the current context. I assume that it is optimal to implement the policy if its cost-effectiveness is smaller than \$83.77, which is the cost-effectiveness of a school construction program in Indonesia (Duflo, 2001; Kazianga *et al.*, 2013). Specifically, it is optimal to implement the policy if

$$\frac{\$4{,}758}{416 \cdot \frac{1}{\tilde{n}} \sum_{i=1}^{n} \mathbf{1}\{-0.256 \le x_i < 0\}[f(x_i, 1) - f(x_i, 0)]} \le \$83.77,$$

[28] The observed enrollment rate is zero in 21 out of 287 villages. I exclude these villages from the analysis since the standard error of $Y_i$ is zero.

[29] Figure D.1 in Appendix D reports the results when I use 10% and 30% instead of 20%. As predicted by the result in Section 4, the minimax regret rule switches from a nonrandomized rule to a randomized rule at a smaller Lipschitz constant $C$ when the fraction of the target villages is larger.

[30] I assume that the cost estimate is a known quantity and is constant across villages. It is, however, natural to think of the policy cost as unknown and heterogeneous across villages and to introduce the cost model on top of the outcome model. I leave this for future work.

where 416 is the number of children per village and $\tilde{n} = \sum_{i=1}^{n} \mathbf{1}\{-0.256 \leq x_i < 0\}$ is the number of villages that would receive a school under the new policy.[31] The denominator represents the increase in the average enrollment across villages that would receive a BRIGHT school under the new policy.

Simple calculations show that the above condition is equivalent to

$$\frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{-0.256 \leq x_i < 0\}[f(x_i, 1) - 0.137 - f(x_i, 0)] \geq 0.$$

My method can be used to consider this decision problem by setting the outcome to $Y_i - 0.137 d_i$, where 0.137 can be viewed as the policy cost measured in the unit of the enrollment rate. I present the results for this scenario with the cost of 0.137 as well as for the scenario where we ignore the policy cost.

I implement my method assuming that the counterfactual outcome function $f$ belongs to the Lipschitz class $\mathcal{F}_{\text{Lip}}(C)$. Since the relative score $x_i$ is computed based on several village-level characteristics, it is difficult to interpret and specify the Lipschitz constant $C$ using domain-specific knowledge. To obtain a reasonable range of $C$, I estimate a lower bound on $C$ using the method described in Section 4.3, which yields the lower bound estimate of 0.149.[32] I present the results for $C \in \{0.05, 0.1, ..., 0.95, 1\}$ and examine their sensitivity to the choice of $C$.

## 5.3 Results

Figure 2 plots $\delta^*(\boldsymbol{Y})$, the probability of choosing the new policy computed by the minimax regret rule, against the Lipschitz constant $C$. When $C < 0.6$, the minimax regret rule is nonrandomized. It chooses the new policy in the no-cost scenario and

[31] The cost per village and the cost-effectiveness of a school construction program in Indonesia are found in Tables A18 and A20, respectively, in Online Appendix of Kazianga *et al.* (2013). I compute the number of children per village by dividing the total enrollment by the enrollment rate reported in Table A17 in Online Appendix of Kazianga *et al.* (2013).

[32] I estimate $\frac{\partial f(x,0)}{\partial x}$ at $x \in \{-2.5, -2.45, ..., -0.05\}$ and $\frac{\partial f(x,1)}{\partial x}$ at $x \in \{0.05, 0.1, ..., 2.5\}$ by local quadratic regression and take the maximum of their absolute values. For local quadratic regression, I use the MSE-optimal bandwidth selection procedure by Calonico, Cattaneo and Farrell (2018), which can be implemented by R package "nprobust."

Figure 2: Optimal Decisions: Probability of Choosing New Policy

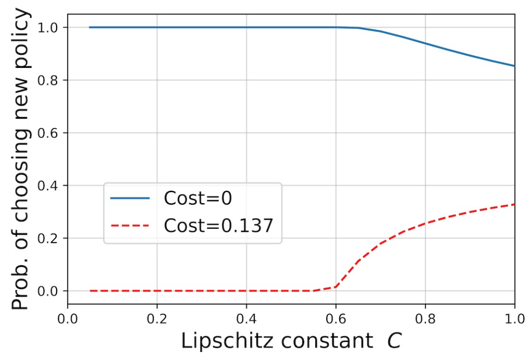


*Notes*: This figure shows the probability of choosing the new policy computed by the minimax regret rule. The new policy is to construct BRIGHT schools in previously ineligible villages whose relative scores are in the top 20%. The solid line shows the results for the scenario where we ignore the policy cost. The dashed line shows the results for the scenario where the policy cost measured in the unit of the enrollment rate is 0.137. I report the results for the range $[0.05, 0.1, ..., 0.95, 1]$ of the Lipschitz constant $C$.

maintains the status quo in the scenario where the policy cost is 0.137.[33] When $C \geq 0.6$, on the other hand, the minimax regret rule is randomized. The decisions become more mixed as $C$ increases.

Given that the estimate of the lower bound on $C$ is 0.149, the minimax regret rule is nonrandomized when $C$ is less than four times the estimated lower bound. Under this reasonable range of $C$, the optimal decision is the same in each scenario. If the policy maker wants to be more conservative about the choice of $C$, they need to randomize their decisions.

If the minimax regret rule is nonrandomized, the rule is of the form $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{\sum_{i=1}^n w_i Y_i \geq 0\}$ for some weights $w_i$'s. Panels (a) and (b) of Figure 3 plot the weight $w_i$ attached to each village against the relative score $x_i$ for $C = 0.1$ and $C = 0.5$, respectively. In the plots, the size of circles is proportional to the inverse of the standard error of the enrollment rate $Y_i$. For both $C = 0.1$ and $C = 0.5$, a few treated units just above the original cutoff (the solid vertical line) receive a positive weight, the untreated units between the original cutoff and the new cutoff (the dashed

[33] To examine the sensitivity of the result to the policy cost, I also compute the maximum of the cost values under which the minimax regret rule chooses the new policy with probability one. The result shows that, when $C$ is above its estimated lower bound of 0.149, it is optimal to maintain the status quo as long as the policy cost is higher than 0.10.

Figure 3: Weight to Each Village Attached by Minimax Regret Rule

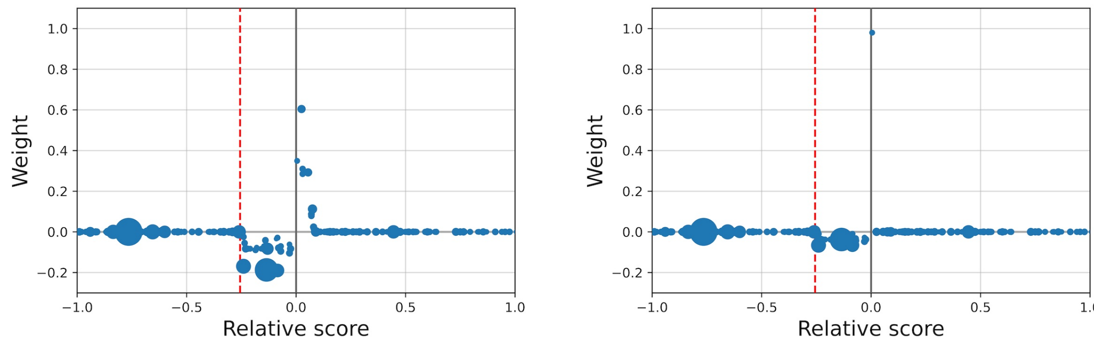


(a) $C = 0.1$ (b) $C = 0.5$

*Notes*: This figure shows the weight $w_i$ attached to each village by the minimax regret rule of the form $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{\sum_{i=1}^n w_i Y_i \geq 0\}$. The weights are normalized so that $\sum_{i=1}^n w_i^2 = 1$. The horizontal axis indicates the relative score of each village. Each circle corresponds to each village. The size of circles is proportional to the inverse of the standard error of the enrollment rate $Y_i$. The vertical dashed line corresponds to the new cutoff $-0.256$. Panels (a) and (b) show the results when the Lipschitz constant $C$ is 0.1 and 0.5, respectively.

vertical line) receive a negative weight, and no other units receive any weight. When $C = 0.1$, the weight tends to be larger for units with a smaller standard error. When $C = 0.5$, a positive weight is attached only to the treated unit closest to the original cutoff. Additionally, the weights on the untreated units between the two cutoffs are almost identical. This situation corresponds to the minimax regret rule of the form $\delta^*(\boldsymbol{Y}) = \mathbf{1}\left\{Y_{+,\min} - \frac{1}{\tilde{n}} \sum_{i:c_1 \leq x_i < c_0} Y_i \geq 0\right\}$ discussed in Section 4.

#### 5.3.1 Comparison with Plug-in Rules

I compare the minimax regret rule with plug-in decision rules that make a decision according to the sign of an estimator of the policy effect. I consider the following three estimators of the policy effect. (1) The linear minimax MSE estimator (Donoho, 1994), described in Section 3.4, under the Lipschitz class $\mathcal{F}_{\text{Lip}}(C)$; (2) The linear minimax MSE estimator under the additional assumption of constant conditional treatment effects. In other words, I construct the estimator assuming that $\mathcal{F} = \{f \in \mathcal{F}_{\text{Lip}}(C) : f(x, 1) - f(x, 0) = f(\tilde{x}, 1) - f(\tilde{x}, 0) \text{ for all } x, \tilde{x}\}$. This estimation corresponds to first nonparametrically estimating the average treatment effect at the original cutoff and then extrapolating the effects on the units between the two cutoffs

Figure 4: Estimated Effects of New Policy on Enrollment Rate

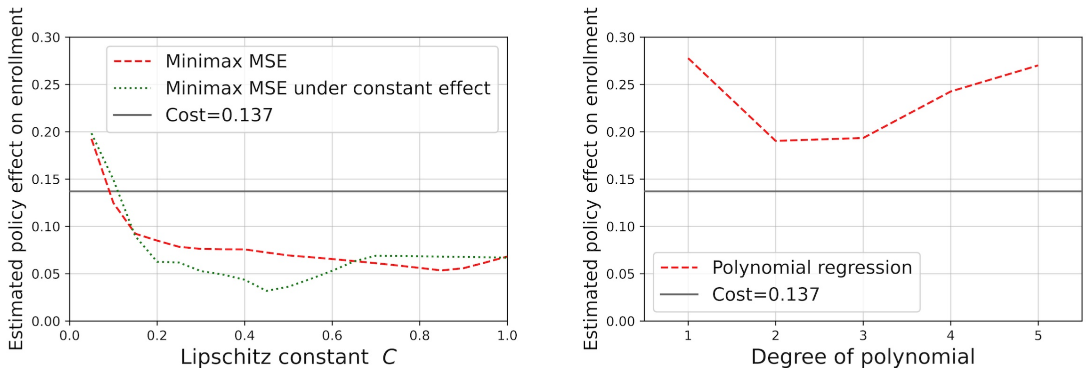


(a) Minimax MSE Estimator Under Lipschitz Class

(b) Polynomial Regression Estimator

*Notes*: This figure shows the average effect of the new policy on the enrollment rate across the villages that would receive a school under the new policy. Panel (a) reports the estimates from the linear minimax MSE estimators with and without the assumption of constant conditional treatment effects. I report the results for the range $[0.05, 0.1, ..., 0.95, 1]$ of the Lipschitz constant $C$. Panel (b) reports the estimates from the polynomial regression estimators of degrees 1 to 5. The horizontal line shows the cost of 0.137, which is my main specification of the policy cost.

by the constant effects assumption; and (3) The polynomial regression estimator (Kazianga *et al.*, 2013).[34] Given the degree of polynomial $p$, I first estimate the model $f(x,d) = \alpha_0 + \alpha_1 x + \cdots + \alpha_p x^p + \beta_0 d + \beta_1 d \cdot x + \cdots \beta_p d \cdot x^p$ by the weighted least squares regression using $1/\sigma^2(x_i, d_i)$ as the weight.[35] I then estimate $L(f)$ by $\frac{1}{n}\sum_{i=1}^n \mathbf{1}\{-0.256 \leq x_i < 0\}[\hat{f}(x_i, 1) - \hat{f}(x_i, 0)]$, where $\hat{f}$ is the estimated polynomial function. This estimator relies on the functional form of $f$ to extrapolate $f(x_i, 1)$ for the untreated units.

Panel (a) of Figure 4 reports the estimated policy effects from the linear minimax MSE estimators with and without constant conditional treatment effects. Overall, these two estimators exhibit a similar pattern. While the estimated policy effects are larger than the policy cost when $C$ is close to zero, they are smaller than the policy cost when $C$ is moderate or large. For $C \geq 0.2$, the resulting decisions about whether to choose the new policy are the same as the decision made by the minimax regret rule until $C$ reaches 0.6, where the minimax regret rule starts to randomize.

[34] Kazianga *et al.* (2013) estimate the treatment effect at the cutoff, not the effect on the units away from the cutoff. They apply global polynomial regression RD estimators to child-level data.

[35] This is equivalent to the OLS regression of $Y_i/\sigma(x_i, d_i)$ on $(1, x, ..., x^p, d, d\cdot x, ..., d\cdot x^p)'/\sigma(x_i, d_i)$.

In contrast, the estimated policy effects from the polynomial regression estimators of degrees 1 to 5 exceed the policy cost, as reported in Panel (b) of Figure 4. The estimates appear to be close to the simple mean outcome difference between eligible and ineligible villages that can be computed from Table 1. The resulting decisions are different from the decision made by the minimax regret rule.[36]

The above estimates and resulting decisions are computed from a particular realization of the sample. To assess the ex ante performance of different decision rules, I compute the maximum regret of these rules when the true function class is $\mathcal{F}_{\mathrm{Lip}}(C)$.[37] Panel (a) of Figure 5 reports the result for the minimax regret rule and the plug-in rules based on the linear minimax MSE estimators with and without constant conditional treatment effects.[38] The maximum regret of the plug-in MSE rule with constant conditional treatment effects is much larger than that of the other two, especially when the Lipschitz constant $C$ is large. The plug-in MSE rule without constant conditional treatment effects performs worse than the minimax regret rule, as predicted by the theoretical analysis. The ratio of the maximum regret between the two rules is maximized at $C = 0.6$, where the minimax regret rule starts to randomize.

### 5.3.2 Sensitivity to Misspecification of Lipschitz Constant $C$

So far, I have constructed decision rules assuming that the Lipschitz constant $C$ is known, which is a crucial assumption in my theoretical analysis. To assess the sensitivity of the performance to misspecification of $C$, I construct decision rules assuming $C = 0.3$ and then compute their maximum regret when the true value of $C$ lies in $\{0.05, 0.1, ..., 0.95, 1\}$.

Panel (b) of Figure 5 reports the result. The solid line indicates the "oracle" maximum regret, which can be achieved if we correctly specify $C$. The result shows that the plug-in MSE rule without constant conditional treatment effects performs

[36] The estimators presented here can be written as $\sum_{i=1}^{n} w_i Y_i$ for some weights $w_i$'s. See Figure D.2 in Appendix D for the plots of these weights. While the linear minimax MSE estimators attach weights to units just above the original cutoff and to units between the two cutoffs, polynomial regression estimators attach weights even to units further away from the cutoffs.

[37] I compute the maximum regret of the minimax regret rule using the formula in Theorem 1. For the other rules, I adapt the approach by Ishihara and Kitagawa (2021) to numerically calculate the maximum regret in this setup.

[38] The result for the plug-in rules based on polynomial regression estimators is omitted since these rules turn out to have significantly larger maximum regret than the other rules.

Figure 5: Maximum Regret of Minimax Regret Rule and Plug-in MSE Rules

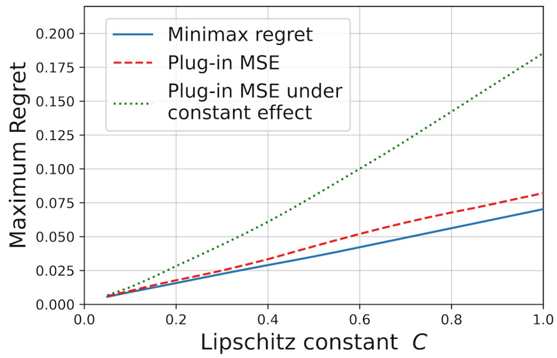


(a) Correctly Specified Lipschitz Constant

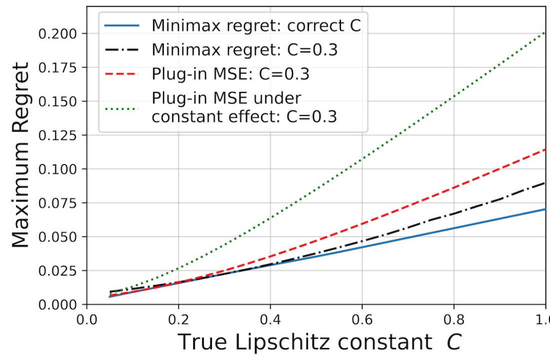


(b) Misspecified Lipschitz Constant

*Notes*: This figure shows the maximum regret of the minimax regret rule and the plug-in rules based on the linear minimax MSE estimators with and without the assumption of constant conditional treatment effects. Panel (a) reports the results for the rules that are constructed using the true Lipschitz constant $C$. Panel (b) reports the results for the rules that are constructed assuming $C = 0.3$, where the maximum regret is computed by setting the true $C$ to the value on the horizontal axis. The solid line in Panel (b) indicates the maximum regret that can be achieved if $C$ is correctly specified. In both panels, the maximum regret is normalized so that the unit is the same as that of the enrollment rate. I report the results for the range $[0.05, 0.1, ..., 0.95, 1]$ of the Lipschitz constant $C$.

slightly better than the minimax regret rule when the true $C$ is close to zero. On the other hand, the minimax regret rule outperforms the plug-in MSE rule with nonnegligible differences for any value of the true $C$ greater than 0.3. The result suggests that the minimax regret rule is more robust to misspecification of $C$ toward zero than the plug-in MSE rule.

The potential superiority of the minimax regret rule seems consistent with the theoretical results in the following way. As shown in Section 4, when the true value of $C$ is large, the oracle minimax regret rule only uses the treated units just above the original cutoff and the untreated units between the original and new cutoffs (see Panel (b) of Figure 3). If the specified $C$ is smaller than the true value, the resulting minimax regret rule is closer to the oracle rule than the plug-in MSE rule since the minimax regret rule places more importance on the bias than the plug-in MSE rule as discussed in Section 3.4. Therefore, it is expected that the minimax regret rule performs better than the plug-in MSE rule under misspecification of $C$ toward zero.

## 6 Conclusion

This paper develops an optimal procedure for using data to make policy decisions in settings where social welfare under each counterfactual policy is only partially identified. I derive a decision rule that achieves the minimax regret optimality in finite samples and within the class of all decision rules. I apply the result to the problem of eligibility cutoff choice and illustrate it in an empirical application to a school construction program in Burkina Faso.

Several extensions of my work are possible. First, my approach only covers a binary choice problem. It is both theoretically and practically important to extend the analysis to a multiple or continuous policy space. Second, while this paper shows that a randomized decision rule can be minimax regret, randomization may not be permitted in some cases due to ethical or legislative constraints. It would be interesting to consider the minimax regret problem over the class of all nonrandomized decision rules, including ones based on nonlinear functions of the sample.

## References


Abadie, A. and Imbens, G. W. (2006). Large Sample Properties of Matching Estimators for Average Treatment Effects. *Econometrica*, **74** (1), 235–267.

Agrawal, A., Verschueren, R., Diamond, S. and Boyd, S. (2018). A Rewriting System for Convex Optimization Problems. *Journal of Control and Decision*, **5** (1), 42–60.

Andrews, I., Kitagawa, T. and McCloskey, A. (2023). Inference on Winners. *Working Paper*.

Angrist, J. D. and Rokkanen, M. (2015). Wanna Get Away? Regression Discontinuity Estimation of Exam School Effects Away From the Cutoff. *Journal of the American Statistical Association*, **110** (512), 1331–1344.

Armstrong, T. B. and Kolesár, M. (2018). Optimal Inference in a Class of Regression Models. *Econometrica*, **86** (2), 655–683.

— and Kolesár, M. (2021). Finite-Sample Optimal Estimation and Inference on

Average Treatment Effects Under Unconfoundedness. *Econometrica*, **89** (3), 1141–1177.

Athey, S. and Wager, S. (2021). Policy Learning With Observational Data. *Econometrica*, **89** (1), 133–161.

Beliakov, G. (2006). Interpolation of Lipschitz Functions. *Journal of Computational and Applied Mathematics*, **196** (1), 20–44.

Ben-Michael, E., Greiner, D. J., Imai, K. and Jiang, Z. (2022). Safe Policy Learning through Extrapolation: Application to Pre-trial Risk Assessment. *arXiv:2109.11679*.

Bennett, M. (2020). How Far is Too Far? Estimation of an Interval for Generalization of a Regression Discontinuity Design Away from the Cutoff. *Working Paper*.

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis; 2nd ed.* Springer Series in Statistics, New York: Springer.

Bertanha, M. (2020). Regression Discontinuity Design with Many Thresholds. *Journal of Econometrics*, **218** (1), 216–241.

— and Imbens, G. W. (2020). External Validity in Fuzzy Regression Discontinuity Designs. *Journal of Business & Economic Statistics*, **38** (3), 593–612.

Bhattacharya, D. and Dupas, P. (2012). Inferring Welfare Maximizing Treatment Assignment Under Budget Constraints. *Journal of Econometrics*, **167** (1), 168–196.

Cai, T. T. and Low, M. G. (2004). An Adaptation Theory for Nonparametric Confidence Intervals. *The Annals of Statistics*, **32** (5), 1805–1840.

Calonico, S., Cattaneo, M. D. and Farrell, M. H. (2018). On the Effect of Bias Estimation on Coverage Accuracy in Nonparametric Inference. *Journal of the American Statistical Association*, **113** (522), 767–779.

Cattaneo, M. D., Keele, L., Titiunik, R. and Vazquez-Bare, G. (2021). Extrapolating Treatment Effects in Multi-Cutoff Regression Discontinuity Designs. *Journal of the American Statistical Association*, **116** (536), 1941–1952.

Chamberlain, G. (2011). Bayesian Aspects of Treatment Choice. In *The Oxford Handbook of Bayesian Econometrics*, Oxford University Press.

Chernozhukov, V., Lee, S. and Rosen, A. M. (2013). Intersection Bounds: Estimation and Inference. *Econometrica*, **81** (2), 667–737.

Christensen, T., Moon, H. R. and Schorfheide, F. (2023). Optimal Discrete

Decisions when Payoffs are Partially Identified. *arXiv:2204.11748*.

D'Adamo, R. (2022). Orthogonal Policy Learning Under Ambiguity. *arXiv:2111.10904*.

de Chaisemartin, C. (2021). The Minimax Estimator of the Average Treatment Effect, among Linear Combinations of Estimators of Bounded Conditional Average Treatment Effects. *arXiv:2105.08766*.

Dehejia, R. H. (2005). Program Evaluation as a Decision Problem. *Journal of Econometrics*, **125** (1), 141–173.

Diamond, S. and Boyd, S. (2016). CVXPY: A Python-embedded Modeling Language for Convex Optimization. *Journal of Machine Learning Research*, **17** (83), 1–5.

Dong, Y. and Lewbel, A. (2015). Identifying The Effect of Changing The Policy Threshold in Regression Discontinuity Models. *The Review of Economics and Statistics*, **97** (5), 1081–1092.

Donoho, D. L. (1994). Statistical Estimation and Optimal Recovery. *The Annals of Statistics*, **22** (1), 238–270.

Duflo, E. (2001). Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *American Economic Review*, **91** (4), 795–813.

Fan, J. and Yao, Q. (1998). Efficient Estimation of Conditional Variance Functions in Stochastic Regression. *Biometrika*, **85** (3), 645–660.

Hirano, K. and Porter, J. R. (2009). Asymptotics for Statistical Treatment Rules. *Econometrica*, **77** (5), 1683–1701.

Imbens, G. and Wager, S. (2019). Optimized Regression Discontinuity Designs. *The Review of Economics and Statistics*, **101** (2), 264–278.

Ishihara, T. and Kitagawa, T. (2021). Evidence Aggregation for Treatment Choice. *arXiv:2108.06473*.

Kallus, N. and Zhou, A. (2021). Minimax-Optimal Policy Learning Under Unobserved Confounding. *Management Science*, **67** (5), 2870–2890.

Kasy, M. (2018). Optimal Taxation and Insurance Using Machine Learning — Sufficient Statistics and Beyond. *Journal of Public Economics*, **167**, 205–219.

Kazianga, H., Levy, D., Linden, L. L. and Sloan, M. (2013). The Effects of

"Girl-Friendly" Schools: Evidence from the BRIGHT School Construction Program in Burkina Faso. *American Economic Journal: Applied Economics*, **5** (3), 41–62.

—, —, — and — (2019). Replication Data for: The Effects of "Girl-Friendly" Schools: Evidence from the BRIGHT School Construction Program in Burkina Faso, TN: American Economic Association [publisher], 2013. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2019-10-12. https://doi.org/10.3886/E113862V1.

Kitagawa, T., Lee, S. and Qiu, C. (2023). Treatment Choice with Nonlinear Regret. *arXiv:2205.08586*.

— and Tetenov, A. (2018). Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice. *Econometrica*, **86** (2), 591–616.

— and — (2021). Equality-Minded Treatment Choice. *Journal of Business & Economic Statistics*, **39** (2), 561–574.

Kwon, K. and Kwon, S. (2020). Inference in Regression Discontinuity Designs under Monotonicity. *arXiv:2011.14216*.

Levy, D., Sloan, M., Linden, L. L. and Kazianga, H. (2009). Impact Evaluation of Burkina Faso's BRIGHT Program: Final Report. Mathematica Policy Research, Washington, DC, USA.

Low, M. G. (1995). Bias-Variance Tradeoffs in Functional Estimation Problems. *The Annals of Statistics*, **23** (3), 824–835.

Manski, C. F. (2000). Identification Problems and Decisions Under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice. *Journal of Econometrics*, **95** (2), 415–442.

— (2004). Statistical Treatment Rules for Heterogeneous Populations. *Econometrica*, **72** (4), 1221–1246.

— (2007). Minimax-regret Treatment Choice with Missing Outcome Data. *Journal of Econometrics*, **139** (1), 105–115.

— (2009). The 2009 Lawrence R. Klein Lecture: Diversified Treatment Under Ambiguity. *International Economic Review*, **50** (4), 1013–1041.

— (2010). Vaccination with Partial Knowledge of External Effectiveness. *Proceedings of the National Academy of Sciences*, **107** (9), 3953–3960.

— (2011a). Choosing Treatment Policies Under Ambiguity. *Annual Review of Eco-*

*nomics*, **3** (1), 25–49.

— (2011b). Policy Choice with Partial Knowledge of Policy Effectiveness. *Journal of Experimental Criminology*, **7**, 111–125.

— (2021). Econometrics For Decision Making: Building Foundations Sketched By Haavelmo And Wald. *Econometrica*, **89** (6), 2827–2853.

Mbakop, E. and Tabord-Meehan, M. (2021). Model Selection for Treatment Choice: Penalized Welfare Maximization. *Econometrica*, **89** (2), 825–848.

Mo, W., Qi, Z. and Liu, Y. (2021). Learning Optimal Distributionally Robust Individualized Treatment Rules. *Journal of the American Statistical Association*, **116** (534), 659–674.

Qian, M. and Murphy, S. A. (2011). Performance Guarantees for Individualized Treatment Rules. *The Annals of Statistics*, **39** (2), 1180–1210.

Rambachan, A. and Roth, J. (2023). A More Credible Approach to Parallel Trends. *The Review of Economic Studies*, **rdad018**.

Rokkanen, M. (2015). Exam Schools, Ability, and the Effects of Affirmative Action: Latent Factor Extrapolation in the Regression Discontinuity Design. Discussion Paper 1415-03, Columbia University, Department of Economics, New York, NY.

Russell, T. M. (2020). Policy Transforms and Learning Optimal Policies. *arXiv:2012.11046*.

Song, Z. (2020). Potential Implications of Lowering the Medicare Eligibility Age to 60. *Journal of the American Medical Association (JAMA)*, **323** (24), 2472–2473.

Stoye, J. (2009). Minimax Regret Treatment Choice with Finite Samples. *Journal of Econometrics*, **151** (1), 70–81.

— (2012). Minimax Regret Treatment Choice with Covariates or with Limited Validity of Experiments. *Journal of Econometrics*, **166** (1), 138–156.

Tetenov, A. (2012). Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria. *Journal of Econometrics*, **166** (1), 157–165.

Wald, A. (1950). *Statistical Decision Functions*. New York: Wiley.

Zhang, Y., Ben-Michael, E. and Imai, K. (2022). Safe Policy Learning under Regression Discontinuity Designs. *arXiv:2208.13323*.

Online Appendix to
"Optimal Decision Rules Under Partial Identification"

Kohei Yata

# A Auxiliary Lemmas and Proofs of Main Results

## A.1 Auxiliary Lemmas

**Lemma A.1.** *Let $g(t) = h(t)\Phi\left(\frac{b-t}{a}\right)$, where $h(t)$ is nonnegative, nonconstant, nondecreasing, concave, and differentiable on $[\underline{t}, \bar{t}]$, $a > 0$, and $b \in \mathbb{R}$. If $a\frac{1-\Phi\left(\frac{\underline{t}-b}{a}\right)}{\phi\left(\frac{\underline{t}-b}{a}\right)} \leq \frac{h(\underline{t})}{h'(\underline{t})}$, then $g'(t) \leq 0$ for all $t \in [\underline{t}, \bar{t}]$ with the inequality strict for all $t \in (\underline{t}, \bar{t}]$, and hence $g(t)$ is strictly decreasing on $[\underline{t}, \bar{t}]$. If $a\frac{1-\Phi\left(\frac{\bar{t}-b}{a}\right)}{\phi\left(\frac{\bar{t}-b}{a}\right)} \geq \frac{h(\bar{t})}{h'(\bar{t})}$, then $g'(t) \geq 0$ for all $t \in [\underline{t}, \bar{t}]$ with the inequality strict for all $t \in [\underline{t}, \bar{t})$, and hence $g(t)$ is strictly increasing on $[\underline{t}, \bar{t}]$. If $a\frac{1-\Phi\left(\frac{\underline{t}-b}{a}\right)}{\phi\left(\frac{\underline{t}-b}{a}\right)} > \frac{h(\underline{t})}{h'(\underline{t})}$ and $a\frac{1-\Phi\left(\frac{\bar{t}-b}{a}\right)}{\phi\left(\frac{\bar{t}-b}{a}\right)} < \frac{h(\bar{t})}{h'(\bar{t})}$ (including the case where $h'(\bar{t}) = 0$), then $g(t)$ is strictly increasing on $[\underline{t}, t^*)$ and strictly decreasing on $(t^*, \bar{t}]$, where $t^*$ is the unique solution to $a\frac{1-\Phi\left(\frac{t-b}{a}\right)}{\phi\left(\frac{t-b}{a}\right)} = \frac{h(t)}{h'(t)}$.*

*Proof.* Note first that $h(t)$ is continuously differentiable since it is concave and differentiable. Note also that $h'(\underline{t}) > 0$; if $h'(\underline{t}) \leq 0$, then $h'(t) = 0$ for all $t \in [\underline{t}, \bar{t}]$ since $h(t)$ is nondecreasing and concave, but this contradicts the assumption that $h(t)$ is nonconstant. In addition, $h(t) > 0$ for all $t \in (\underline{t}, \bar{t}]$ since $h(t)$ is nonnegative and nondecreasing and $h'(\underline{t}) > 0$.

Suppose that $h'(\bar{t}) > 0$, which implies $h'(t) > 0$ for all $t \in [\underline{t}, \bar{t}]$ since $h(t)$ is concave. By differentiating $g(t)$, we have for $t \in [\underline{t}, \bar{t}]$,

$$g'(t) = h'(t)\Phi\left(\frac{b-t}{a}\right) - \frac{h(t)}{a}\phi\left(\frac{b-t}{a}\right) = \left[a\frac{1-\Phi\left(\frac{t-b}{a}\right)}{\phi\left(\frac{t-b}{a}\right)} - \frac{h(t)}{h'(t)}\right]\frac{h'(t)}{a}\phi\left(\frac{t-b}{a}\right),$$

where the second equality holds since $\Phi(x) = 1 - \Phi(-x)$ and $\phi(x) = \phi(-x)$. By the fact that the Mills ratio $\frac{1-\Phi(x)}{\phi(x)}$ of a standard normal random variable is strictly decreasing, $a\frac{1-\Phi\left(\frac{t-b}{a}\right)}{\phi\left(\frac{t-b}{a}\right)}$ is strictly decreasing in $t$. In addition, $a\frac{1-\Phi\left(\frac{t-b}{a}\right)}{\phi\left(\frac{t-b}{a}\right)}$ is continuous.

Furthermore, since $h(t)$ is nonnegative, nondecreasing, and concave on $[\underline{t},\bar{t}]$, $\frac{h(t)}{h'(t)}$ is nondecreasing on $[\underline{t},\bar{t}]$. In addition, $\frac{h(t)}{h'(t)}$ is continuous since $h(t)$ is continuously differentiable. Therefore, if $a\frac{1-\Phi\left(\frac{\underline{t}-b}{a}\right)}{\phi\left(\frac{\underline{t}-b}{a}\right)} \leq \frac{h(\underline{t})}{h'(\underline{t})}$, then $g'(t) \leq 0$ for all $t \in [\underline{t},\bar{t}]$ with the inequality strict for all $t \in (\underline{t},\bar{t}]$. If $a\frac{1-\Phi\left(\frac{\bar{t}-b}{a}\right)}{\phi\left(\frac{\bar{t}-b}{a}\right)} \geq \frac{h(\bar{t})}{h'(\bar{t})}$, then $g'(t) \geq 0$ for all $t \in [\underline{t},\bar{t}]$ with the inequality strict for all $t \in [\underline{t},\bar{t})$. If $a\frac{1-\Phi\left(\frac{\underline{t}-b}{a}\right)}{\phi\left(\frac{\underline{t}-b}{a}\right)} > \frac{h(\underline{t})}{h'(\underline{t})}$ and $a\frac{1-\Phi\left(\frac{\bar{t}-b}{a}\right)}{\phi\left(\frac{\bar{t}-b}{a}\right)} < \frac{h(\bar{t})}{h'(\bar{t})}$, then $g'(t) > 0$ for $t \in [\underline{t},t^*)$, $g'(t^*) = 0$, and $g'(t) < 0$ for $t \in (t^*,\bar{t}]$, where $t^*$ is the unique solution to $a\frac{1-\Phi\left(\frac{t-b}{a}\right)}{\phi\left(\frac{t-b}{a}\right)} = \frac{h(t)}{h'(t)}$.

Suppose that $h'(\bar{t}) = 0$. Since $h(t)$ is concave and continuously differentiable and $h'(\underline{t}) > 0$, $h'(t) > 0$ for $t \in [\underline{t},\tilde{t})$ and $h'(t) = 0$ for $t \in [\tilde{t},\bar{t}]$, where $\tilde{t} = \sup\{t \in [\underline{t},\bar{t}] : h'(t) > 0\}$. It follows that $g'(t) = -h(t)\phi\left(\frac{b-t}{a}\right)/a < 0$ for all $t \in [\tilde{t},\bar{t}]$. For $t \in [\underline{t},\tilde{t})$, $g'(t) = \left[a\frac{1-\Phi\left(\frac{t-b}{a}\right)}{\phi\left(\frac{t-b}{a}\right)} - \frac{h(t)}{h'(t)}\right] h'(t)\phi\left(\frac{t-b}{a}\right)/a$. By the same argument as the one for the case where $h'(\bar{t}) > 0$, if $a\frac{1-\Phi\left(\frac{\underline{t}-b}{a}\right)}{\phi\left(\frac{\underline{t}-b}{a}\right)} \leq \frac{h(\underline{t})}{h'(\underline{t})}$, then $g'(t) \leq 0$ for all $t \in [\underline{t},\tilde{t})$ with the inequality strict for all $t \in (\underline{t},\tilde{t})$. If $a\frac{1-\Phi\left(\frac{\underline{t}-b}{a}\right)}{\phi\left(\frac{\underline{t}-b}{a}\right)} > \frac{h(\underline{t})}{h'(\underline{t})}$, then $g'(t) > 0$ for $t \in [\underline{t},t^*)$, $g'(t^*) = 0$, and $g'(t) < 0$ for $t \in (t^*,\tilde{t})$, where $t^*$ solves $a\frac{1-\Phi\left(\frac{t-b}{a}\right)}{\phi\left(\frac{t-b}{a}\right)} = \frac{h(t)}{h'(t)}$. □

**Lemma A.2.** *Let $\psi(a,b) = a\Phi(-b)$. Then, $\psi(a,b)$ is strictly quasi-concave on $(0,\infty)\times\mathbb{R}$.*

*Proof.* Take any $a_0, a_1 > 0$ and $b_0, b_1 \in \mathbb{R}$ such that $(a_0,b_0) \neq (a_1,b_1)$. I show that $\psi(a_0+\lambda(a_1-a_0), b_0+\lambda(b_1-b_0)) > \min\{\psi(a_0,b_0),\psi(a_1,b_1)\}$ for all $\lambda \in (0,1)$.

First, suppose that $a_0 \leq a_1$ and $b_0 \geq b_1$. Since either $a_0 < a_1$ or $b_0 > b_1$ or both must hold, $\psi(a_0+\lambda(a_1-a_0), b_0+\lambda(b_1-b_0)) = (a_0+\lambda(a_1-a_0))\Phi(-b_0-\lambda(b_1-b_0))$ is strictly increasing in $\lambda$. It then follows that $\psi(a_0+\lambda(a_1-a_0), b_0+\lambda(b_1-b_0)) > \psi(a_0,b_0)$. Likewise, if $a_0 \geq a_1$ and $b_0 \leq b_1$, then $\psi(a_0+\lambda(a_1-a_0), b_0+\lambda(b_1-b_0)) > \psi(a_1,b_1)$.

Now suppose that $a_0 < a_1$ and $b_0 < b_1$. Note that the set $\{(a_0,b_0)+\lambda(a_1-a_0,b_1-b_0) : \lambda \in (0,1)\}$ is equivalent to $\left\{\left(0, b_0 - a_0\frac{b_1-b_0}{a_1-a_0}\right) + t\left(\frac{a_1-a_0}{b_1-b_0},1\right) : t \in \left(a_0\frac{b_1-b_0}{a_1-a_0}, a_1\frac{b_1-b_0}{a_1-a_0}\right)\right\}$. We have $\psi\left(\left(0, b_0 - a_0\frac{b_1-b_0}{a_1-a_0}\right) + t\left(\frac{a_1-a_0}{b_1-b_0},1\right)\right) = t\left(\frac{a_1-a_0}{b_1-b_0}\right)\Phi\left(-b_0 + a_0\frac{b_1-b_0}{a_1-a_0} - t\right) = \left(\frac{a_1-a_0}{b_1-b_0}\right)g(t)$, where $g(t) = t\Phi\left(-b_0 + a_0\frac{b_1-b_0}{a_1-a_0} - t\right)$. Lemma A.1 implies that the minimum of $g(t)$ over an interval $[\underline{t}, \bar{t}] \subset \mathbb{R}_+$ is attained only at $\underline{t}$ or $\bar{t}$ or both. Hence, for all $t \in \left(a_0 \frac{b_1-b_0}{a_1-a_0}, a_1 \frac{b_1-b_0}{a_1-a_0}\right)$, $g(t) > \min\left\{g\left(a_0 \frac{b_1-b_0}{a_1-a_0}\right), g\left(a_1 \frac{b_1-b_0}{a_1-a_0}\right)\right\}$. Thus, for all $t \in \left(a_0 \frac{b_1-b_0}{a_1-a_0}, a_1 \frac{b_1-b_0}{a_1-a_0}\right)$, $\psi\left(\left(0, b_0 - a_0 \frac{b_1-b_0}{a_1-a_0}\right) + t\left(\frac{a_1-a_0}{b_1-b_0}, 1\right)\right) > \left(\frac{a_1-a_0}{b_1-b_0}\right) \min\left\{g\left(a_0 \frac{b_1-b_0}{a_1-a_0}\right), g\left(a_1 \frac{b_1-b_0}{a_1-a_0}\right)\right\} = \min\{\psi(a_0, b_0), \psi(a_1, b_1)\}$. Therefore, $\psi(a_0 + \lambda(a_1 - a_0), b_0 + \lambda(b_1 - b_0)) > \min\{\psi(a_0, b_0), \psi(a_1, b_1)\}$ for all $\lambda \in (0, 1)$. The same argument holds for the case where $a_0 > a_1$ and $b_0 > b_1$. $\square$

The lemma below immediately follows from Lemma D.1 in Supplemental Appendix D of Armstrong and Kolesár (2018) in the case where $\mathcal{F} = \mathcal{G}$ in their notation, and hence the proof is omitted.[a]

**Lemma A.3.** *Let $\Theta$ be convex. Let $\theta_\epsilon$ attain the modulus of continuity at $\epsilon > 0$ with $\|\boldsymbol{m}(\theta_\epsilon)\| = \epsilon$, and suppose that there exists $\iota \in \Theta$ such that $L(\iota) = 1$ and $\theta_\epsilon + c\iota \in \Theta$ for all $c$ in a neighborhood of zero. Then, $\omega(\cdot)$ is differentiable at $\epsilon$ with $\omega'(\epsilon) = \frac{\epsilon}{\boldsymbol{m}(\iota)'\boldsymbol{m}(\theta_\epsilon)}$.*

Assuming $\boldsymbol{w}^*$ exists, let $\rho(\epsilon) = \sup\{L(\theta) : (\boldsymbol{w}^*)'\boldsymbol{m}(\theta) = \epsilon, \theta \in \Theta\}$ as defined in Assumption 1(i)(d). By the convexity of $\Theta$ and linearity of $\boldsymbol{m}$, $\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\}$ is convex.

**Lemma A.4.** *Let $\Theta$ be convex. Then, $\rho(\cdot)$ is concave on $\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\}$.*

*Proof.* Pick any $\epsilon, \epsilon' \in \{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\}$, and let $\{\theta_{\epsilon,n}\}_{n=1}^\infty$ and $\{\theta_{\epsilon',n}\}_{n=1}^\infty$ be sequences in $\Theta$ such that $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_{\epsilon,n}) = \epsilon$ and $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_{\epsilon',n}) = \epsilon'$ for all $n$ and that $\lim_{n\to\infty} L(\theta_{\epsilon,n}) = \rho(\epsilon)$ and $\lim_{n\to\infty} L(\theta_{\epsilon',n}) = \rho(\epsilon')$. Then, for each $\lambda \in [0, 1]$, $\lambda\theta_{\epsilon,n} + (1-\lambda)\theta_{\epsilon',n} \in \Theta$ by the convexity of $\Theta$, and $(\boldsymbol{w}^*)'\boldsymbol{m}(\lambda\theta_{\epsilon,n} + (1-\lambda)\theta_{\epsilon',n}) = \lambda\epsilon + (1-\lambda)\epsilon'$ so that $\rho(\lambda\epsilon + (1-\lambda)\epsilon') \geq L(\lambda\theta_{\epsilon,n} + (1-\lambda)\theta_{\epsilon',n})$ by the definition of $\rho$. Taking the limit of the right-hand side as $n \to \infty$ gives $\rho(\lambda\epsilon + (1-\lambda)\epsilon') \geq \lambda\rho(\epsilon) + (1-\lambda)\rho(\epsilon')$. $\square$

For a set $S \subset \mathbb{R}$, $\mathrm{int}(S)$ denotes the interior of $S$.

**Lemma A.5.** *Let $\Theta$ be convex. Let $\theta_\epsilon \in \Theta$ satisfy $L(\theta_\epsilon) = \rho(\epsilon)$ and $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_\epsilon) = \epsilon \in \mathrm{int}(\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\})$, and suppose that there exists $\iota \in \Theta$ such that $L(\iota) = 1$ and $\theta_\epsilon + c\iota \in \Theta$ for all $c$ in a neighborhood of zero. Then, $\rho(\cdot)$ is differentiable at $\epsilon$ with $\rho'(\epsilon) = \frac{1}{(\boldsymbol{w}^*)'\boldsymbol{m}(\iota)}$.*

[a] Note that their definition of the modulus of continuity when $\mathcal{F} = \mathcal{G}$ is the same as Donoho (1994)'s definition (i.e., $\tilde{\omega}(\epsilon) = \sup\{|L(\theta) - L(\tilde{\theta})| : \|\boldsymbol{m}(\theta - \tilde{\theta})\| \leq \epsilon, \theta, \tilde{\theta} \in \Theta\}$), which is different from my definition.

*Proof.* Let $I$ denote $\mathrm{int}(\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\})$. Since $\rho(\cdot)$ is concave on $I$ by Lemma A.4, the superdifferential of $\rho(\cdot)$ at $\epsilon$, $\partial\rho(\epsilon) = \{d : \rho(\eta) \leq \rho(\epsilon) + d(\eta - \epsilon) \text{ for all } \eta \in I\}$, is nonempty for all $\epsilon \in I$. Let $\theta_\epsilon$ satisfy $L(\theta_\epsilon) = \rho(\epsilon)$ and $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_\epsilon) = \epsilon$ for some $\epsilon \in I$, and suppose that there exists $\iota \in \Theta$ such that $L(\iota) = 1$ and $\theta_\epsilon + c\iota \in \Theta$ for all $c$ in a neighborhood of zero. Then, for any $d \in \partial\rho(\epsilon)$ and for any $c$ in a neighborhood of zero such that $\theta_\epsilon + c\iota \in \Theta$ and $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_\epsilon + c\iota) \in I$,

$$\rho(\epsilon) + d[(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_\epsilon + c\iota) - \epsilon] \geq \rho((\boldsymbol{w}^*)'\boldsymbol{m}(\theta_\epsilon + c\iota)) \geq L(\theta_\epsilon + c\iota) = L(\theta_\epsilon) + c = \rho(\epsilon) + c,$$

where the second inequality follows from the definition of $\rho$. Since $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_\epsilon + c\iota) = \epsilon + c(\boldsymbol{w}^*)'\boldsymbol{m}(\iota)$, it follows that $cd(\boldsymbol{w}^*)'\boldsymbol{m}(\iota) \geq c$ for all $c$ in a neighborhood of zero. This implies that $d(\boldsymbol{w}^*)'\boldsymbol{m}(\iota) = 1$. The result then follows. $\square$

**Lemma A.6.** *Let $\Theta$ be convex, and suppose that Assumption A.2(i)(a)–(b) and (ii) in Appendix A.3.2 hold. Then, $\omega(0) = \rho(0)$.*

*Proof.* Since $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) = 0$ for any $\theta \in \Theta$ such that $\boldsymbol{m}(\theta) = \boldsymbol{0}$, it follows that $\rho(0) \geq \omega(0)$ by the definition of $\omega(\cdot)$ and $\rho(\cdot)$. Suppose to the contrary that $\rho(0) > \omega(0)$. Then there exists $\theta \in \Theta$ such that $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) = 0$, $\boldsymbol{m}(\theta) \neq \boldsymbol{0}$, and $L(\theta) > \omega(0)$. For $\epsilon \in (0, \bar{\epsilon}]$, by the Cauchy-Schwarz inequality,

$$\epsilon^{-1}\left|\frac{\boldsymbol{m}(\theta_\epsilon)'}{\epsilon}\boldsymbol{m}(\theta)\right| = \epsilon^{-1}\left|\left(\frac{\boldsymbol{m}(\theta_\epsilon)}{\|\boldsymbol{m}(\theta_\epsilon)\|} - \boldsymbol{w}^*\right)'\boldsymbol{m}(\theta)\right| \leq \epsilon^{-1}\left\|\frac{\boldsymbol{m}(\theta_\epsilon)}{\|\boldsymbol{m}(\theta_\epsilon)\|} - \boldsymbol{w}^*\right\|\|\boldsymbol{m}(\theta)\|,$$

where $\theta_\epsilon$ attains the modulus of continuity at $\epsilon$, and the first equality follows from Assumption A.2(i)(a) and the fact that $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) = 0$. The right-hand side converges to zero as $\epsilon \to 0$ by Assumption A.2(i)(b). Together with the fact that $L(\theta) > \omega(0)$ and the continuity of $\omega(\cdot)$ at $\epsilon = 0$ (Assumption A.2(ii)), this implies that $\epsilon^{-1}\frac{\boldsymbol{m}(\theta_\epsilon)'}{\epsilon}\boldsymbol{m}(\theta) \leq 1/2$ and $L(\theta) > \omega(\epsilon)$ for any sufficiently small $\epsilon > 0$. Pick such an $\epsilon \in (0, 2\|\boldsymbol{m}(\theta)\|)$, and let $\theta_\lambda = \lambda\theta_\epsilon + (1-\lambda)\theta$ for $\lambda \in [0,1]$. By the convexity of $\Theta$, $\theta_\lambda \in \Theta$. By simple algebra,

$$\begin{aligned}
\|\boldsymbol{m}(\theta_\lambda)\|^2 &= \lambda^2\|\boldsymbol{m}(\theta_\epsilon)\|^2 + 2\lambda(1-\lambda)\boldsymbol{m}(\theta_\epsilon)'\boldsymbol{m}(\theta) + (1-\lambda)^2\|\boldsymbol{m}(\theta)\|^2 \\
&\leq \lambda^2\epsilon^2 + \lambda(1-\lambda)\epsilon^2 + (1-\lambda)^2\|\boldsymbol{m}(\theta)\|^2 \\
&= \|\boldsymbol{m}(\theta)\|^2\lambda^2 - (2\|\boldsymbol{m}(\theta)\|^2 - \epsilon^2)\lambda + \|\boldsymbol{m}(\theta)\|^2.
\end{aligned}$$

Observe that the right-hand side is quadratic in $\lambda$, minimized at $\lambda = \frac{2\|\boldsymbol{m}(\theta)\|^2 - \epsilon^2}{2\|\boldsymbol{m}(\theta)\|^2} \in (0,1)$, and equal to $\epsilon^2$ when $\lambda = 1$. This implies that $\|\boldsymbol{m}(\theta_\lambda)\|^2 < \epsilon^2$ for any $\lambda$ close to one. However, $L(\theta_\lambda) = \lambda L(\theta_\epsilon) + (1-\lambda)L(\theta) > \omega(\epsilon)$ for all $\lambda \in (0,1)$, which contradicts the assumption that $\theta_\epsilon$ attains the modulus of continuity at $\epsilon$. $\square$

The next lemma derives a minimax regret rule for a class of univariate problems.

**Lemma A.7** (Minimax Rules for Univariate Problems)**.** *Suppose that* $\Theta = [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1]$ *for some* $\tau_0, \tau_1$ *such that* $\tau_1 \geq \tau_0 \geq 0$ *and* $\tau_1 > 0$*, that* $\boldsymbol{m}(\theta) = \theta$*, and that* $L(\theta) = \theta$*. Then, the decision rule* $\delta^*(Y) = \mathbf{1}\{Y \geq 0\}$ *is minimax regret. The minimax risk is given by*

$$\mathcal{R}_{\text{uni}}(\sigma; [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1]) = \begin{cases} \tau_1 \Phi(-\tau_1/\sigma) & \text{if } \tau_1 < a^*\sigma, \\ a^*\sigma\Phi(-a^*) & \text{if } a^*\sigma \in [\tau_0, \tau_1], \\ \tau_0\Phi(-\tau_0/\sigma) & \text{if } a^*\sigma < \tau_0. \end{cases}$$

*Proof.* Following Stoye (2009), I use a statistical game to solve the minimax regret problem. Consider the following two-person zero-sum game between the decision maker and nature. The strategy space for the decision maker is $\mathcal{D}$, the set of all decision rules. The strategy space for nature is $\Delta(\Theta)$, the set of probability distributions on $\Theta = [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1]$. If the decision maker chooses $\delta \in \mathcal{D}$ and nature chooses $\pi \in \Delta(\Theta)$, nature's expected payoff (and the decision maker's expected loss) is given by $r(\delta, \pi) = \int R(\delta, \theta) d\pi(\theta)$, the Bayes risk of $\delta$ with respect to prior $\pi$. By Theorem 17 in Chapter 5 of Berger (1985), if $(\delta^*, \pi^*)$ satisfies

$$\delta^* \in \arg\min_{\delta \in \mathcal{D}} r(\delta, \pi^*), \text{ and } R(\delta^*, \theta) \leq r(\delta^*, \pi^*) \text{ for all } \theta \in \Theta, \tag{A.1}$$

then $\delta^*$ is a minimax regret rule. Below I construct $(\delta^*, \pi^*)$ that satisfies (A.1).

I first restrict the search space of decision rules to an essentially complete class of decision rules, following Tetenov (2012).[b] Since $Y$ has monotone likelihood ratio and the loss function satisfies $l(1,\theta) - l(0,\theta) \geq 0$ if $\theta < 0$ and $l(1,\theta) - l(0,\theta) \leq 0$ if $\theta > 0$, it follows from Theorem 5 in Chapter 8 of Berger (1985) (which is originally

[b] A class $\mathcal{C}$ of decision rules is essentially complete if, for any decision rule $\delta \notin \mathcal{C}$, there is a decision rule $\delta' \in \mathcal{C}$ such that $R(\delta, \theta) \geq R(\delta', \theta)$ for all $\theta \in \Theta$.

from Karlin and Rubin (1956)) that the class of monotone decision rules $\delta(Y) = 0 \cdot \mathbf{1}\{Y < t\} + \lambda \cdot \mathbf{1}\{Y = t\} + 1 \cdot \mathbf{1}\{Y > t\}$, where $t \in \mathbb{R}$ and $\lambda \in [0,1]$, is essentially complete. Furthermore, since $\mathbb{P}_\theta(Y = t) = 0$, a smaller class of threshold decision rules $\delta(Y) = \mathbf{1}\{Y \geq t\}$, $t \in \mathbb{R}$, is also essentially complete.

Let $\delta_t$ denote the threshold rule with threshold $t$. Since $Y \sim \mathcal{N}(\theta, \sigma^2)$, $R(\delta_t, \theta) = \theta^+ \Phi(\sigma^{-1}(t-\theta)) + (-\theta)^+(1 - \Phi(\sigma^{-1}(t-\theta)))$, where $x^+ = \max\{x, 0\}$. Let $\bar{R}_0(t, \tau_0, \tau_1) = \max_{\theta \in [-\tau_1, -\tau_0]} R(\delta_t, \theta) = \max_{\theta \in [-\tau_1, -\tau_0]} -\theta(1 - \Phi(\sigma^{-1}(t-\theta)))$ and $\bar{R}_1(t, \tau_0, \tau_1) = \max_{\theta \in [\tau_0, \tau_1]} R(\delta_t, \theta) = \max_{\theta \in [\tau_0, \tau_1]} \theta \Phi(\sigma^{-1}(t-\theta))$. By symmetry of $\bar{R}_0(t, \tau_0, \tau_1)$ and $\bar{R}_1(t, \tau_0, \tau_1)$, $\bar{R}_0(0, \tau_0, \tau_1) = \bar{R}_1(0, \tau_0, \tau_1)$.

Now let $\theta_0^* \in \arg\max_{\theta \in [-\tau_1, -\tau_0]} R(\delta_0, \theta)$ and $\theta_1^* \in \arg\max_{\theta \in [\tau_0, \tau_1]} R(\delta_0, \theta)$, where $\delta_0(Y) = \mathbf{1}\{Y \geq 0\}$. By symmetry, we can pick $(\theta_0^*, \theta_1^*)$ such that $\theta_0^* = -\theta_1^*$. Let $\pi^* \in \Delta(\Theta)$ be such that $\pi^*(\theta_0^*) = \pi^*(\theta_1^*) = \frac{1}{2}$.

I show that $(\delta_0, \pi^*)$ satisfies (A.1). Since $\bar{R}_0(0, \tau_0, \tau_1) = \bar{R}_1(0, \tau_0, \tau_1) = R(\delta_0, \theta_0^*) = R(\delta_0, \theta_1^*)$, $r(\delta_0, \pi^*) = R(\delta_0, \theta_0^*) = R(\delta_0, \theta_1^*) \geq R(\delta_0, \theta)$ for all $\theta \in [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1] = \Theta$.

Since the class of threshold rules is essentially complete, $\delta_0 \in \arg\min_{\delta \in \mathcal{D}} r(\delta, \pi^*)$ if $0 \in \arg\min_{t \in \mathbb{R}} r(\delta_t, \pi^*)$. Observe

$$
\begin{aligned}
r(\delta_t, \pi^*) &= \left[R(\delta_t, \theta_0^*) + R(\delta_t, \theta_1^*)\right]/2 \\
&= \left[-\theta_0^*(1 - \Phi(\sigma^{-1}(t - \theta_0^*))) + \theta_1^* \Phi(\sigma^{-1}(t - \theta_1^*))\right]/2, \\
\frac{\partial r(\delta_t, \pi^*)}{\partial t} &= \sigma^{-1}\phi(\sigma^{-1}(t - \theta_0^*)) \left[\theta_0^* + \theta_1^* \frac{\phi(\sigma^{-1}(t - \theta_1^*))}{\phi(\sigma^{-1}(t - \theta_0^*))}\right]/2.
\end{aligned}
$$

Since $\frac{\phi(\sigma^{-1}(t-\theta_1^*))}{\phi(\sigma^{-1}(t-\theta_0^*))}$ is nondecreasing in $t$ by the monotone likelihood ratio property and $\theta_0^* + \theta_1^* \frac{\phi(\sigma^{-1}(t-\theta_1^*))}{\phi(\sigma^{-1}(t-\theta_0^*))} = 0$ for $t = 0$ by construction, it follows that $\frac{\partial r(\delta_t, \pi^*)}{\partial t} \geq 0$ if $t > 0$, $\frac{\partial r(\delta_t, \pi^*)}{\partial t} = 0$ if $t = 0$, and $\frac{\partial r(\delta_t, \pi^*)}{\partial t} \leq 0$ if $t < 0$. Therefore, $0 \in \arg\min_{t \in \mathbb{R}} r(\delta_t, \pi^*)$. Thus, $\delta_0$ is minimax regret.

The minimax risk is given by $\mathcal{R}_{\text{uni}}(\sigma; [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1]) = \max_{\theta \in [\tau_0, \tau_1]} \theta \Phi(-\theta/\sigma) = \max_{a \in [\tau_0/\sigma, \tau_1/\sigma]} \sigma a \Phi(-a)$. Let $g(a) = a\Phi(-a)$. We have $g'(a) = \Phi(-a) - a\phi(-a) = \left(\frac{1-\Phi(a)}{\phi(a)} - a\right)\phi(a)$. By the fact that the Mills ratio $\frac{1-\Phi(a)}{\phi(a)}$ of a standard normal random variable is strictly decreasing and continuous, $g'(a) > 0$ for $a \in [0, a^*)$, $g'(a^*) = 0$, and $g'(a) < 0$ for $a > a^*$, where $a^* > 0$

uniquely solves $\frac{1-\Phi(a)}{\phi(a)} = a$. The expression for $\mathcal{R}_{\text{uni}}(\sigma; [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1])$ then follows from the fact that $g(a)$ is strictly increasing on $[0, a^*)$ and strictly decreasing on $(a^*, \infty)$. $\square$

## A.2 Proof of Lemma 1

Note first that if $L(\bar{\theta}) = 0$, then $L(\theta) = 0$ for all $\theta \in [-\bar{\theta}, \bar{\theta}]$ by the linearity of $L$, and therefore any decision rule is minimax regret with maximum regret equal to zero.

Consider the case where $L(\bar{\theta}) > 0$ and $\boldsymbol{m}(\bar{\theta}) = \mathbf{0}$. Since $\boldsymbol{Y} \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I}_n)$ under $\theta$ for all $\theta \in [-\bar{\theta}, \bar{\theta}]$, $\mathbb{E}_\theta[\delta(\boldsymbol{Y})] = \mathbb{E}[\delta(\boldsymbol{Y}^*)]$, where $\boldsymbol{Y}^* \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I}_n)$, for all $\theta \in [-\bar{\theta}, \bar{\theta}]$. Therefore, the maximum regret of decision rule $\delta$ is

$$\sup_{\theta \in [-\bar{\theta}, \bar{\theta}]} R(\delta, \theta) = \begin{cases} L(\bar{\theta})(1 - \mathbb{E}[\delta(\boldsymbol{Y}^*)]) & \text{if } \mathbb{E}[\delta(\boldsymbol{Y}^*)] < 1/2, \\ L(\bar{\theta})/2 & \text{if } \mathbb{E}[\delta(\boldsymbol{Y}^*)] = 1/2, \\ (-L(-\bar{\theta}))\mathbb{E}[\delta(\boldsymbol{Y}^*)] & \text{if } \mathbb{E}[\delta(\boldsymbol{Y}^*)] > 1/2. \end{cases}$$

Thus, any decision rule $\delta^*$ such that $\mathbb{E}[\delta^*(\boldsymbol{Y}^*)] = \frac{1}{2}$ is minimax regret. The minimax risk is given by $\mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}]) = \frac{L(\bar{\theta})}{2}$.

Below, I prove the following result, which covers the result for the case where $L(\bar{\theta}) > 0$ and $\boldsymbol{m}(\bar{\theta}) \neq \mathbf{0}$ in Lemma 1 as a special case.

**Lemma A.8** (Minimax Rule for Informative One-dimensional Subproblems)**.** *Suppose that $\Theta = [-\bar{\theta}, -t\bar{\theta}] \cup [t\bar{\theta}, \bar{\theta}]$, where $\bar{\theta} \in \mathbb{V}$, $L(\bar{\theta}) > 0$, $\boldsymbol{m}(\bar{\theta}) \neq \mathbf{0}$, and $t \in [0, 1]$. Then, the decision rule $\delta^*(\boldsymbol{Y}) = \mathbf{1}\left\{\boldsymbol{m}(\bar{\theta})'\boldsymbol{Y} \geq 0\right\}$ is minimax regret. The minimax risk is given by $\mathcal{R}(\sigma; [-\bar{\theta}, -t\bar{\theta}] \cup [t\bar{\theta}, \bar{\theta}]) = \frac{L(\bar{\theta})}{\|\boldsymbol{m}(\bar{\theta})\|}\mathcal{R}_{\text{uni}}(\sigma; [-\|\boldsymbol{m}(\bar{\theta})\|, -t\|\boldsymbol{m}(\bar{\theta})\|] \cup [t\|\boldsymbol{m}(\bar{\theta})\|, \|\boldsymbol{m}(\bar{\theta})\|])$, where $\mathcal{R}_{\text{uni}}(\sigma; [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1])$ is given in Lemma A.7.*

*Proof.* Fix $\bar{\theta} \in \mathbb{V}$, where $L(\bar{\theta}) > 0$ and $\boldsymbol{m}(\bar{\theta}) \neq \mathbf{0}$, and $t \in [0, 1]$. We can write $[-\bar{\theta}, -t\bar{\theta}] \cup [t\bar{\theta}, \bar{\theta}] = \{\lambda\bar{\theta} : \lambda \in [-1, -t] \cup [t, 1]\}$. For $\lambda \in [-1, -t] \cup [t, 1]$, regret of decision rule $\delta$ under $\lambda\bar{\theta}$ equals $R(\delta, \lambda\bar{\theta}) = (L(\lambda\bar{\theta}))^+(1 - \mathbb{E}_{\lambda\bar{\theta}}[\delta(\boldsymbol{Y})]) + (-L(\lambda\bar{\theta}))^+\mathbb{E}_{\lambda\bar{\theta}}[\delta(\boldsymbol{Y})] = L(\bar{\theta})\left(\lambda^+(1 - \mathbb{E}_{\lambda\bar{\theta}}[\delta(\boldsymbol{Y})]) + (-\lambda)^+\mathbb{E}_{\lambda\bar{\theta}}[\delta(\boldsymbol{Y})]\right)$, where $x^+ = \max\{x, 0\}$. Minimax regret decision rules thus solve $\inf_\delta \sup_{\lambda \in [-1,-t] \cup [t,1]} \left(\lambda^+(1 - \mathbb{E}_{\lambda\bar{\theta}}[\delta(\boldsymbol{Y})]) + (-\lambda)^+\mathbb{E}_{\lambda\bar{\theta}}[\delta(\boldsymbol{Y})]\right)$.

Viewing $\lambda$ as a parameter, I derive a sufficient statistic of $\boldsymbol{Y}$ for $\lambda$. Let $h(\boldsymbol{y}) = \frac{1}{\sqrt{(2\pi)^n\sigma^n}} \exp(-\frac{1}{2\sigma^2}\|\boldsymbol{y}\|^2)$, $g(t,\lambda) = \exp(-\frac{1}{2\sigma^2}(-2\lambda t + \lambda^2)\|\boldsymbol{m}(\bar{\theta})\|^2)$, and $T(\boldsymbol{y}) = \frac{\boldsymbol{m}(\bar{\theta})'\boldsymbol{y}}{\|\boldsymbol{m}(\bar{\theta})\|^2}$. For $\lambda \in [-1,-t] \cup [t,1]$, $\boldsymbol{Y} \sim \mathcal{N}(\lambda \boldsymbol{m}(\bar{\theta}), \sigma^2 \boldsymbol{I}_n)$ under $\lambda\bar{\theta}$. It follows that the probability density of $\boldsymbol{Y}$ is $p(\boldsymbol{y}) = \frac{1}{\sqrt{(2\pi)^n\sigma^n}} \exp\left(-\frac{1}{2\sigma^2}\|\boldsymbol{y} - \lambda\boldsymbol{m}(\bar{\theta})\|^2\right) = \frac{1}{\sqrt{(2\pi)^n\sigma^n}} \exp\left(-\frac{1}{2\sigma^2}(\|\boldsymbol{y}\|^2 - 2\lambda\boldsymbol{m}(\bar{\theta})'\boldsymbol{y} + \lambda^2\|\boldsymbol{m}(\bar{\theta})\|^2)\right) = h(\boldsymbol{y})g(T(\boldsymbol{y}),\lambda)$. By the factorization theorem, $T(\boldsymbol{Y})$ is a sufficient statistic for $\lambda$.

By Theorem 1 in Chapter 1 of Berger (1985), the class of decision rules that only depend on $T(\boldsymbol{Y})$ is essentially complete. Since $T(\boldsymbol{Y}) \sim \mathcal{N}(\lambda, \frac{\sigma^2}{\|\boldsymbol{m}(\bar{\theta})\|^2})$ under $\lambda\bar{\theta}$, minimax regret decision rules that only depend on $T(\boldsymbol{Y})$ solve $\inf_\delta \sup_{\lambda\in[-1,-t]\cup[t,1]} (\lambda^+(1 - \mathbb{E}_\lambda[\delta(T)]) + (-\lambda)^+\mathbb{E}_\lambda[\delta(T)])$, where the expectation is taken with respect to $T \sim \mathcal{N}(\lambda, \frac{\sigma^2}{\|\boldsymbol{m}(\bar{\theta})\|^2})$. This problem is equivalent to the univariate problem in Lemma A.7 where $\Theta = [-1,-t] \cup [t,1]$, $\boldsymbol{m}(\theta) = \theta$, $L(\theta) = \theta$, and the variance of the observed normal random variable is $\frac{\sigma^2}{\|\boldsymbol{m}(\bar{\theta})\|^2}$. Thus, by Lemma A.7, the decision rule $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{T(\boldsymbol{Y}) \geq 0\} = \mathbf{1}\left\{\boldsymbol{m}(\bar{\theta})'\boldsymbol{Y} \geq 0\right\}$ is minimax regret. The minimax risk is given by $\mathcal{R}(\sigma; [-\bar{\theta}, -t\bar{\theta}] \cup [t\bar{\theta}, \bar{\theta}]) = L(\bar{\theta})\mathcal{R}_{\text{uni}}\left(\frac{\sigma}{\|\boldsymbol{m}(\bar{\theta})\|}; [-1,-t] \cup [t,1]\right) = \frac{L(\bar{\theta})}{\|\boldsymbol{m}(\bar{\theta})\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{m}(\bar{\theta})\| - t\|\boldsymbol{m}(\bar{\theta})\|] \cup [t\|\boldsymbol{m}(\bar{\theta})\|, \|\boldsymbol{m}(\bar{\theta})\|]\right)$, where the second equality follows from the fact that $\mathcal{R}_{\text{uni}}(\alpha\sigma; [-\alpha\tau_1, -\alpha\tau_0] \cup [\alpha\tau_0, \alpha\tau_1]) = \alpha\mathcal{R}_{\text{uni}}(\sigma; [-\tau_1, -\tau_0] \cup [\tau_0, \tau_1])$ for all $\alpha > 0$. ☐

## A.3 Proof of Theorem 1

The following lemma proves part (i).

**Lemma A.9.** *Suppose that $\omega(\cdot)$ is differentiable at any $\epsilon \in [0, a^*\sigma]$. Then, there exists a unique solution to the maximization problem $\max_{\epsilon\in[0,a^*\sigma]} \omega(\epsilon)\Phi(-\epsilon/\sigma)$. The solution is nonzero if and only if $\sigma > 2\phi(0)\frac{\omega(0)}{\omega'(0)}$.*

*Proof.* See Appendix A.4.1. ☐

Below, I prove part (ii). I provide separate arguments for the nonrandomized and randomized rules. For each case, I first state an assumption weaker than the conditions of Theorem 1, present a result under the relaxed assumption, and then provide the proof for the more general result.

### A.3.1 Nonrandomized Rule

Consider the following assumption, which holds if $\sigma > 2\phi(0)\frac{\omega(0)}{\omega'(0)}$ under Assumption 1(ii) by Lemma A.9.

**Assumption A.1** (Informative Worst Case)**.** *There exists a unique, nonzero solution to the maximization problem* $\max_{\epsilon\in[0,a^*\sigma]}\omega(\epsilon)\Phi(-\epsilon/\sigma)$.

I obtain a minimax regret rule under Assumption A.1. The statement on the nonrandomized rule in Theorem 1 immediately follows from the result below.

**Theorem A.1** (Nonrandomized Minimax Regret Rule)**.** *Let $\Theta$ be convex and centrosymmetric, and suppose that Assumption A.1 holds. Let* $\epsilon^* \in \arg\max_{\epsilon\in[0,a^*\sigma]}\omega(\epsilon)\Phi(-\epsilon/\sigma)$*, and suppose that there exists $\theta_{\epsilon^*}$ that attains the modulus of continuity at $\epsilon^*$. Then, the decision rule $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y} \geq 0\}$ is minimax regret. Here, $\boldsymbol{m}(\theta_{\epsilon^*})$ does not depend on the choice of $\theta_{\epsilon^*}$ among those that attain the modulus of continuity at $\epsilon^*$, and $\|\boldsymbol{m}(\theta_{\epsilon^*})\| = \epsilon^*$. The minimax risk is given by* $\mathcal{R}(\sigma;\Theta) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$.

Below, I provide the proof of Theorem A.1. The lemma below shows that $\boldsymbol{m}(\theta_{\epsilon^*})$ does not depend on the choice of $\theta_{\epsilon^*}$ among those that attain the modulus of continuity at $\epsilon^*$ and $\|\boldsymbol{m}(\theta_{\epsilon^*})\| = \epsilon^*$.

**Lemma A.10.** *Suppose that the conditions of Theorem A.1 hold, and let* $\Theta_{\epsilon^*} = \arg\max_{\theta\in\Theta:\|\boldsymbol{m}(\theta)\|\leq\epsilon^*} L(\theta)$*. Then, $\|\boldsymbol{m}(\theta)\| = \epsilon^*$ for any $\theta \in \Theta_{\epsilon^*}$, and $\boldsymbol{m}(\theta) = \boldsymbol{m}(\tilde{\theta})$ for any $\theta, \tilde{\theta} \in \Theta_{\epsilon^*}$.*

*Proof.* See Appendix A.4.2. □

Next, I characterize the supremum of the minimax risk $\mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}])$ over all one-dimensional subfamilies of the form $[-\bar{\theta},\bar{\theta}]$, where $\bar{\theta}\in\Theta$, $L(\bar{\theta}) > 0$, and $\boldsymbol{m}(\bar{\theta}) \neq \mathbf{0}$.

**Lemma A.11.** *Under the conditions of Theorem A.1,* $\mathcal{R}(\sigma;[-\theta_{\epsilon^*},\theta_{\epsilon^*}]) = \sup_{\bar{\theta}\in\Theta:L(\bar{\theta})>0,\boldsymbol{m}(\bar{\theta})\neq\mathbf{0}}\mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}]) = \sup_{\epsilon>0}\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}(\sigma;[-\epsilon,\epsilon]) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$*, where $\mathcal{R}_{\text{uni}}(\sigma;[-\epsilon,\epsilon])$ is given by Lemma A.7.*

*Proof.* See Appendix A.4.3. □

By Lemma 1, the decision rule $\delta^*(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y} \geq 0\}$ is minimax regret for the one-dimensional subproblem $[-\theta_{\epsilon^*}, \theta_{\epsilon^*}]$. Since $\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y} \sim \mathcal{N}(\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta), \sigma^2\|\boldsymbol{m}(\theta_{\epsilon^*})\|^2)$ under $\theta$, the maximum regret of $\delta^*$ over $\Theta$ is given by

$$
\begin{aligned}
&\sup_{\theta\in\Theta} R(\delta^*, \theta) \\
&= \sup_{\theta\in\Theta}\left[(L(\theta))^+\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right) + (-L(\theta))^+\left(1-\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)\right)\right] \\
&= \sup_{\theta\in\Theta:L(\theta)>0} L(\theta)\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right),
\end{aligned}
$$

where $x^+ = \max\{x, 0\}$ and the second equality holds by the symmetry of the objective function and the centrosymmetry of $\Theta$.

The following lemma is fundamental to characterizing minimax regret rules for the original problem.

**Lemma A.12** (Worst Case for Nonrandomized Rule)**.** *Under the conditions of Theorem A.1,* $\theta_{\epsilon^*} \in \arg\max_{\theta\in\Theta:L(\theta)>0} L(\theta)\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$.

*Proof.* See Appendix A.4.4. □

By Lemma A.12, the maximum regret of the decision rule $\delta^*$ over $\Theta$ is attained at $\theta_{\epsilon^*}$. Therefore, $\delta^*$ and $[-\theta_{\epsilon^*}, \theta_{\epsilon^*}]$ satisfy Property 1 in Section 3.1.2, and $\delta^*$ is minimax regret for $\Theta$. The minimax risk is $\mathcal{R}(\sigma; [-\theta_{\epsilon^*}, \theta_{\epsilon^*}]) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$.

### A.3.2 Randomized Rule

Consider the following assumption. I adopt the convention that $x/\infty = 0$ for all $x \in \mathbb{R}$.

**Assumption A.2** (Regularity for Randomized Rule)**.**

*(i) For some $\bar{\epsilon} > 0$, there exists $\{\theta_\epsilon\}_{\epsilon\in[0,\bar{\epsilon}]}$ with $\theta_\epsilon \in \Theta$ such that the following holds.*

*(a) For all $\epsilon \in [0, \bar{\epsilon}]$, $\theta_\epsilon$ attains the modulus of continuity at $\epsilon$ with $\|\boldsymbol{m}(\theta_\epsilon)\| = \epsilon$.*

*(b) There exists $\boldsymbol{w}^* \in \mathbb{R}^n$ such that $\lim_{\epsilon\to 0}\epsilon^{-1}\left(\boldsymbol{w}^* - \frac{\boldsymbol{m}(\theta_\epsilon)}{\|\boldsymbol{m}(\theta_\epsilon)\|}\right) = \mathbf{0}$.*

*(c) There exists $\sigma^* \in [\sigma, \infty]$ such that $0 \in \arg\max_{\epsilon \in \mathbb{R}} \rho(\epsilon)\Phi\left(-\epsilon/\sigma^*\right)$, where $\rho(\epsilon) \coloneqq \sup\left\{L(\theta) : (\boldsymbol{w}^*)'\boldsymbol{m}(\theta) = \epsilon, \theta \in \Theta\right\}$ for $\epsilon \in \mathbb{R}$.*[c]

*(ii) $\omega(\cdot)$ is continuous at $\epsilon = 0$.*

Assumption A.2(i)(a) is slightly stronger than Assumption 1(i)(a) since it requires that the constraint $\|\boldsymbol{m}(\theta_\epsilon)\| \leq \epsilon$ of the modulus problem hold with equality. Assumption A.2(i)(b) is the same as Assumption 1(i)(b). Assumption A.2(ii) is implied by Assumption 1(ii).

Given $\sigma^*$, the maximization problem $\max_{\epsilon \in \mathbb{R}} \rho(\epsilon)\Phi\left(-\epsilon/\sigma^*\right)$ in Assumption A.2(i)(c) corresponds to the problem of finding the worst-case parameter values for a randomized decision rule $\delta(\boldsymbol{Y}) = \mathbb{P}\left((\boldsymbol{w}^*)'\boldsymbol{Y} + \xi \geq 0 | \boldsymbol{Y}\right)$, where $\xi | \boldsymbol{Y} \sim \mathcal{N}(0, (\sigma^*)^2 - \sigma^2)$. I will later show that under Assumption A.2(i)(c), we can find the variance of $\xi$ such that the maximum regret of $\delta$ is attained at a value of $\theta$ that attains the modulus of continuity at $\epsilon = 0$.

The lemma below shows that Assumption 1 implies Assumption A.2 if $\sigma \leq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$.

**Lemma A.13.** *Let $\Theta$ be convex and centrosymmetric. Suppose that Assumption 1 holds, and let $\{\theta_\epsilon\}_{\epsilon \in [0, \bar{\epsilon}]}$, $\iota$, and $\boldsymbol{w}^*$ satisfy Assumption 1(i). Then, $\|\boldsymbol{m}(\theta_\epsilon)\| = \epsilon$ for all $\epsilon \in [0, \bar{\epsilon}]$ and $\omega'(0) = \rho'(0) = \frac{L(\iota)}{\boldsymbol{m}(\iota)'\boldsymbol{w}^*} > 0$. Suppose in addition that $2\phi(0)\frac{\omega(0)}{\omega'(0)} > 0$. Then, $0 \in \arg\max_{\epsilon \in \mathbb{R}} \rho(\epsilon)\Phi\left(-\epsilon/\sigma^*\right)$ with $\sigma^* = 2\phi(0)\frac{\omega(0)}{\omega'(0)}$.*

*Proof.* See Appendix A.4.5. □

I obtain a minimax regret rule under Assumption A.2. The statement on the randomized rule in Theorem 1 immediately follows from the following result.

**Theorem A.2** (Randomized Minimax Regret Rule)**.** *Let $\Theta$ be convex and centrosymmetric, and suppose that Assumption A.2 holds. Then, the following decision rule is minimax regret:*

$$\delta^*(\boldsymbol{Y}) = \begin{cases} \mathbf{1}\{(\boldsymbol{w}^*)'\boldsymbol{Y} \geq 0\} & \text{if } \sigma^* = \sigma, \\ \Phi\left(\dfrac{(\boldsymbol{w}^*)'\boldsymbol{Y}}{((\sigma^*)^2 - \sigma^2)^{1/2}}\right) & \text{if } \sigma^* > \sigma. \end{cases}$$

[c] I allow the search space of $\sigma^*$ to contain $\infty$, letting $\Phi(x/\infty) = 1/2$ for all $x \in \mathbb{R}$. Assumption A.2(i)(c) then holds with $\sigma^* = \infty$ in Stoye (2012)'s setup described in Section 2.3.2 when $b \geq 1$.

*The minimax risk is given by $\mathcal{R}(\sigma;\Theta) = \omega(0)/2$.*

Note that if $\omega(\cdot)$ is differentiable at any $\epsilon \in [0, a^*\sigma]$ and $\sigma \le 2\phi(0)\frac{\omega(0)}{\omega'(0)}$, $\epsilon^* = 0$ by Lemma A.9, where $\epsilon^* \in \arg\max_{\epsilon \in [0,a^*\sigma]} \omega(\epsilon)\Phi(-\epsilon/\sigma)$. The minimax risk $\omega(0)/2$ can then be written as $\omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$, leading to the expression in Theorem 1.

Below, I prove Theorem A.2 by showing that $\delta^*$ and $[-\theta_0, \theta_0]$ satisfy Property 1 in Section 3.1.2, where $\theta_0$ attains the modulus of continuity at 0 and hence $\boldsymbol{m}(\theta_0) = \mathbf{0}$. Note first that $\boldsymbol{w}^*$ is a unit vector by construction. Hence, $(\boldsymbol{w}^*)'\boldsymbol{Y} \sim \mathcal{N}\left((\boldsymbol{w}^*)'\boldsymbol{m}(\theta), \sigma^2\right)$. Simple calculations show that we can write $\delta^*(\boldsymbol{Y}) = \mathbb{P}\left((\boldsymbol{w}^*)'\boldsymbol{Y} + \xi \ge 0 | \boldsymbol{Y}\right)$, where $\xi | \boldsymbol{Y} \sim \mathcal{N}(0, (\sigma^*)^2 - \sigma^2)$. Since $(\boldsymbol{w}^*)'\boldsymbol{Y} + \xi \sim \mathcal{N}\left(0, (\sigma^*)^2\right)$ if $\boldsymbol{Y} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I}_n)$, it follows that $\mathbb{E}[\delta^*(\boldsymbol{Y}^*)] = \frac{1}{2}$, where $\boldsymbol{Y}^* \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I}_n)$. By Lemma 1, $\delta^*$ is minimax regret for the one-dimensional subproblem $[-\theta_0, \theta_0]$.

Since $(\boldsymbol{w}^*)'\boldsymbol{Y} + \xi \sim \mathcal{N}\left((\boldsymbol{w}^*)'\boldsymbol{m}(\theta), (\sigma^*)^2\right)$ under $\theta$, the maximum regret of $\delta^*$ over $\Theta$ is given by

$$
\begin{aligned}
&\sup_{\theta \in \Theta} R(\delta^*, \theta) \\
&= \sup_{\theta \in \Theta} \left[ (L(\theta))^+ \Phi\left( -\frac{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta)}{\sigma^*} \right) + (-L(\theta))^+ \left( 1 - \Phi\left( -\frac{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta)}{\sigma^*} \right) \right) \right] \\
&= \sup_{\theta \in \Theta} L(\theta) \Phi\left( -\frac{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta)}{\sigma^*} \right),
\end{aligned}
$$

where the second equality holds by the symmetry of the objective function and the centrosymmetry of $\Theta$. The following lemma shows that the maximum regret is attained at $\theta_0$, and hence Property 1 is satisfied.

**Lemma A.14** (Worst Case for Randomized Rule)**.** *Under the conditions of Theorem A.2, $\theta_0 \in \arg\max_{\theta \in \Theta} L(\theta)\Phi\left(-\frac{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta)}{\sigma^*}\right)$.*

*Proof.* See Appendix A.4.6. □

## A.4 Proofs of Lemmas in Appendix A.3

### A.4.1 Proof of Lemma A.9

Let $g(\epsilon) = \omega(\epsilon)\Phi(-\epsilon/\sigma)$. Suppose that $\omega'(0) = 0$, in which case $\sigma \le 2\phi(0)\frac{\omega(0)}{\omega'(0)}$. Since $\omega(\epsilon)$ is nondecreasing and concave, $\omega(\epsilon)$ is constant. Then $g$ is uniquely maximized at 0 over $[0, a^*\sigma]$.

Suppose that $\omega'(0) > 0$. $\omega(\epsilon)$ is nonnegative, nonconstant, nondecreasing, concave, and differentiable on $[0, a^*\sigma]$. By Lemma A.1, if $\sigma \frac{1-\Phi\left(\frac{0}{\sigma}\right)}{\phi\left(\frac{0}{\sigma}\right)} > \frac{\omega(0)}{\omega'(0)}$, equivalently $\sigma > 2\phi(0)\frac{\omega(0)}{\omega'(0)}$, then $g$ is uniquely maximized at some $\epsilon \in (0, a^*\sigma]$. If $\sigma \frac{1-\Phi\left(\frac{0}{\sigma}\right)}{\phi\left(\frac{0}{\sigma}\right)} \leq \frac{\omega(0)}{\omega'(0)}$, equivalently $\sigma \leq 2\phi(0)\frac{\omega(0)}{\omega'(0)}$, then $g$ is uniquely maximized at 0 over $[0, a^*\sigma]$.

### A.4.2 Proof of Lemma A.10

First, pick any $\theta \in \Theta_{\epsilon^*}$. Since $\theta$ attains the modulus of continuity at $\epsilon^*$, it also attains the modulus at $\|\boldsymbol{m}(\theta)\|$, so that $\omega(\|\boldsymbol{m}(\theta)\|) = \omega(\epsilon^*)$. It follows that $\|\boldsymbol{m}(\theta)\| = \epsilon^*$, since if $\|\boldsymbol{m}(\theta)\| < \epsilon^*$, $\omega(\|\boldsymbol{m}(\theta)\|)\Phi(-\|\boldsymbol{m}(\theta)\|/\sigma) = \omega(\epsilon^*)\Phi(-\|\boldsymbol{m}(\theta)\|/\sigma) > \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$, which contradicts the assumption that $\epsilon^*$ maximizes $\omega(\epsilon)\Phi(-\epsilon/\sigma)$ over $[0, a^*\sigma]$.

Now, pick any $\theta, \tilde{\theta} \in \Theta_{\epsilon^*}$. By the above argument, $\|\boldsymbol{m}(\theta)\| = \|\boldsymbol{m}(\tilde{\theta})\| = \epsilon^*$, and hence $\boldsymbol{m}(\theta), \boldsymbol{m}(\tilde{\theta}) \in \{\boldsymbol{\beta} \in \mathbb{R}^n : \|\boldsymbol{\beta}\| \leq \epsilon^*\}$. Suppose that $\boldsymbol{m}(\theta) \neq \boldsymbol{m}(\tilde{\theta})$, and let $\bar{\theta} = \lambda\theta + (1-\lambda)\tilde{\theta}$ for some $\lambda \in (0,1)$. Then $L(\bar{\theta}) = \lambda L(\theta) + (1-\lambda)L(\tilde{\theta}) = \omega(\epsilon^*)$. By the convexity of $\Theta$, $\bar{\theta} \in \Theta$. Furthermore, since $\{\boldsymbol{\beta} \in \mathbb{R}^n : \|\boldsymbol{\beta}\| \leq \epsilon^*\}$ is strictly convex, $\boldsymbol{m}(\bar{\theta}) = \lambda\boldsymbol{m}(\theta) + (1-\lambda)\boldsymbol{m}(\tilde{\theta})$ is an interior point of $\{\boldsymbol{\beta} \in \mathbb{R}^n : \|\boldsymbol{\beta}\| \leq \epsilon^*\}$, which implies that $\|\boldsymbol{m}(\bar{\theta})\| < \epsilon^*$. Thus, $\bar{\theta}$ attains the modulus at $\epsilon^*$, but then it must be the case that $\|\boldsymbol{m}(\bar{\theta})\| = \epsilon^*$.

### A.4.3 Proof of Lemma A.11

By Lemma A.8,

$$\sup_{\bar{\theta}\in\Theta: L(\bar{\theta})>0, \boldsymbol{m}(\bar{\theta})\neq\boldsymbol{0}} \mathcal{R}(\sigma; [-\bar{\theta}, \bar{\theta}]) = \sup_{\bar{\theta}\in\Theta: L(\bar{\theta})>0, \boldsymbol{m}(\bar{\theta})\neq\boldsymbol{0}} \frac{L(\bar{\theta})}{\|\boldsymbol{m}(\bar{\theta})\|}\mathcal{R}_{\text{uni}}(\sigma; [-\|\boldsymbol{m}(\bar{\theta})\|, \|\boldsymbol{m}(\bar{\theta})\|])$$
$$= \sup_{\epsilon>0}\left\{\sup_{\bar{\theta}\in\Theta:\|\boldsymbol{m}(\bar{\theta})\|=\epsilon} \frac{L(\bar{\theta})}{\|\boldsymbol{m}(\bar{\theta})\|}\mathcal{R}_{\text{uni}}(\sigma; [-\|\boldsymbol{m}(\bar{\theta})\|, \|\boldsymbol{m}(\bar{\theta})\|])\right\}$$
$$= \sup_{\epsilon>0}\left\{\frac{\sup_{\bar{\theta}\in\Theta:\|\boldsymbol{m}(\bar{\theta})\|=\epsilon} L(\bar{\theta})}{\epsilon}\mathcal{R}_{\text{uni}}(\sigma; [-\epsilon, \epsilon])\right\} \leq \sup_{\epsilon>0}\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}(\sigma; [-\epsilon, \epsilon]),$$

where the last inequality holds by the definition of $\omega(\epsilon)$. By Lemma A.7,

$$\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}(\sigma;[-\epsilon,\epsilon]) = \begin{cases} \omega(\epsilon)\Phi(-\epsilon/\sigma) & \text{if } \epsilon \le a^*\sigma, \\ \frac{\omega(\epsilon)}{\epsilon}a^*\sigma\Phi(-a^*) & \text{if } \epsilon > a^*\sigma. \end{cases}$$

Since $\omega(\epsilon)$ is concave, $\frac{\omega(\epsilon)}{\epsilon}$ is continuous and nonincreasing on $(0,\infty)$, so that $\sup_{\epsilon>a^*\sigma}\frac{\omega(\epsilon)}{\epsilon}a^*\sigma\Phi(-a^*) = \omega(a^*\sigma)\Phi(-a^*)$. Therefore, $\sup_{\epsilon>0}\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}(\sigma;[-\epsilon,\epsilon]) = \sup_{0<\epsilon\le a^*\sigma}\omega(\epsilon)\Phi(-\epsilon/\sigma) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$. Hence, $\sup_{\bar{\theta}\in\Theta:L(\bar{\theta})>0,\boldsymbol{m}(\bar{\theta})\neq\boldsymbol{0}}\mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}]) \le \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$.

On the other hand, by Lemmas A.8 and A.10, $\mathcal{R}(\sigma;[-\theta_{\epsilon^*},\theta_{\epsilon^*}]) = \frac{L(\theta_{\epsilon^*})}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\mathcal{R}_{\text{uni}}(\sigma;[-\|\boldsymbol{m}(\theta_{\epsilon^*})\|,\|\boldsymbol{m}(\theta_{\epsilon^*})\|]) = \frac{\omega(\epsilon^*)}{\epsilon^*}\mathcal{R}_{\text{uni}}(\sigma;[-\epsilon^*,\epsilon^*]) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$, where the last equality follows from Lemma A.7 and the fact that $\epsilon^* \le a^*\sigma$. Since $\mathcal{R}(\sigma;[-\theta_{\epsilon^*},\theta_{\epsilon^*}]) \le \sup_{\bar{\theta}\in\Theta:L(\bar{\theta})>0,\boldsymbol{m}(\bar{\theta})\neq\boldsymbol{0}}\mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}])$, it follows that $\mathcal{R}(\sigma;[-\theta_{\epsilon^*},\theta_{\epsilon^*}]) = \sup_{\bar{\theta}\in\Theta:L(\bar{\theta})>0,\boldsymbol{m}(\bar{\theta})\neq\boldsymbol{0}}\mathcal{R}(\sigma;[-\bar{\theta},\bar{\theta}])$.

### A.4.4 Proof of Lemma A.12

Let $\epsilon^*$ be the unique nonzero solution to $\max_{0\le\epsilon\le a^*\sigma}\omega(\epsilon)\Phi(-\epsilon/\sigma)$, and let $\theta_{\epsilon^*}$ attain the modulus of continuity at $\epsilon^*$. By Lemma A.10, $\|\boldsymbol{m}(\theta_{\epsilon^*})\| = \epsilon^*$. I first introduce some notation. Pick any $\underline{\eta} \in (0,\min\{\omega(\epsilon^*),\epsilon^*\})$ and any $\bar{\epsilon} > \epsilon^*$. Since $\omega(\epsilon^*)$ is finite, $\omega(\bar{\epsilon})$ is also finite by the convexity of $\omega(\epsilon)$. Define $\Theta_{+,\underline{\eta},\bar{\epsilon}} = \left\{\theta\in\Theta : L(\theta)\ge\underline{\eta}, \frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\ge\underline{\eta}, \|\boldsymbol{m}(\theta)\|\le\bar{\epsilon}\right\}$,

$$\begin{aligned}\Theta_{\underline{\eta},\bar{\epsilon}} &= \{\theta\in\Theta : (\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}) \text{ or } (-\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}})\} \\ &= \Bigg\{\theta\in\Theta : \Bigg[\left(L(\theta)\ge\underline{\eta}, \frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\ge\underline{\eta}\right) \\ &\qquad \text{or } \left(L(\theta)\le-\underline{\eta}, \frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\le-\underline{\eta}\right)\Bigg] \text{ and } \|\boldsymbol{m}(\theta)\|\le\bar{\epsilon}\Bigg\},\end{aligned}$$

and $\Gamma_{+,\underline{\eta},\bar{\epsilon}} = \{(L(\theta),\boldsymbol{m}(\theta)')'\in\mathbb{R}^{n+1} : \theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}\}$. Note that $\Gamma_{+,\underline{\eta},\bar{\epsilon}}$ is bounded, since $\underline{\eta}\le\alpha\le\omega(\bar{\epsilon})$ and $\|\boldsymbol{\beta}\|\le\bar{\epsilon}$ for all $\boldsymbol{\gamma} = (\alpha,\boldsymbol{\beta}')'\in\Gamma_{+,\underline{\eta},\bar{\epsilon}}$.

Now, let $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ denote the closure of $\Gamma_{+,\underline{\eta},\bar{\epsilon}}$. Define a set-valued function $\Psi$ :

$\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}} \to 2^{\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}}$ as follows: for $\boldsymbol{\gamma} = (\alpha, \boldsymbol{\beta}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$,

$$\Psi(\boldsymbol{\gamma}) = \arg \max_{\bar{\boldsymbol{\gamma}}=(\bar{\alpha},\bar{\boldsymbol{\beta}}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}} \bar{\alpha}\Phi\left(-\frac{\boldsymbol{\beta}'\bar{\boldsymbol{\beta}}}{\sigma\|\boldsymbol{\beta}\|}\right).$$

Note that $\alpha > 0$ and $\boldsymbol{\beta} \neq 0$ for all $(\alpha, \boldsymbol{\beta}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$, since $\tilde{\alpha} \geq \underline{\eta} > 0$ and $\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\tilde{\boldsymbol{\beta}}}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} \geq \underline{\eta} > 0$ for all $\tilde{\boldsymbol{\gamma}} = (\tilde{\alpha}, \tilde{\boldsymbol{\beta}}')' \in \Gamma_{+,\underline{\eta},\bar{\epsilon}}$, and $(\alpha, \boldsymbol{\beta}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ is a point or a limit point of $\Gamma_{+,\underline{\eta},\bar{\epsilon}}$.

The proof consists of six steps.

**Step 1.** *$\Psi$ has a fixed point, i.e., there exists $\boldsymbol{\gamma} \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ such that $\boldsymbol{\gamma} \in \Psi(\boldsymbol{\gamma})$.*

*Proof.* I apply Kakutani's fixed point theorem. First of all, $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ is nonempty, since $(L(\theta_{\epsilon^*}), \boldsymbol{m}(\theta_{\epsilon^*}))')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$. Furthermore, $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ is closed and bounded by construction.

I show that $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ is convex. It suffices to show that $\Gamma_{+,\underline{\eta},\bar{\epsilon}}$ is convex, since the closure of a convex subset of $\mathbb{R}^{n+1}$ is convex. Pick any $\boldsymbol{\gamma}, \tilde{\boldsymbol{\gamma}} \in \Gamma_{+,\underline{\eta},\bar{\epsilon}}$. Let $\theta, \tilde{\theta} \in \Theta_{+,\underline{\eta},\bar{\epsilon}}$ be such that $(L(\theta), \boldsymbol{m}(\theta)')' = \boldsymbol{\gamma}$ and $(L(\tilde{\theta}), \boldsymbol{m}(\tilde{\theta})')' = \tilde{\boldsymbol{\gamma}}$. Fix $\lambda \in [0, 1]$. By the linearity of $L$ and $\boldsymbol{m}$, $\lambda\boldsymbol{\gamma} + (1-\lambda)\tilde{\boldsymbol{\gamma}} = (L(\lambda\theta + (1-\lambda)\tilde{\theta}), \boldsymbol{m}(\lambda\theta + (1-\lambda)\tilde{\theta})')'$. We have $\lambda\theta + (1-\lambda)\tilde{\theta} \in \Theta$ by the convexity of $\Theta$, $L(\lambda\theta + (1-\lambda)\tilde{\theta}) = \lambda L(\theta) + (1-\lambda)L(\tilde{\theta}) \geq \underline{\eta}$, $\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\lambda\theta+(1-\lambda)\tilde{\theta})}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} = \lambda\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} + (1-\lambda)\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\tilde{\theta})}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} \geq \underline{\eta}$, and $\|\boldsymbol{m}(\lambda\theta + (1-\lambda)\tilde{\theta})\| \leq \|\boldsymbol{m}(\lambda\theta)\| + \|\boldsymbol{m}((1-\lambda)\tilde{\theta})\| = \lambda\|\boldsymbol{m}(\theta)\| + (1-\lambda)\|\boldsymbol{m}(\tilde{\theta})\| \leq \bar{\epsilon}$. Therefore, $\lambda\theta + (1-\lambda)\tilde{\theta} \in \Theta_{+,\underline{\eta},\bar{\epsilon}}$ and $\lambda\boldsymbol{\gamma} + (1-\lambda)\tilde{\boldsymbol{\gamma}} \in \Gamma_{+,\underline{\eta},\bar{\epsilon}}$.

Next, I show that $\Psi(\boldsymbol{\gamma})$ is nonempty and convex for all $\boldsymbol{\gamma} \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$. Fix $\boldsymbol{\gamma} = (\alpha, \boldsymbol{\beta}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$. Let $S_{\boldsymbol{\beta}} = \left\{\left(\bar{\alpha}, \frac{\boldsymbol{\beta}'\bar{\boldsymbol{\beta}}}{\sigma\|\boldsymbol{\beta}\|}\right) \in \mathbb{R}^2 : \bar{\boldsymbol{\gamma}} = (\bar{\alpha}, \bar{\boldsymbol{\beta}}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}\right\}$, which is a subset of $(0, \infty) \times \mathbb{R}$. Using $S_{\boldsymbol{\beta}}$, we can write $\Psi(\boldsymbol{\gamma}) = \left\{(\bar{\alpha}, \bar{\boldsymbol{\beta}}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}} : \left(\bar{\alpha}, \frac{\boldsymbol{\beta}'\bar{\boldsymbol{\beta}}}{\sigma\|\boldsymbol{\beta}\|}\right) \in \arg\max_{(a,b) \in S_{\boldsymbol{\beta}}} a\Phi(-b)\right\}$. Since the mapping $\bar{\boldsymbol{\gamma}} \mapsto \left(\bar{\alpha}, \frac{\boldsymbol{\beta}'\bar{\boldsymbol{\beta}}}{\sigma\|\boldsymbol{\beta}\|}\right)$ is continuous and $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ is compact, $S_{\boldsymbol{\beta}}$ is compact. Furthermore, since $\bar{\boldsymbol{\gamma}} \mapsto \left(\bar{\alpha}, \frac{\boldsymbol{\beta}'\bar{\boldsymbol{\beta}}}{\sigma\|\boldsymbol{\beta}\|}\right)$ is linear and $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ is convex, $S_{\boldsymbol{\beta}}$ is convex. It then follows that $\arg\max_{(a,b) \in S_{\boldsymbol{\beta}}} a\Phi(-b)$ is nonempty and singleton, since $a\Phi(-b)$ is continuous and is strictly quasi-concave on $(0, \infty) \times \mathbb{R}$ by Lemma A.2. Let $(a^*_{\boldsymbol{\beta}}, b^*_{\boldsymbol{\beta}}) \in \arg\max_{(a,b) \in S_{\boldsymbol{\beta}}} a\Phi(-b)$. We can then write $\Psi(\boldsymbol{\gamma}) = \left\{(\bar{\alpha}, \bar{\boldsymbol{\beta}}')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}} : \bar{\alpha} = a^*_{\boldsymbol{\beta}}, \frac{\boldsymbol{\beta}'\bar{\boldsymbol{\beta}}}{\sigma\|\boldsymbol{\beta}\|} = b^*_{\boldsymbol{\beta}}\right\}$, which is nonempty and convex.

Lastly, I show that $\Psi$ has a closed graph. Take any sequence $\{(\boldsymbol{\gamma}_n, \boldsymbol{\gamma}_n^*)\}_{n=1}^{\infty}$ such that $\boldsymbol{\gamma}_n, \boldsymbol{\gamma}_n^* \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ for all $n$, $\lim_{n\to\infty}(\boldsymbol{\gamma}_n, \boldsymbol{\gamma}_n^*) = (\boldsymbol{\gamma}, \boldsymbol{\gamma}^*)$, and $\boldsymbol{\gamma}_n^* \in \Psi(\boldsymbol{\gamma}_n)$ for all $n$. Since $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ is closed, $\boldsymbol{\gamma} \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$. It then suffices to show that $\boldsymbol{\gamma}^* \in \Psi(\boldsymbol{\gamma})$. Suppose $\boldsymbol{\gamma}^* \notin \Psi(\boldsymbol{\gamma})$. Let $f(\bar{\boldsymbol{\gamma}}, \boldsymbol{\gamma}) = \bar{\alpha}\Phi\left(-\frac{\boldsymbol{\beta}'\bar{\boldsymbol{\beta}}}{\sigma\|\boldsymbol{\beta}\|}\right)$. Then there exist $\boldsymbol{\gamma}^{**} \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ and $\epsilon > 0$ such that $f(\boldsymbol{\gamma}^{**}, \boldsymbol{\gamma}) > f(\boldsymbol{\gamma}^*, \boldsymbol{\gamma}) + 3\epsilon$. Also, since $f$ is continuous on $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}} \times \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ and $(\boldsymbol{\gamma}_n, \boldsymbol{\gamma}_n^*) \to (\boldsymbol{\gamma}, \boldsymbol{\gamma}^*)$, we have $f(\boldsymbol{\gamma}^{**}, \boldsymbol{\gamma}_n) > f(\boldsymbol{\gamma}^{**}, \boldsymbol{\gamma}) - \epsilon$ and $f(\boldsymbol{\gamma}^*, \boldsymbol{\gamma}) > f(\boldsymbol{\gamma}_n^*, \boldsymbol{\gamma}_n) - \epsilon$ for any sufficiently large $n$. Combining the preceding inequalities, we obtain for any sufficiently large $n$, $f(\boldsymbol{\gamma}^{**}, \boldsymbol{\gamma}_n) > f(\boldsymbol{\gamma}^*, \boldsymbol{\gamma}) + 2\epsilon > f(\boldsymbol{\gamma}_n^*, \boldsymbol{\gamma}_n) + \epsilon$. This contradicts the assumption that $\boldsymbol{\gamma}_n^* \in \Psi(\boldsymbol{\gamma}_n)$ for all $n$.

Application of Kakutani's fixed point theorem proves the statement. $\square$

Let $\boldsymbol{\gamma}^* = (\alpha^*, (\boldsymbol{\beta}^*)')'$ be a fixed point of $\Psi$. In Steps 2–5, I prove that $L(\theta_{\epsilon^*}) = \alpha^*$ and $\boldsymbol{m}(\theta_{\epsilon^*}) = \boldsymbol{\beta}^*$. $\boldsymbol{\gamma}^*$ may not be an element of $\Gamma_{+,\underline{\eta},\bar{\epsilon}}$, but since it is an element of $\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$, we can take sequences $\{\boldsymbol{\gamma}_n = (\alpha_n, \boldsymbol{\beta}_n')'\}_{n=1}^{\infty}$ and $\{\theta_n\}_{n=1}^{\infty}$ such that $\theta_n \in \Theta_{+,\underline{\eta},\bar{\epsilon}}$ and $\boldsymbol{\gamma}_n = (L(\theta_n), \boldsymbol{m}(\theta_n)')' \in \Gamma_{+,\underline{\eta},\bar{\epsilon}}$ for all $n$ and that $\lim_{n\to\infty} \boldsymbol{\gamma}_n = \boldsymbol{\gamma}^*$. Below, I suppress the argument $\sigma$ of the minimax risk $\mathcal{R}(\sigma;\cdot)$ for notational brevity.

**Step 2.** $\lim_{n\to\infty} \mathcal{R}([-\theta_n, \theta_n] \cap \Theta_{\underline{\eta},\bar{\epsilon}}) = \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta, \theta] \cap \Theta_{\underline{\eta},\bar{\epsilon}})$.

*Proof.* Let $\tilde{\delta}(\boldsymbol{Y}) = \mathbf{1}\{(\boldsymbol{\beta}^*)'\boldsymbol{Y} \geq 0\}$. I first show that $\sup_{\theta\in\Theta_{\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta}, \theta) = \alpha^*\Phi\left(-\frac{\|\boldsymbol{\beta}^*\|}{\sigma}\right)$. Since $\boldsymbol{\gamma}^* \in \arg\max_{\boldsymbol{\gamma}\in\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}} \alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right)$,

$$\begin{aligned}
&\sup_{\boldsymbol{\gamma}\in\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}} \alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right) = \alpha^*\Phi\left(-\frac{(\boldsymbol{\beta}^*)'(\boldsymbol{\beta}^*)}{\sigma\|\boldsymbol{\beta}^*\|}\right) \\
&= \lim_{n\to\infty} \alpha_n\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}_n}{\sigma\|\boldsymbol{\beta}^*\|}\right) = \lim_{n\to\infty} L(\theta_n)\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{m}(\theta_n)}{\sigma\|\boldsymbol{\beta}^*\|}\right) = \lim_{n\to\infty} R(\tilde{\delta}, \theta_n),
\end{aligned}$$

where the second equality follows by the fact that the mapping $\boldsymbol{\gamma} \mapsto \alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right)$ is continuous. On the other hand, by definition,

$$\begin{aligned}
&\sup_{\boldsymbol{\gamma}\in\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}} \alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right) \geq \sup_{\boldsymbol{\gamma}\in\Gamma_{+,\underline{\eta},\bar{\epsilon}}} \alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right) \\
&= \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} L(\theta)\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{\beta}^*\|}\right) = \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta}, \theta) \geq \lim_{n\to\infty} R(\tilde{\delta}, \theta_n).
\end{aligned}$$

Therefore, $\sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta},\theta) = \lim_{n\to\infty} R(\tilde{\delta},\theta_n) = \alpha^*\Phi\left(-\frac{(\boldsymbol{\beta}^*)'(\boldsymbol{\beta}^*)}{\sigma\|\boldsymbol{\beta}^*\|}\right)$. Note that $\sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta},\theta) = \sup_{\theta\in\Theta_{\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta},\theta)$ by the symmetry of the regret function and the centrosymmetry of $\Theta_{\underline{\eta},\bar{\epsilon}}$.

Next, let $\delta_n(\boldsymbol{Y}) = \mathbf{1}\left\{\boldsymbol{m}(\theta_n)'\boldsymbol{Y} \geq 0\right\}$ for all $n$. Observe that, by continuity of the mapping $\boldsymbol{\gamma} \mapsto \alpha\Phi\left(-\frac{(\boldsymbol{\beta})'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}\|}\right)$ at $\boldsymbol{\gamma} = \boldsymbol{\gamma}^*$,

$$\begin{aligned}\lim_{n\to\infty} R(\delta_n,\theta_n) &= \lim_{n\to\infty} L(\theta_n)\Phi\left(-\frac{\boldsymbol{m}(\theta_n)'\boldsymbol{m}(\theta_n)}{\sigma\|\boldsymbol{m}(\theta_n)\|}\right)\\ &= \lim_{n\to\infty} \alpha_n\Phi\left(-\frac{\boldsymbol{\beta}_n'\boldsymbol{\beta}_n}{\sigma\|\boldsymbol{\beta}_n\|}\right) = \alpha^*\Phi\left(-\frac{(\boldsymbol{\beta}^*)'(\boldsymbol{\beta}^*)}{\sigma\|\boldsymbol{\beta}^*\|}\right) = \sup_{\theta\in\Theta_{\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta},\theta).\end{aligned}$$

Now, I show that $\lim_{n\to\infty} \mathcal{R}([-\theta_n,\theta_n]\cap\Theta_{\underline{\eta},\bar{\epsilon}}) = \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta,\theta]\cap\Theta_{\underline{\eta},\bar{\epsilon}})$. We have that for all $n$,

$$R(\delta_n,\theta_n) \leq \sup_{\theta\in[-\theta_n,\theta_n]\cap\Theta_{\underline{\eta},\bar{\epsilon}}} R(\delta_n,\theta) = \mathcal{R}([-\theta_n,\theta_n]\cap\Theta_{\underline{\eta},\bar{\epsilon}}) \leq \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta,\theta]\cap\Theta_{\underline{\eta},\bar{\epsilon}}), \tag{A.2}$$

where the equality holds since $[-\theta_n,\theta_n]\cap\Theta_{\underline{\eta},\bar{\epsilon}} = [-\theta_n,-t_n\theta_n]\cup[t_n\theta_n,\theta_n]$ with $t_n = \max\{\frac{\eta}{L(\theta_n)}, \underline{\eta}\frac{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta_n)}\}$ and $\delta_n$ is minimax regret for $[-\theta_n,-t_n\theta_n]\cup[t_n\theta_n,\theta_n]$ by Lemma A.8. On the other hand,

$$\lim_{n\to\infty} R(\delta_n,\theta_n) = \sup_{\theta\in\Theta_{\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta},\theta) \geq \mathcal{R}(\Theta_{\underline{\eta},\bar{\epsilon}}) \geq \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta,\theta]\cap\Theta_{\underline{\eta},\bar{\epsilon}}), \tag{A.3}$$

where the inequalities hold by the definition of the minimax risk. (A.2) and (A.3) imply that $\lim_{n\to\infty} R(\delta_n,\theta_n) = \sup_{\theta\in\Theta_{\underline{\eta},\bar{\epsilon}}} R(\tilde{\delta},\theta) = \lim_{n\to\infty} \mathcal{R}([-\theta_n,\theta_n]\cap\Theta_{\underline{\eta},\bar{\epsilon}}) = \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta,\theta]\cap\Theta_{\underline{\eta},\bar{\epsilon}})$. □

**Step 3.** $\sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta,\theta]\cap\Theta_{\underline{\eta},\bar{\epsilon}}) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$.

*Proof.* By Lemma A.8,

$$\mathcal{R}([-\theta_{\epsilon^*},\theta_{\epsilon^*}]) = \frac{L(\theta_{\epsilon^*})}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\mathcal{R}_{\text{uni}}\left(\sigma;[-\|\boldsymbol{m}(\theta_{\epsilon^*})\|,\|\boldsymbol{m}(\theta_{\epsilon^*})\|]\right) = \frac{\omega(\epsilon^*)}{\epsilon^*}\mathcal{R}_{\text{uni}}\left(\sigma;[-\epsilon^*,\epsilon^*]\right),$$

and with $t = \max\{\frac{\underline{\eta}}{L(\theta_{\epsilon^*})}, \frac{\underline{\eta}}{\epsilon^*}\}$,

$$\begin{aligned}
&\mathcal{R}([-\theta_{\epsilon^*}, \theta_{\epsilon^*}] \cap \Theta_{\underline{\eta},\bar{\epsilon}}) = \mathcal{R}([-\theta_{\epsilon^*}, -t\theta_{\epsilon^*}] \cup [t\theta_{\epsilon^*}, \theta_{\epsilon^*}]) \\
&= \frac{L(\theta_{\epsilon^*})}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{m}(\theta_{\epsilon^*})\|, -t\|\boldsymbol{m}(\theta_{\epsilon^*})\|] \cup [t\|\boldsymbol{m}(\theta_{\epsilon^*})\|, \|\boldsymbol{m}(\theta_{\epsilon^*})\|]\right) \\
&= \frac{\omega(\epsilon^*)}{\epsilon^*}\mathcal{R}_{\text{uni}}\left(\sigma; [-\epsilon^*, -t\epsilon^*] \cup [t\epsilon^*, \epsilon^*]\right).
\end{aligned}$$

Since $\epsilon^* \leq a^*\sigma$, $\mathcal{R}_{\text{uni}}\left(\sigma; [-\epsilon^*, \epsilon^*]\right) = \mathcal{R}_{\text{uni}}\left(\sigma; [-\epsilon^*, -t\epsilon^*] \cup [t\epsilon^*, \epsilon^*]\right)$ by Lemma A.7. Therefore, $\mathcal{R}([-\theta_{\epsilon^*}, \theta_{\epsilon^*}]) = \mathcal{R}([-\theta_{\epsilon^*}, \theta_{\epsilon^*}] \cap \Theta_{\underline{\eta},\bar{\epsilon}})$. Note that this equals $\sup_{\theta\in\Theta: L(\theta)>0, \boldsymbol{m}(\theta)\neq\boldsymbol{0}} \mathcal{R}([-\theta, \theta]) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$ by Lemma A.11. On the other hand, we have

$$\sup_{\theta\in\Theta: L(\theta)>0, \boldsymbol{m}(\theta)\neq\boldsymbol{0}} \mathcal{R}([-\theta, \theta]) \geq \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta, \theta] \cap \Theta_{\underline{\eta},\bar{\epsilon}}) \geq \mathcal{R}([-\theta_{\epsilon^*}, \theta_{\epsilon^*}] \cap \Theta_{\underline{\eta},\bar{\epsilon}}),$$

where the first inequality holds by construction and the second holds since $\theta_{\epsilon^*} \in \Theta_{+,\underline{\eta},\bar{\epsilon}}$. It follows that $\sup_{\theta\in\Theta: L(\theta)>0, \boldsymbol{m}(\theta)\neq\boldsymbol{0}} \mathcal{R}([-\theta, \theta]) = \sup_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} \mathcal{R}([-\theta, \theta] \cap \Theta_{\underline{\eta},\bar{\epsilon}}) = \mathcal{R}([-\theta_{\epsilon^*}, \theta_{\epsilon^*}] \cap \Theta_{\underline{\eta},\bar{\epsilon}}) = \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$. □

**Step 4.** *$\alpha^* = \omega(\epsilon^*)$ and $\|\boldsymbol{\beta}^*\| = \epsilon^*$.*

*Proof.* Note first that $\alpha_n \leq \omega(\|\boldsymbol{\beta}_n\|)$ for all $n$ by the definition of $\omega(\cdot)$. Since $\omega(\cdot)$ is continuous on $(0, \infty]$ by the concavity, taking the limit of both sides yields $\alpha^* \leq \omega(\|\boldsymbol{\beta}^*\|)$.

With $t_n = \max\{\frac{\underline{\eta}}{L(\theta_n)}, \underline{\eta}\frac{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta_n)}\}$ and $t^* = \lim_{n\to\infty} t_n = \max\{\frac{\underline{\eta}}{\alpha^*}, \underline{\eta}\frac{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}\}$,

$$\begin{aligned}
&\lim_{n\to\infty} \mathcal{R}([-\theta_n, \theta_n] \cap \Theta_{\underline{\eta},\bar{\epsilon}}) = \lim_{n\to\infty} \mathcal{R}([-\theta_n, -t_n\theta_n] \cup [t_n\theta_n, \theta_n]) \\
&= \lim_{n\to\infty} \frac{L(\theta_n)}{\|\boldsymbol{m}(\theta_n)\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{m}(\theta_n)\|, -t_n\|\boldsymbol{m}(\theta_n)\|] \cup [t_n\|\boldsymbol{m}(\theta_n)\|, \|\boldsymbol{m}(\theta_n)\|]\right) \\
&= \frac{\alpha^*}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, -t^*\|\boldsymbol{\beta}^*\|] \cup [t^*\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right),
\end{aligned}$$

where the last equality follows from the fact that the mapping $(\tau, t) \mapsto \mathcal{R}_{\text{uni}}(\sigma; [-\tau, -t\tau] \cup [t\tau, \tau])$ is continuous. By Steps 2–3, $\lim_{n\to\infty} \mathcal{R}([-\theta_n, \theta_n] \cap \Theta_{\underline{\eta},\bar{\epsilon}}) =$

$\omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$. Therefore,

$$\begin{aligned}\omega(\epsilon^*)\Phi(-\epsilon^*/\sigma) &= \frac{\alpha^*}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, -t^*\|\boldsymbol{\beta}^*\|] \cup [t^*\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right) \\ &\leq \frac{\alpha^*}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right) \leq \frac{\omega(\|\boldsymbol{\beta}^*\|)}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right),\end{aligned}$$

where the last inequality holds since $\alpha^* \leq \omega(\|\boldsymbol{\beta}^*\|)$. On the other hand, by Lemma A.11,

$$\omega(\epsilon^*)\Phi(-\epsilon^*/\sigma) = \sup_{\epsilon>0}\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}\left(\sigma; [-\epsilon, \epsilon]\right) \geq \frac{\omega(\|\boldsymbol{\beta}^*\|)}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right).$$

It follows that

$$\begin{aligned}&\sup_{\epsilon>0}\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}\left(\sigma; [-\epsilon, \epsilon]\right) = \frac{\alpha^*}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, -t^*\|\boldsymbol{\beta}^*\|] \cup [t^*\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right) \\ &= \frac{\alpha^*}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right) = \frac{\omega(\|\boldsymbol{\beta}^*\|)}{\|\boldsymbol{\beta}^*\|}\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right).\end{aligned}$$

Therefore, $\alpha^* = \omega(\|\boldsymbol{\beta}^*\|)$ and $\|\boldsymbol{\beta}^*\| \in \arg\max_{\epsilon>0}\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}\left(\sigma; [-\epsilon, \epsilon]\right)$. Furthermore,

$$\mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, -t^*\|\boldsymbol{\beta}^*\|] \cup [t^*\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right) = \mathcal{R}_{\text{uni}}\left(\sigma; [-\|\boldsymbol{\beta}^*\|, \|\boldsymbol{\beta}^*\|]\right). \tag{A.4}$$

If it is shown that $\|\boldsymbol{\beta}^*\| \leq a^*\sigma$, then $\|\boldsymbol{\beta}^*\| \in \arg\max_{0<\epsilon\leq a^*\sigma}\frac{\omega(\epsilon)}{\epsilon}\mathcal{R}_{\text{uni}}\left(\sigma; [-\epsilon, \epsilon]\right) = \arg\max_{0<\epsilon\leq a^*\sigma}\omega(\epsilon)\Phi(-\epsilon/\sigma) = \{\epsilon^*\}$ by Assumption A.1. Suppose to the contrary that $\|\boldsymbol{\beta}^*\| > a^*\sigma$. By the fact that $g(a) = a\Phi(-a)$ is strictly increasing on $[0, a^*)$ and strictly decreasing on $(a^*, \infty)$ as shown in the proof of Lemma A.7, the form of $\mathcal{R}_{\text{uni}}$ given in Lemma A.7 implies that it is necessary that $t^*\|\boldsymbol{\beta}^*\| \leq a^*\sigma$ for Eq. (A.4) to hold. Therefore, $t^* \leq \frac{a^*\sigma}{\|\boldsymbol{\beta}^*\|} < 1$ since $\|\boldsymbol{\beta}^*\| > a^*\sigma$. It follows that $t^* = \max\{\frac{\underline{\eta}}{\alpha^*}, \underline{\eta}\frac{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}\} < 1$, so $\bar{\eta} < \alpha^*$ and $\underline{\eta} < \frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}$. We can pick $\underline{t} \in (0, 1)$ sufficiently close to 1 so that for all $t \in [\underline{t}, 1]$, $\underline{\eta} < t\alpha^* = \lim_{n\to\infty} L(t\theta_n)$ and $\underline{\eta} < \frac{\boldsymbol{m}(\theta_{\epsilon^*})'t\boldsymbol{\beta}^*}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} = \lim_{n\to\infty}\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(t\theta_n)}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|}$. It follows that $t\boldsymbol{\gamma}^* = (t\alpha^*, t(\boldsymbol{\beta}^*)')' = \lim_{n\to\infty}(L(t\theta_n), \boldsymbol{m}(t\theta_n)')' \in \bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}$ for all $t \in [\underline{t}, 1]$. Since $\boldsymbol{\gamma}^* \in \arg\max_{\boldsymbol{\gamma}\in\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}}\alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right)$, we have $\boldsymbol{\gamma}^* \in \arg\max_{\boldsymbol{\gamma}\in[\underline{t}\boldsymbol{\gamma}^*, \boldsymbol{\gamma}^*]}\alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right)$, which implies that $1 \in \arg\max_{t\in[\underline{t},1]} t\alpha^*\Phi\left(-\frac{(\boldsymbol{\beta}^*)'t\boldsymbol{\beta}^*}{\sigma\|\boldsymbol{\beta}^*\|}\right) = \arg\max_{t\in[\underline{t},1]} t\Phi\left(-\frac{t\|\boldsymbol{\beta}^*\|}{\sigma}\right)$. By Lemma

A.1, $t \mapsto t\Phi\left(-\frac{t\|\boldsymbol{\beta}^*\|}{\sigma}\right)$ is strictly increasing on $[0, \frac{a^*\sigma}{\|\boldsymbol{\beta}^*\|})$ and strictly decreasing on $(\frac{a^*\sigma}{\|\boldsymbol{\beta}^*\|}, 1]$. Hence, $1 \notin \arg\max_{t\in[\underline{t},1]} t\Phi\left(-\frac{t\|\boldsymbol{\beta}^*\|}{\sigma}\right)$, which is a contradiction. $\square$

**Step 5.** $\boldsymbol{m}(\theta_{\epsilon^*}) = \boldsymbol{\beta}^*$.

*Proof.* Suppose that $\boldsymbol{m}(\theta_{\epsilon^*}) \neq \boldsymbol{\beta}^*$. Pick any $\lambda \in (0,1)$, and consider the sequence $\{\theta_{\lambda,n}\}_{n=1}^{\infty}$, where $\theta_{\lambda,n} = \lambda\theta_{\epsilon^*} + (1-\lambda)\theta_n$. Since $\Theta_{+,\underline{\eta},\bar{\epsilon}}$ is convex, $\theta_{\lambda,n} \in \Theta_{+,\underline{\eta},\bar{\epsilon}}$ for all $n$. We have $\lim_{n\to\infty} L(\theta_{\lambda,n}) = \lambda L(\theta_{\epsilon^*}) + (1-\lambda)\lim_{n\to\infty} L(\theta_n) = \lambda\omega(\epsilon^*) + (1-\lambda)\alpha^* = \omega(\epsilon^*)$, and $\lim_{n\to\infty} \|\boldsymbol{m}(\theta_{\lambda,n})\| = \|\lambda\boldsymbol{m}(\theta_{\epsilon^*}) + (1-\lambda)\lim_{n\to\infty}\boldsymbol{m}(\theta_n)\| = \|\lambda\boldsymbol{m}(\theta_{\epsilon^*}) + (1-\lambda)\boldsymbol{\beta}^*\| < \epsilon^*$, where the last inequality holds since $\{\boldsymbol{\beta} \in \mathbb{R}^n : \|\boldsymbol{\beta}\| \le \epsilon^*\}$ is strictly convex. This implies that $\omega(\tilde{\epsilon}) = \sup\{L(\theta) : \theta \in \Theta, \|\boldsymbol{m}(\theta)\| \le \tilde{\epsilon}\} = \omega(\epsilon^*)$ for any $\tilde{\epsilon} \in (\lim_{n\to\infty}\|\boldsymbol{m}(\theta_{\lambda,n})\|, \epsilon^*)$. It follows that $\omega(\tilde{\epsilon})\Phi(-\tilde{\epsilon}/\sigma) = \omega(\epsilon^*)\Phi(-\tilde{\epsilon}/\sigma) > \omega(\epsilon^*)\Phi(-\epsilon^*/\sigma)$, which contradicts the fact that $\epsilon^*$ maximizes $\omega(\epsilon)\Phi(-\epsilon/\sigma)$ over $[0, a^*\sigma]$. $\square$

**Step 6.** $\theta_{\epsilon^*} \in \arg\max_{\theta\in\Theta:L(\theta)>0} L(\theta)\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$.

*Proof.* Note first that, since $\boldsymbol{\gamma}^* = (L(\theta_{\epsilon^*}), \boldsymbol{m}(\theta_{\epsilon^*})')' \in \arg\max_{\boldsymbol{\gamma}\in\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}} \alpha\Phi\left(-\frac{(\boldsymbol{\beta}^*)'\boldsymbol{\beta}}{\sigma\|\boldsymbol{\beta}^*\|}\right)$, we have $\boldsymbol{\gamma}^* \in \arg\max_{\boldsymbol{\gamma}\in\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}} \alpha\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$ and $\theta_{\epsilon^*} \in \arg\max_{\theta\in\Theta_{+,\underline{\eta},\bar{\epsilon}}} L(\theta)\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$.

Pick any $\theta \in \Theta$ such that $L(\theta) > 0$ and $\theta \notin \Theta_{+,\underline{\eta},\bar{\epsilon}}$. Let $\boldsymbol{\gamma} = (\alpha, \boldsymbol{\beta}')' = (L(\theta), \boldsymbol{m}(\theta)')'$. Since $\underline{\eta} < \omega(\epsilon^*) = \alpha^*$ and $\underline{\eta} < \epsilon^* < \bar{\epsilon}$ by the choice of $\underline{\eta}$ and $\bar{\epsilon}$, we can pick $t \in (0,1)$ sufficiently close to 1 so that $L((1-t)\theta + t\theta_{\epsilon^*}) = (1-t)\alpha + t\alpha^* > \underline{\eta}$, $\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}((1-t)\theta+t\theta_{\epsilon^*})}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} = \frac{\boldsymbol{m}(\theta_{\epsilon^*})'((1-t)\boldsymbol{\beta}+t\boldsymbol{\beta}^*)}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} = (1-t)\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}}{\|\boldsymbol{m}(\theta_{\epsilon^*})\|} + t\epsilon^* > \underline{\eta}$, and $\|\boldsymbol{m}((1-t)\theta + t\theta_{\epsilon^*})\| = \|(1-t)\boldsymbol{\beta} + t\boldsymbol{\beta}^*\| \le (1-t)\|\boldsymbol{\beta}\| + t\epsilon^* < \bar{\epsilon}$. It follows that $(1-t)\theta + t\theta_{\epsilon^*} \in \Theta_{+,\underline{\eta},\bar{\epsilon}}$ and $(1-t)\boldsymbol{\gamma} + t\boldsymbol{\gamma}^* \in \Gamma_{+,\underline{\eta},\bar{\epsilon}}$. Since $\boldsymbol{\gamma}^* \in \arg\max_{\boldsymbol{\gamma}\in\bar{\Gamma}_{+,\underline{\eta},\bar{\epsilon}}} \alpha\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$, this implies that $\alpha^*\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right) \ge [(1-t)\alpha + t\alpha^*]\Phi\left(-(1-t)\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|} - t\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$. Since the function $(a,b) \mapsto a\Phi(-b)$ is strictly quasi-concave on $(0,\infty)\times\mathbb{R}$ by Lemma A.2, $[(1-t)\alpha + t\alpha^*]\Phi\left(-(1-t)\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|} - t\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right) > \min\left\{\alpha\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right), \alpha^*\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)\right\}$. Therefore, $\alpha^*\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}^*}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right) > \alpha\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{\beta}}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$, which implies that $L(\theta_{\epsilon^*})\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta_{\epsilon^*})}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right) > L(\theta)\Phi\left(-\frac{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{m}(\theta)}{\sigma\|\boldsymbol{m}(\theta_{\epsilon^*})\|}\right)$. The conclusion then follows. $\square$

### A.4.5 Proof of Lemma A.13

First, $\|\boldsymbol{m}(\theta_\epsilon)\| = \epsilon$ for all $\epsilon \in [0, \bar{\epsilon}]$; if $\|\boldsymbol{m}(\theta_\epsilon)\| < \epsilon$, then there exists $c$ in a neighborhood of zero such that $\theta_\epsilon + c\iota \in \Theta$, $L(\theta_\epsilon + c\iota) > L(\theta_\epsilon)$, and $\|\boldsymbol{m}(\theta_\epsilon + c\iota)\| < \epsilon$ by Assumption 1(i)(c), which contradicts Assumption 1(i)(a).

Without loss of generality, I normalize $\iota$ so that $L(\iota) = 1$. By Lemma A.3, for any $\epsilon \in (0, \bar{\epsilon}]$, $\omega'(\epsilon) = \frac{\epsilon}{\boldsymbol{m}(\iota)'\boldsymbol{m}(\theta_\epsilon)}$. Since $\omega(\cdot)$ is differentiable on $[0, a^*\sigma]$ and concave, it is continuously differentiable at 0. Therefore, $\omega'(0) = \lim_{\epsilon\to 0}\omega'(\epsilon) = \lim_{\epsilon\to 0}\frac{1}{\boldsymbol{m}(\iota)'\boldsymbol{m}(\theta_\epsilon)/\epsilon} = \lim_{\epsilon\to 0}\frac{1}{\boldsymbol{m}(\iota)'\boldsymbol{m}(\theta_\epsilon)/\|\boldsymbol{m}(\theta_\epsilon)\|} = \frac{1}{\boldsymbol{m}(\iota)'\boldsymbol{w}^*}$ by Assumption 1(i)(b).

Next, I use Lemma A.5 to derive $\rho'(0)$. Since Assumption A.2(i)(a)–(b) and (ii) hold, Lemma A.6 implies that $\omega(0) = \rho(0)$. Therefore, $L(\theta_0) = \omega(0) = \rho(0)$. Additionally, $(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_0) = 0$ by construction. Also, $0 \in \mathrm{int}(\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\})$; otherwise, it is necessary that $\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\} = \{0\}$ by the convexity and centrosymmetry of $\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\}$, but then $\rho(\epsilon) = -\infty$ for any $\epsilon \neq 0$, which contradicts the assumption that $\rho(\cdot)$ is differentiable at any $\epsilon \in \{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\}$. Applying Lemma A.5, we have that $\rho'(0) = \frac{1}{(\boldsymbol{w}^*)'\boldsymbol{m}(\iota)}$.

Since $\omega'(0)$ is nonnegative by the fact that the modulus is nondecreasing and $\boldsymbol{m}(\iota)'\boldsymbol{w}^*$ is finite, $\omega'(0) = \rho'(0) > 0$. Moreover, by the assumption that $2\phi(0)\frac{\omega(0)}{\omega'(0)} > 0$, $\omega(0) = \rho(0) > 0$.

Now, let $g(\epsilon) = \rho(\epsilon)\Phi\left(-\frac{\epsilon}{\sigma^*}\right)$, where $\sigma^* = 2\phi(0)\frac{\omega(0)}{\omega'(0)} = 2\phi(0)\frac{\rho(0)}{\rho'(0)}$. I show that $0 \in \arg\max_{\epsilon\in\mathbb{R}} g(\epsilon)$. Here I only consider the case where $\{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta) : \theta \in \Theta\} = \mathbb{R}$. The proof for other cases is analogous. Let $\underline{\epsilon} = \inf\{\epsilon : \rho(\epsilon) > 0\}$, $\tilde{\epsilon} = \sup\{\epsilon : \rho'(\epsilon) > 0\}$, and $\bar{\epsilon} = \sup\{\epsilon : \rho(\epsilon) > 0\}$. Since $\rho(0) > 0$ and $\rho(\cdot)$ is continuous, $\underline{\epsilon} < 0$ and $\bar{\epsilon} > 0$. Moreover, since $\rho'(0) > 0$ and $\rho'(\cdot)$ is continuous by the concavity (Lemma A.4) and differentiability of $\rho(\cdot)$, we have $\tilde{\epsilon} > 0$. In addition, since $\rho(0) > 0$, $\rho'(0) > 0$, and $\rho'(\cdot)$ is nonincreasing by the concavity of $\rho(\cdot)$, we have that $\rho'(\cdot)$ is positive on $(\underline{\epsilon}, \tilde{\epsilon})$, $\rho(\cdot)$ is nondecreasing and positive on $(\underline{\epsilon}, \tilde{\epsilon})$, and $\frac{\rho(\epsilon)}{\rho'(\epsilon)}$ is nondecreasing on $(\underline{\epsilon}, \tilde{\epsilon})$. By differentiating $g(\cdot)$, we have for $\epsilon \in (\underline{\epsilon}, \tilde{\epsilon})$,

$$g'(\epsilon) = \rho'(\epsilon)\Phi\left(-\frac{\epsilon}{\sigma^*}\right) - \rho(\epsilon)\phi\left(-\frac{\epsilon}{\sigma^*}\right)/\sigma^* = \left[\sigma^*\frac{1-\Phi\left(\frac{\epsilon}{\sigma^*}\right)}{\phi\left(\frac{\epsilon}{\sigma^*}\right)} - \frac{\rho(\epsilon)}{\rho'(\epsilon)}\right]\frac{\rho'(\epsilon)}{\sigma^*}\phi\left(\frac{\epsilon}{\sigma^*}\right).$$

$g'(0) = 0$ by the choice of $\sigma^*$. By the fact that the Mills ratio $\frac{1-\Phi(x)}{\phi(x)}$ of a standard normal random variable is strictly decreasing and continuous, $g'(\epsilon) > 0$ for $\epsilon \in (\underline{\epsilon}, 0)$ and $g'(\epsilon) < 0$ for $\epsilon \in (0, \tilde{\epsilon})$. Therefore, $g(\cdot)$ is maximized at 0 over $\epsilon \in (\underline{\epsilon}, \tilde{\epsilon})$. If $\tilde{\epsilon} < \bar{\epsilon}$, then $\rho'(\epsilon) \leq 0$ and hence $g'(\epsilon) \leq 0$ for $\epsilon \in [\tilde{\epsilon}, \bar{\epsilon})$, so that $g(0) > g(\epsilon)$ for $\epsilon \in [\tilde{\epsilon}, \bar{\epsilon})$. Since $\rho(\cdot)$ is continuous, if $-\infty < \epsilon \leq \underline{\epsilon}$ or $\bar{\epsilon} \leq \epsilon < \infty$, then $\rho(\epsilon) \leq 0$ by the construction of $\underline{\epsilon}$ and $\bar{\epsilon}$, and hence $g(\epsilon) \leq 0 < g(0)$. Thus, $g(\cdot)$ is maximized at 0 globally.

### A.4.6 Proof of Lemma A.14

By Lemma A.6, $\rho(0) = \omega(0)$. We have

$$
\begin{aligned}
&\sup_{\theta \in \Theta} L(\theta) \Phi\left(-\frac{(\boldsymbol{w}^*)' \boldsymbol{m}(\theta)}{\sigma^*}\right) = \sup_{\epsilon \in \mathbb{R}} \sup_{\theta \in \Theta : (\boldsymbol{w}^*)' \boldsymbol{m}(\theta) = \epsilon} L(\theta) \Phi\left(-\frac{\epsilon}{\sigma^*}\right) \\
&= \sup_{\epsilon \in \mathbb{R}} \rho(\epsilon) \Phi\left(-\frac{\epsilon}{\sigma^*}\right) = \frac{1}{2}\rho(0) = \frac{1}{2}\omega(0),
\end{aligned}
$$

where the second last equality holds by Assumption A.2(i)(c). Since $L(\theta_0)\Phi\left(-\frac{(\boldsymbol{w}^*)'\boldsymbol{m}(\theta_0)}{\sigma^*}\right) = \frac{1}{2}\omega(0)$, the conclusion follows.

## A.5 Proof of Proposition 1

The statement immediately holds when $\epsilon^* = 0$, so let $\epsilon^* > 0$. I first show that $\sigma \frac{1-\Phi(\epsilon/\sigma)}{\phi(\epsilon/\sigma)} \geq \frac{\omega(\epsilon)}{\omega'(\epsilon)}$ for all $\epsilon \in [0, \epsilon^*]$. Since $\omega'(0) > 0$, $\omega(\cdot)$ is nonconstant on $[0, \epsilon^*]$. Application of Lemma A.1 to $g(\epsilon) = \omega(\epsilon)\Phi(-\epsilon/\sigma)$ on $[0, \epsilon^*]$ implies that $\sigma \frac{1-\Phi(\epsilon^*/\sigma)}{\phi(\epsilon^*/\sigma)} \geq \frac{\omega(\epsilon^*)}{\omega'(\epsilon^*)}$; otherwise, $\epsilon^*$ does not solve $\max_{\epsilon \in [0,\epsilon^*]} g(\epsilon)$. By the fact that the Mills ratio $\frac{1-\Phi(x)}{\phi(x)}$ of a standard normal random variable is strictly decreasing, $\frac{1-\Phi(\epsilon/\sigma)}{\phi(\epsilon/\sigma)}$ is strictly decreasing in $\epsilon$. Moreover, since $\omega(\cdot)$ is nonnegative, nondecreasing, and concave, $\frac{\omega(\epsilon)}{\omega'(\epsilon)}$ is nondecreasing. Hence, $\sigma \frac{1-\Phi(\epsilon/\sigma)}{\phi(\epsilon/\sigma)} \geq \frac{\omega(\epsilon)}{\omega'(\epsilon)}$ for all $\epsilon \in [0, \epsilon^*]$.

Now, note that $\omega(\epsilon_{\text{MSE}}) > 0$ since $\omega'(0) > 0$ and $\omega(\cdot)$ is nonnegative and nondecreasing. Since $\epsilon_{\text{MSE}} > 0$ solve $\frac{\epsilon^2}{\epsilon^2+\sigma^2} = \frac{\omega'(\epsilon)\epsilon}{\omega(\epsilon)}$, $\omega'(\epsilon_{\text{MSE}}) > 0$ and $\frac{\epsilon_{\text{MSE}}^2+\sigma^2}{\epsilon_{\text{MSE}}} = \frac{\omega(\epsilon_{\text{MSE}})}{\omega'(\epsilon_{\text{MSE}})}$. If $\frac{\epsilon^2+\sigma^2}{\epsilon} > \sigma \frac{1-\Phi(\epsilon/\sigma)}{\phi(\epsilon/\sigma)}$ for all $\epsilon > 0$, then $\frac{\epsilon^2+\sigma^2}{\epsilon} > \frac{\omega(\epsilon)}{\omega'(\epsilon)}$ for all $\epsilon \leq \epsilon^*$, which implies that $\epsilon^* < \epsilon_{\text{MSE}}$. Below, I show that $\frac{\epsilon^2+\sigma^2}{\epsilon} > \sigma \frac{1-\Phi(\epsilon/\sigma)}{\phi(\epsilon/\sigma)}$ for all $\epsilon > 0$. Let $g(\epsilon) = \frac{\epsilon^2+\sigma^2}{\epsilon}$. We have $g'(\epsilon) = 1 - \frac{\sigma^2}{\epsilon^2}$, which is strictly increasing in $\epsilon$. Therefore, $g(\epsilon)$ is minimized at $\epsilon = \sigma$, at which $g'(\epsilon) = 0$. For any $\epsilon > 0$, $g(\epsilon) \geq g(\sigma) = 2\sigma > \sigma \frac{1-\Phi(0)}{\phi(0)} > \sigma \frac{1-\Phi(\epsilon/\sigma)}{\phi(\epsilon/\sigma)}$, where the second last inequality holds since $\frac{1-\Phi(0)}{\phi(0)} \approx 1.253$, and the last holds since $\frac{1-\Phi(\epsilon/\sigma)}{\phi(\epsilon/\sigma)}$ is strictly decreasing.

### A.6 Proof of Proposition 2

Let $\bar{\epsilon} = (C/\max_i \sigma(x_i, d_i)) \min_{i,j:i\neq j} |x_i - x_j| > 0$. Fix $\epsilon \in [0, \bar{\epsilon}]$ and consider the following problem:

$$\max_{(f(x_i,0),f(x_i,1))_{i=1,\ldots,n}\in\mathbb{R}^{2n}} \frac{1}{n} \sum_{i:c_1\leq x_i<c_0} [f(x_i,1) - f(x_i,0)] \tag{A.5}$$
$$s.t. \quad \sum_{i:c_1\leq x_i<c_0} \frac{f(x_i,0)^2}{\sigma^2(x_i,0)} + \frac{f(x_{+,\min},1)^2}{\sigma^2_{+,\min}} \leq \epsilon^2,$$
$$f(x_i,1) - f(x_{+,\min},1) \leq C|x_i - x_{+,\min}|, \quad i \in \{j : c_1 \leq x_j < c_0\}.$$

I first provide a solution to (A.5) and then show that the solution also solves the problem (4). Observe that, given a value of $f(x_{+,\min},1)$, the objective is maximized only when $f(x_i,1) = C(x_{+,\min} - x_i) + f(x_{+,\min},1)$ for any $i$ with $c_1 \leq x_i < c_0$ under the constraints of (A.5). Additionally, none of the objective and constraints depends on $f(x_i,0)$ with $x_i < c_1$ or $x_i \geq c_0$ or on $f(x_i,1)$ with $x_i < c_1$ or $x_i > x_{+,\min}$. Thus, a solution to (A.5) is given by $(f_\epsilon(x_i,0), f_\epsilon(x_i,1))_{i=1,\ldots,n}$, where $f_\epsilon(x_i,0) = 0$ for any $i$ with $x_i < c_1$ or $x_i \geq c_0$, $f_\epsilon(x_i,1) = C(x_{+,\min} - x_i) + f(x_{+,\min},1)$ for any $i$ with $x_i < c_0$, $f_\epsilon(x_i,1) = 0$ for any $i$ with $x_i > x_{+,\min}$, and $((f_\epsilon(x_i,0))_{i:c_1\leq x_i<c_0}, f_\epsilon(x_{+,\min},1))$ solves

$$\max_{\boldsymbol{f}\in\mathbb{R}^{\tilde{n}+1}} \frac{C}{n} \sum_{i:c_1\leq x_i<c_0} (x_{+,\min} - x_i) + \frac{\tilde{n}}{n} f(x_{+,\min},1) - \frac{1}{n} \sum_{i:c_1\leq x_i<c_0} f(x_i,0)$$
$$s.t. \quad \sum_{i:c_1\leq x_i<c_0} \frac{f(x_i,0)^2}{\sigma^2(x_i,0)} + \frac{f(x_{+,\min},1)^2}{\sigma^2_{+,\min}} \leq \epsilon^2,$$

where $\boldsymbol{f}$ denotes $((f(x_i,0))_{i:c_1\leq x_i<c_0}, f(x_{+,\min},1))$. This is a convex optimization problem that maximizes a weighted sum of $\tilde{n}+1$ unknowns under the constraint on the upper bound on a weighted Euclidean norm of the unknowns. Simple calculations show that the solution is given by $f_\epsilon(x_i,0) = -\frac{\sigma^2(x_i,0)\epsilon}{(\tilde{n}^2\sigma^2_{+,\min} + \sum_{i:c_1\leq x_i<c_0}\sigma^2(x_i,0))^{1/2}}$ for any $i$ with $c_1 \leq x_i < c_0$ and $f_\epsilon(x_{+,\min},1) = \frac{\tilde{n}\sigma^2_{+,\min}\epsilon}{(\tilde{n}^2\sigma^2_{+,\min} + \sum_{i:c_1\leq x_i<c_0}\sigma^2(x_i,0))^{1/2}}$. Note that $(f_\epsilon(x_i,0), f_\epsilon(x_i,1))_{i=1,\ldots,n}$ equals the expression in the statement of Proposition 2.

The constraints of (A.5) are weaker than those of (4). To show that $(f_\epsilon(x_i,0), f_\epsilon(x_i,1))_{i=1,\ldots,n}$ also solves (4), it suffices to show that

$(f_\epsilon(x_i,0), f_\epsilon(x_i,1))_{i=1,\dots,n}$ satisfies the constraints of (4). Clearly, the norm constraint is satisfied. To see that the Lipschitz constraint is satisfied for $d=0$, observe that for any $i,j$ with $i \neq j$,

$$
\begin{aligned}
&|f_\epsilon(x_i,0) - f_\epsilon(x_j,0)|^2 = f_\epsilon(x_i,0)^2 + f_\epsilon(x_j,0)^2 - 2f_\epsilon(x_i,0)f_\epsilon(x_j,0) \\
&\leq f_\epsilon(x_i,0)^2 + f_\epsilon(x_j,0)^2 \leq (\max_k \sigma(x_k,d_k))^2\bar{\epsilon}^2 = C^2 \min_{k,l:k\neq l} |x_k - x_l|^2 \leq C^2|x_i - x_j|^2,
\end{aligned}
$$

where the first inequality holds since $f_\epsilon(x_i,0) \leq 0$ for all $i$, the second inequality follows from the norm constraint of (A.5) and the fact that $f_\epsilon(x_i,0) = 0$ for all $i$ with $x_i < c_1$ or $x_i \geq c_0$, and the equality in the second line from the definition of $\bar{\epsilon}$. For $d=1$, by construction, $|f_\epsilon(x_i,1) - f_\epsilon(x_{+,\min},1)| \leq C|x_i - x_{+,\min}|$ for any $i$ with $x_i < x_{+,\min}$. For any $i$ with $x_i > x_{+,\min}$, $|f_\epsilon(x_i,1) - f_\epsilon(x_{+,\min},1)|^2 = f_\epsilon(x_{+,\min},1)^2 \leq (\max_k \sigma(x_k,d_k))^2\bar{\epsilon}^2 = C^2 \min_{k,l:k\neq l}|x_k - x_l|^2 \leq C^2|x_i - x_{+,\min}|^2$. Therefore, $(f_\epsilon(x_i,1))_{i=1,\dots,n}$ satisfies the Lipschitz constraint.

The expression for $\omega(\epsilon)$ is obtained by plugging the solution into the objective function.

# B Additional Details for Section 4

## B.1 Differentiability of $\omega(\cdot)$ and $\rho(\cdot)$

The differentiability of $\omega(\cdot)$ at $\epsilon = 0$ follows from the results in Section 4. I apply Lemma A.3 in Appendix A.1 to show the differentiability of $\omega(\cdot)$ at any $\epsilon > 0$. Consider the problem (4). The norm constraint must hold with equality; otherwise, we can increase the objective by increasing $f(x_i,1)$ for all $i$ by a small amount. It follows that there exists $f_\epsilon \in \mathcal{F}_{\text{Lip}}(C)$ that attains the modulus of continuity with $\|\tilde{\boldsymbol{m}}(f_\epsilon)\| = \epsilon$. The differentiability of $\omega(\cdot)$ at any $\epsilon > 0$ then holds by Lemma A.3.

I next show the differentiability of $\rho(\cdot)$ at any $\epsilon \in \mathbb{R}$ by deriving its closed-form expression. Observe that $\rho(\epsilon)$ is the value of the following problem:

$$\max_{(f(x_i,0),f(x_i,1))_{i=1,...,n}\in\mathbb{R}^{2n}} \frac{1}{n}\sum_{i:c_1\leq x_i<c_0}[f(x_i,1)-f(x_i,0)] \tag{B.1}$$
$$s.t. \sum_{i=1}^{n} w_i^* \frac{f(x_i,d_i)}{\sigma(x_i,d_i)} = \epsilon,\ \ f(x_i,d)-f(x_j,d)\leq C|x_i-x_j|,\quad d\in\{0,1\}, i,j\in\{1,...,n\},$$

where $w_i^*$ is the $i$th element of $\boldsymbol{w}^*$ given in Section 4. Here I use the same argument used to obtain (4) in Section 4 and impose the Lipschitz constraint only among $x$'s in the sample.

The norm constraint $\sum_{i=1}^n w_i^* \frac{f(x_i,d_i)}{\sigma(x_i,d_i)} = \epsilon$ of (B.1) is equivalent to $-\sum_{i:c_1\leq x_i<c_0} f(x_i,0) + \tilde{n}f(x_{+,\min},1) = \bar{\sigma}\epsilon$, where $\bar{\sigma} = (\tilde{n}^2\sigma_{+,\min}^2 + \sum_{i:c_1\leq x_i<c_0}\sigma^2(x_i,0))^{1/2}$. Note that the objective and the norm constraint depend only on $((f(x_i,0),f(x_i,1))_{i:c_1\leq x_i<c_0}, f(x_{+,\min},1))$, so we can set the other values arbitrarily as long as the Lipschitz constraint holds. Given a value of $f(x_{+,\min},1)$, the objective is maximized only when $f(x_i,1)=C(x_{+,\min}-x_i)+f(x_{+,\min},1)$ for all $i$ with $c_1\leq x_i<c_0$. Therefore, the value of the problem (B.1) equals to the value of the following problem:

$$\max_{\boldsymbol{f}\in\mathbb{R}^{\tilde{n}+1}} \frac{C}{n}\sum_{i:c_1\leq x_i<c_0}(x_{+,\min}-x_i)+\frac{\tilde{n}}{n}f(x_{+,\min},1)-\frac{1}{n}\sum_{i:c_1\leq x_i<c_0} f(x_i,0)$$
$$s.t.\ -\sum_{i:c_1\leq x_i<c_0} f(x_i,0)+\tilde{n}f(x_{+,\min},1)=\bar{\sigma}\epsilon,$$
$$f(x_i,0)-f(x_j,0)\leq C|x_i-x_j|,\quad i,j\in\{k: c_1\leq x_k<c_0\},$$

where $\boldsymbol{f}$ denotes $((f(x_i,0))_{i:c_1\leq x_i<c_0}, f(x_{+,\min},1))$. By plugging the norm constraint into the objective, we obtain that the value of the problem is $\rho(\epsilon) = \frac{C}{n}\sum_{i:c_1\leq x_i<c_0}(x_{+,\min}-x_i)+\frac{\bar{\sigma}\epsilon}{n}$. Therefore, $\rho(\cdot)$ is differentiable for all $\epsilon\in\mathbb{R}$.

## B.2 Procedure for Computing $\epsilon^*\in\arg\max_{\epsilon\in[0,a^*]}\omega(\epsilon)\Phi(-\epsilon)$

The procedure is based on the first-order condition. By differentiating $\omega(\epsilon)\Phi(-\epsilon)$, we have $\omega'(\epsilon)\Phi(-\epsilon)-\omega(\epsilon)\phi(-\epsilon) = \left[\frac{1-\Phi(\epsilon)}{\phi(\epsilon)}-\frac{\omega(\epsilon)}{\omega'(\epsilon)}\right]\omega'(\epsilon)\phi(\epsilon)$. $\frac{1-\Phi(\epsilon)}{\phi(\epsilon)}$ is the Mills ratio of a standard normal random variable, which is strictly decreasing in $\epsilon$. Since $\omega(\epsilon)$

is nonnegative, nondecreasing, and concave, $\frac{\omega(\epsilon)}{\omega'(\epsilon)}$ is nondecreasing in $\epsilon$. Therefore, $\frac{1-\Phi(\epsilon)}{\phi(\epsilon)} - \frac{\omega(\epsilon)}{\omega'(\epsilon)}$ is strictly decreasing in $\epsilon$.

I suggest using the following procedure to compute $\epsilon^*$.

1. If $\frac{1-\Phi(a^*)}{\phi(a^*)} - \frac{\omega(a^*)}{\omega'(a^*)} > 0$, $\epsilon^* = a^*$.

2. If not, use the bisection method to find $\epsilon^* \in [0, a^*]$ that solves $\frac{1-\Phi(\epsilon)}{\phi(\epsilon)} - \frac{\omega(\epsilon)}{\omega'(\epsilon)} = 0$.

Note that, for each $\epsilon$, once we solve the problem (4) to compute $\omega(\epsilon)$ and $(f_\epsilon(x_i,0), f_\epsilon(x_i,1))_{i=1,\ldots,n}$, we can compute $\omega'(\epsilon)$ using the closed-form expression $\omega'(\epsilon) = \frac{\epsilon}{\tilde{\boldsymbol{m}}(\iota)'\tilde{\boldsymbol{m}}(f_\epsilon)} = \frac{\epsilon}{\frac{n}{\tilde{n}}\sum_{i=1}^n d_i f_\epsilon(x_i,d_i)/\sigma^2(x_i,d_i)}$, which is obtained by application of Lemma A.3 with $\iota(x,d) = \frac{n}{\tilde{n}}d$ for all $x \in \mathbb{R}$.

# C Additional Results

## C.1 Asymptotic Optimality When the Error Distribution is Unknown

In this section, I propose a feasible version of the minimax regret decision rule when the error distribution is unknown and establish its asymptotic optimality, closely following the asymptotic framework considered by Armstrong and Kolesár (2018).

Suppose that the sample $\boldsymbol{Y}_n = (Y_1, ..., Y_n)'$ is generated by the following model:

$$\boldsymbol{Y}_n = \boldsymbol{m}_n(\theta) + \boldsymbol{U}_n, \quad \theta \in \Theta,\ \boldsymbol{U}_n \sim Q \in \mathcal{Q}_n,$$

where $\boldsymbol{U}_n = (U_1, ..., U_n)'$ and $\mathcal{Q}_n$ is the set of possible distributions of $\boldsymbol{U}_n$. As in the main text, $\Theta$ is a known subset of a vector space $\mathbb{V}$, $\boldsymbol{m}_n : \mathbb{V} \to \mathbb{R}^n$ is a known linear function, and consider a binary policy choice problem in which $L(\theta)$ represents the welfare difference, where $L : \mathbb{V} \to \mathbb{R}$ is a known linear function. For notational simplicity, I hold $\Theta$ and $L$ fixed over $n$ and assume that the variance of $\boldsymbol{U}_n$ is $\sigma^2\boldsymbol{I}_n$ under any $Q \in \mathcal{Q}_n$ for some fixed, unknown $\sigma > 0$. It is straightforward to extend the result to the case where $\Theta$ and $L$ may change as $n$ increases and the variance of $\boldsymbol{U}_n$ may vary across $Q \in \mathcal{Q}_n$.

Regret of decision rule $\delta_n : \mathbb{R}^n \to [0,1]$ now depends on $\theta$ and $Q$:

$$R_n(\delta_n, \theta, Q) := \begin{cases} L(\theta)(1 - \mathbb{E}_{\theta,Q}[\delta_n(\boldsymbol{Y}_n)]) & \text{if } L(\theta) \geq 0, \\ -L(\theta)\mathbb{E}_{\theta,Q}[\delta_n(\boldsymbol{Y}_n)] & \text{if } L(\theta) < 0, \end{cases}$$

where $\mathbb{E}_{\theta,Q}$ denotes the expectation under $(\theta, Q)$. Let $\mathcal{R}_n(\Theta, \mathcal{Q}_n)$ denote the minimax risk:

$$\mathcal{R}_n(\Theta, \mathcal{Q}_n) := \inf_{\delta_n} \sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} R_n(\delta_n, \theta, Q).$$

I say that a sequence of decision rules $\delta_n$ is *asymptotically minimax regret* if

$$\lim_{n\to\infty} \frac{\mathcal{R}_n(\Theta, \mathcal{Q}_n)}{\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} R_n(\delta_n, \theta, Q)} = 1.$$

To construct a feasible decision rule, suppose we have some estimator $\hat{\sigma}_n(\boldsymbol{Y}_n)$ of $\sigma$, which is a function of the sample $\boldsymbol{Y}_n$. Then consider the minimax regret decision rule under normal errors in Theorem 1, where $\sigma$ is replaced with $\hat{\sigma}_n(\boldsymbol{Y}_n)$:

$$\hat{\delta}_n(\boldsymbol{Y}_n) = \begin{cases} \mathbf{1}\left\{\boldsymbol{m}_n(\theta_{n,\hat{\epsilon}_n})'\boldsymbol{Y}_n \geq 0\right\} & \text{if } \hat{\sigma}_n(\boldsymbol{Y}_n) > \sigma_n^*, \\ \mathbf{1}\left\{(\boldsymbol{w}_n^*)'\boldsymbol{Y}_n \geq 0\right\} & \text{if } \hat{\sigma}_n(\boldsymbol{Y}_n) = \sigma_n^*, \\ \Phi\left(\dfrac{(\boldsymbol{w}_n^*)'\boldsymbol{Y}_n}{((\sigma_n^*)^2 - \hat{\sigma}_n^2(\boldsymbol{Y}_n))^{1/2}}\right) & \text{if } \hat{\sigma}_n(\boldsymbol{Y}_n) < \sigma_n^*. \end{cases}$$

Here, $\omega_n(\epsilon) := \sup\{L(\theta) : \|\boldsymbol{m}_n(\theta)\| \leq \epsilon, \theta \in \Theta\}$ for $\epsilon \geq 0$, $\theta_{n,\epsilon} \in \arg\max_{\theta\in\Theta:\|\boldsymbol{m}_n(\theta)\|\leq\epsilon} L(\theta)$ for $\epsilon \geq 0$, $\boldsymbol{w}_n^* = \lim_{\epsilon\to 0} \frac{\boldsymbol{m}_n(\theta_{n,\epsilon})}{\|\boldsymbol{m}_n(\theta_{n,\epsilon})\|}$, $\epsilon_n^* \in \arg\max_{\epsilon\in[0,a^*\sigma]} \omega_n(\epsilon)\Phi(-\epsilon/\sigma)$, $\hat{\epsilon}_n := \hat{\epsilon}_n(\boldsymbol{Y}_n) \in \arg\max_{\epsilon\in[0,a^*\hat{\sigma}_n(\boldsymbol{Y}_n)]} \omega_n(\epsilon)\Phi(-\epsilon/\hat{\sigma}_n(\boldsymbol{Y}_n))$, and $\sigma_n^* := 2\phi(0)\frac{\omega_n(0)}{\omega_n'(0)}$.

I show that $\hat{\delta}_n$ is asymptotically minimax regret under suitable conditions. For a sequence of random variables $W_n$, I use the notation $W_n \underset{\Theta,\mathcal{Q}_n}{\overset{p}{\to}} w$, where $w$ is a constant, to denote the convergence in probability uniformly over $\theta \in \Theta$ and $Q \in \mathcal{Q}_n$, that is, for every $c > 0$, $\lim_{n\to\infty} \sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} \mathbb{P}_{\theta,Q}(|W_n - w| > c) = 0$. I use the notation $W_n \underset{\Theta,\mathcal{Q}_n}{\overset{d}{\to}} \mathcal{L}$, where $\mathcal{L}$ is a probability law, to denote the convergence in distribution uniformly over $\theta \in \Theta$ and $Q \in \mathcal{Q}_n$, that is, $\lim_{n\to\infty} \sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} \sup_{w\in\mathbb{R}} |\mathbb{P}_{\theta,Q}(W_n \leq w) - F_{\mathcal{L}}(w)| = 0$, where $F_{\mathcal{L}}$ is the cumulative distribution function of the law $\mathcal{L}$, which is assumed to be continuous.

**Assumption C.1** (Regularity for Asymptotic Optimality)**.**

*(i)* $\sigma^* := \lim_{n\to\infty} \sigma_n^*$ *exists.*

*(ii)* $\hat{\sigma}_n(\boldsymbol{Y}_n) \underset{\Theta,\mathcal{Q}_n}{\overset{p}{\to}} \sigma$.

*(iii) For any sufficiently large $n$, the following holds with probability one under any $(\theta, Q) \in \Theta \times \mathcal{Q}_n$: there exists $\theta_{n,\hat{\epsilon}_n}$ that attains the modulus of continuity at $\hat{\epsilon}_n$.*

*(iv) Either of the following holds.*

*(a)* $\sigma > \sigma^*$ *and* $\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{U}_n + (\boldsymbol{m}_n(\theta_{n,\hat{\epsilon}_n}) - \boldsymbol{m}_n(\theta_{n,\epsilon_n^*}))\boldsymbol{Y}_n}{\sigma\epsilon_n^*} \underset{\Theta,\mathcal{Q}_n}{\overset{d}{\to}} \mathcal{N}(0, 1)$.

*(b)* $\sigma < \sigma^*$ *and* $\frac{(\boldsymbol{w}_n^*)'\boldsymbol{U}_n + \hat{\xi}_n(\boldsymbol{Y}_n)}{\sigma_n^*} \underset{\Theta,\mathcal{Q}_n}{\overset{d}{\to}} \mathcal{N}(0, 1)$*, where $\hat{\xi}_n(\cdot) : \mathbb{R}^n \to \mathbb{R}$ is a sequence of random functions independent of $\boldsymbol{U}_n$ whose distribution does not depend on $(\theta, Q)$, such that $\hat{\xi}_n(\boldsymbol{y}_n) \sim \mathcal{N}(0, (\sigma_n^*)^2 - \hat{\sigma}^2(\boldsymbol{y}_n))$ for all $\boldsymbol{y}_n \in \mathbb{R}^n$.*

*(v) For any sufficiently large $n$, there exists $Q \in \mathcal{Q}_n$ such that $\boldsymbol{U}_n \sim \mathcal{N}(\boldsymbol{0}, \sigma^2\boldsymbol{I}_n)$ under $Q$.*

*(vi)* $\sup_{\theta\in\Theta} |L(\theta)| < \infty$.

The key conditions are Assumption C.1(iv) and (v). The former assumes the uniform asymptotic normality of the weighted sum of the errors $\boldsymbol{U}_n$ (plus an error due to the estimation of $\sigma$ when $\hat{\delta}_n$ is nonrandomized and a mean-zero normal noise $\hat{\xi}_n(\boldsymbol{Y}_n)$ when $\hat{\delta}_n$ is randomized). This assumption is used to show that the maximum regret of $\hat{\delta}_n$ over $\theta \in \Theta$ and $Q \in \mathcal{Q}_n$ converges to the minimax risk under normal errors in Theorem 1. The uniform asymptotic normality can be established, for example, under conditions on the weight vector $\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})}{\sigma\epsilon_n^*}$ or $\frac{\boldsymbol{w}_n^*}{\sigma_n^*}$ à la the Lindeberg condition for the central limit theorem. In the context of uniform asymptotic valid inference on RD parameters, such conditions are directly imposed (Imbens and Wager, 2019) or verified under low-level conditions (Armstrong and Kolesár, 2018).

Assumption C.1(v) states that the class of possible distributions contains a normal distribution. Under this assumption, the minimax risk under normal errors is no greater than the minimax risk over the class of possible distributions. Since the maximum regret of $\hat{\delta}_n$ is shown to converge to the former, it also converges to the latter, establishing the asymptotic optimality of $\hat{\delta}_n$. Similar assumptions are often used in the literature on minimax estimation and inference in nonparametric regression models (Fan, 1993; Armstrong and Kolesár, 2018).

Let $\mathcal{R}_n(\Theta, \mathcal{N}) = \inf_{\delta_n} \sup_{\theta\in\Theta} R_n(\delta_n, \theta, Q_{\mathcal{N}})$ denote the minimax risk under normal errors, where $Q_{\mathcal{N}}$ is $\mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I}_n)$. Below, I focus on the case where the sequence $\mathcal{R}_n(\Theta, \mathcal{N})$ is bounded below by a positive constant, which holds when the welfare difference is not identified for any sample size or in the limit.

**Theorem C.1** (Asymptotic Optimality)**.** *Suppose that the conditions of Theorem 1 hold for all $Q \in \mathcal{Q}_n$ and any sufficiently large $n$, that $\mathcal{R}_n(\Theta, \mathcal{N})$ is bounded below by a positive constant, and that Assumption C.1 holds. Then, $\hat{\delta}_n$ is asymptotically minimax regret.*

*Proof.* Since $Q_{\mathcal{N}} \in \mathcal{Q}_n$,

$$\mathcal{R}_n(\Theta, \mathcal{N}) \leq \inf_{\delta_n} \sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} R_n(\delta_n, \theta, Q) = \mathcal{R}_n(\Theta, \mathcal{Q}_n). \tag{C.1}$$

Below, I show that

$$\sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} R_n(\hat{\delta}_n, \theta, Q) \leq \mathcal{R}_n(\Theta, \mathcal{N}) + c_n \tag{C.2}$$

for some positive constant $c_n \to 0$. Once it is done, combining inequalities (C.1) and (C.2) yields

$$\frac{\mathcal{R}_n(\Theta, \mathcal{N})}{\mathcal{R}_n(\Theta, \mathcal{N}) + c_n} \leq \frac{\mathcal{R}_n(\Theta, \mathcal{Q}_n)}{\sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} R_n(\hat{\delta}_n, \theta, Q)} \leq 1.$$

The left-hand side converges to one as $n \to \infty$ since $\mathcal{R}_n(\Theta, \mathcal{N})$ is bounded below by a positive constant, which proves that $\hat{\delta}_n$ is asymptotically minimax regret.

In what follows, I show that (C.2) holds for the case where $\sigma > \sigma^* = \lim_{n\to\infty} \sigma_n^*$. The proof for the case where $\sigma < \sigma^*$ is omitted since it is similar, noting that we can write $\hat{\delta}_n(\boldsymbol{Y}_n) = \mathbb{P}_{\theta,Q}\left((\boldsymbol{w}_n^*)'\boldsymbol{Y}_n + \hat{\xi}_n(\boldsymbol{Y}_n) \geq 0 | \boldsymbol{Y}_n\right)$ when $\hat{\sigma}_n(\boldsymbol{Y}_n) < \sigma_n^*$.

Suppose $\sigma > \sigma^*$. Let $\tilde{\delta}_n(\boldsymbol{Y}_n) = \mathbf{1}\{\boldsymbol{m}_n(\theta_{n,\hat{\epsilon}_n})'\boldsymbol{Y}_n \geq 0\}$. By construction, $\sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} \mathbb{P}_{\theta,Q}\left(\hat{\delta}_n(\boldsymbol{Y}_n) \neq \tilde{\delta}_n(\boldsymbol{Y}_n)\right) \leq \sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} \mathbb{P}_{\theta,Q}\left(\hat{\sigma}_n(\boldsymbol{Y}_n) \leq \sigma_n^*\right)$. The right-hand side converges to zero since $\hat{\sigma}_n(\boldsymbol{Y}_n) \underset{\Theta,\mathcal{Q}_n}{\overset{p}{\to}} \sigma > \sigma^* = \lim_{n\to\infty} \sigma_n^*$.

Consider the following inequality:

$$\begin{aligned}\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)(1-\mathbb{E}_{\theta,Q}[\hat{\delta}_n(\boldsymbol{Y}_n)]) \leq &\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)\left(1-\mathbb{E}_{\theta,Q}\left[\tilde{\delta}_n(\boldsymbol{Y}_n)\right]\right)\\ &+\left(\sup_{\theta\in\Theta}|L(\theta)|\right)\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n}\mathbb{E}_{\theta,Q}\left[|\hat{\delta}_n(\boldsymbol{Y}_n)-\tilde{\delta}_n(\boldsymbol{Y}_n)|\right].\end{aligned}$$

The second term on the right-hand side is bounded above by $(\sup_{\theta\in\Theta}|L(\theta)|)\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n}\mathbb{E}_{\theta,Q}\left[\mathbf{1}\left\{\hat{\delta}_n(\boldsymbol{Y}_n)\neq\tilde{\delta}_n(\boldsymbol{Y}_n)\right\}\right]$, which converges to zero as $n\to\infty$ since $\sup_{\theta\in\Theta}|L(\theta)|<\infty$ and $\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n}\mathbb{P}_{\theta,Q}\left(\hat{\delta}_n(\boldsymbol{Y}_n)\neq\tilde{\delta}_n(\boldsymbol{Y}_n)\right)\to 0$. It follows that for some positive constant $c_n\to 0$,

$$\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)(1-\mathbb{E}_{\theta,Q}[\hat{\delta}_n(\boldsymbol{Y}_n)]) \leq \sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)\left(1-\mathbb{E}_{\theta,Q}\left[\tilde{\delta}_n(\boldsymbol{Y}_n)\right]\right)+c_n. \tag{C.3}$$

Next, we bound $\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)\left(1-\mathbb{E}_{\theta,Q}\left[\tilde{\delta}_n(\boldsymbol{Y}_n)\right]\right)$. Let $Z_n = \frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{U}_n+(\boldsymbol{m}_n(\theta_{n,\hat{\epsilon}_n})-\boldsymbol{m}_n(\theta_{n,\epsilon_n^*}))\boldsymbol{Y}_n}{\sigma\epsilon_n^*}$. We have

$$\begin{aligned}&\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)(1-\mathbb{E}_{\theta,Q}[\tilde{\delta}_n(\boldsymbol{Y}_n)]) = \sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)\left(1-\mathbb{P}_{\theta,Q}(\boldsymbol{m}_n(\theta_{n,\hat{\epsilon}_n})'\boldsymbol{Y}_n\geq 0)\right)\\ &=\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)\left(1-\mathbb{P}_{\theta,Q}\left((\boldsymbol{m}_n(\theta_{n,\hat{\epsilon}_n})-\boldsymbol{m}_n(\theta_{n,\epsilon_n^*}))\boldsymbol{Y}_n+\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{Y}_n\geq 0\right)\right)\\ &=\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n} L(\theta)\left(1-\mathbb{P}_{\theta,Q}\left(-Z_n\leq\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{m}_n(\theta)}{\sigma\epsilon_n^*}\right)\right)\\ &\leq\sup_{\theta\in\Theta} L(\theta)\left(1-\Phi\left(\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{m}_n(\theta)}{\sigma\epsilon_n^*}\right)\right)\\ &\quad+\left(\sup_{\theta\in\Theta}|L(\theta)|\right)\\ &\qquad\times\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n}\left|\mathbb{P}_{\theta,Q}\left(-Z_n\leq\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{m}_n(\theta)}{\sigma\epsilon_n^*}\right)-\Phi\left(\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{m}_n(\theta)}{\sigma\epsilon_n^*}\right)\right|.\end{aligned}$$

The second term in the last expression is bounded above by $(\sup_{\theta\in\Theta}|L(\theta)|)\sup_{\theta\in\Theta,Q\in\mathcal{Q}_n}\sup_{z\in\mathbb{R}}|\mathbb{P}_{\theta,Q}(-Z_n\leq z)-\Phi(z)|$, which converges to zero as $n\to\infty$ since $\sup_{\theta\in\Theta}|L(\theta)|<\infty$ and $Z_n \underset{\Theta,\mathcal{Q}_n}{\xrightarrow{d}} \mathcal{N}(0,1)$.

Note that the first term $\sup_{\theta\in\Theta} L(\theta)\left(1-\Phi\left(\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{m}_n(\theta)}{\sigma\epsilon_n^*}\right)\right)$ is the maximum regret of decision rule $\delta_n^*(\boldsymbol{Y}_n) = \mathbf{1}\left\{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{Y}_n \geq 0\right\}$ under normal errors. $\delta_n^*$ is minimax regret under normal errors if $\sigma > \sigma_n^*$. Since $\sigma > \sigma^* = \lim_{n\to\infty}\sigma_n^*$, $\delta_n^*$ is minimax regret under normal errors for any sufficiently large $n$, which implies that $\sup_{\theta\in\Theta} L(\theta)\left(1 - \Phi\left(\frac{\boldsymbol{m}_n(\theta_{n,\epsilon_n^*})'\boldsymbol{m}_n(\theta)}{\sigma\epsilon_n^*}\right)\right) = \mathcal{R}_n(\Theta, \mathcal{N})$ for any sufficiently large $n$. Together with (C.3), the above argument implies that for some positive $c_n \to 0$,

$$\sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} L(\theta)(1 - \mathbb{E}_{\theta,Q}[\hat{\delta}_n(\boldsymbol{Y}_n)]) \leq \mathcal{R}_n(\Theta, \mathcal{N}) + c_n.$$

The same argument follows for $\sup_{\theta\in\Theta, Q\in\mathcal{Q}_n} -L(\theta)\mathbb{E}_{\theta,Q}[\hat{\delta}_n(\boldsymbol{Y}_n)]$, so that inequality (C.2) holds as desired. □

## C.2 Comparison with Hypothesis Testing Rules

Hypothesis testing can be viewed as an alternative procedure for deciding between two policies. Here, I compare the minimax regret rule with a class of hypothesis testing rules. To define it, suppose $\Theta$ is convex and centrosymmetric, and consider testing

$$H_0 : L(\theta) \leq -b \text{ and } \theta \in \Theta \quad \text{vs.} \quad H_1 : L(\theta) \geq b \text{ and } \theta \in \Theta$$

for some $b > 0$. Let $\theta^{(b)}$ solve $\inf_{\theta\in\Theta:L(\theta)\geq b} \|\boldsymbol{m}(\theta)\|$. For any level $\alpha > 0$, the minimax test, which has the largest minimum power under $H_1$, is given by the Neyman-Pearson test of $H_0 : \theta = -\theta^{(b)}$ vs. $H_1 : \theta = \theta^{(b)}$ (Armstrong and Kolesár, 2018, Lemma A.2). It rejects $H_0$ if the test statistic $\boldsymbol{m}(\theta^{(b)})'\boldsymbol{Y}$ is greater than its $1-\alpha$ quantile under $-\theta^{(b)}$. Since $\boldsymbol{m}(\theta^{(b)})'\boldsymbol{Y} \sim \mathcal{N}(-\|\boldsymbol{m}(\theta^{(b)})\|^2, \sigma^2\|\boldsymbol{m}(\theta^{(b)})\|^2)$ under $-\theta^{(b)}$, the critical value is $-\|\boldsymbol{m}(\theta^{(b)})\|^2 + z_{1-\alpha}\sigma\|\boldsymbol{m}(\theta^{(b)})\|$, where $z_{1-\alpha}$ is the $1-\alpha$ quantile of a standard normal random variable. The level-$\alpha$ minimax test is then given by

$$\delta_{\alpha,b}(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\theta^{(b)})'\boldsymbol{Y} \geq -\|\boldsymbol{m}(\theta^{(b)})\|^2 + z_{1-\alpha}\sigma\|\boldsymbol{m}(\theta^{(b)})\|\}.$$

I call such tests *hypothesis testing rules.*

Are there any hypothesis testing rules that exactly match the minimax regret rule? Let $\epsilon^* > 0$ solve $\max_{\epsilon\in[0,a^*\sigma]} \omega(\epsilon)\Phi(-\epsilon/\sigma)$, and let $\theta_{\epsilon^*}$ solve the modulus of continuity at $\epsilon^*$ with $\|\boldsymbol{m}(\theta_{\epsilon^*})\| = \epsilon^*$. By the duality of the problem, $\theta_{\epsilon^*}$ also solves $\inf_{\theta\in\Theta:L(\theta)\geq b^*} \|\boldsymbol{m}(\theta)\|$, where $b^* = \omega(\epsilon^*)$. Let $\alpha^*$ satisfy $-\|\boldsymbol{m}(\theta_{\epsilon^*})\| + z_{1-\alpha^*}\sigma = 0$, i.e.,

$\alpha^* = \Phi(-\epsilon^*/\sigma)$, so that the critical value is zero. For this choice of $\alpha^*$ and $b^*$, the hypothesis testing rule is $\delta_{\alpha^*,b^*}(\boldsymbol{Y}) = \mathbf{1}\{\boldsymbol{m}(\theta_{\epsilon^*})'\boldsymbol{Y} \geq 0\}$, which is identical to the minimax regret rule. Since $\epsilon^* \leq a^*\sigma$, we can obtain a lower bound on $\alpha^*$: $\alpha^* = \Phi(-\epsilon^*/\sigma) \geq \Phi(-a^*) \approx 0.226$. Therefore, the minimax regret rule is less conservative in rejection of the null hypothesis than hypothesis testing rules that use conventional levels such as 0.01 and 0.05. This is consistent with the fact that the minimax regret criterion takes into consideration the potential welfare loss as well as the probability of making a wrong choice.

# D Empirical Policy Application: Additional Figures

Figure D.1: Optimal Decisions for Alternative New Policies

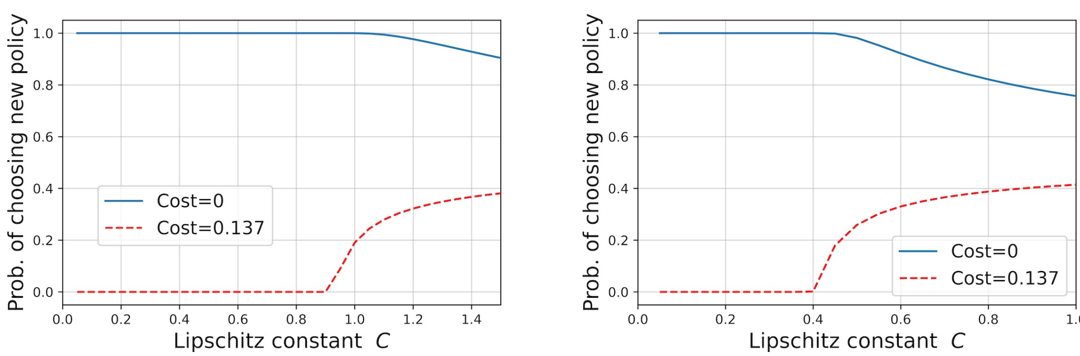


(a) Constructing Schools in 10% of Villages

(b) Constructing Schools in 30% of Villages

*Notes*: This figure shows the probability of choosing the new policy computed by the minimax regret rule. The new policy is to construct BRIGHT schools in previously ineligible villages whose relative scores are in the top 10% (Panel (a)) or in the top 30% (Panel (b)). The solid line shows the results for the scenario where we ignore the policy cost. The dashed line shows the results for the scenario where the policy cost measured in the unit of the enrollment rate is 0.137. I report the results for the range $[0.05, 0.1, ..., 1.45, 1.5]$ of the Lipschitz constant $C$ in Panel (a) and for the range $[0.05, 0.1, ..., 0.95, 1]$ in Panel (b).

Figure D.2: Weight to Each Village Attached by Plug-in Rules

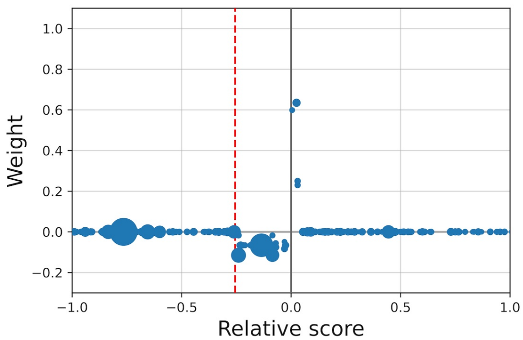


(a) Minimax MSE, $C = 0.5$

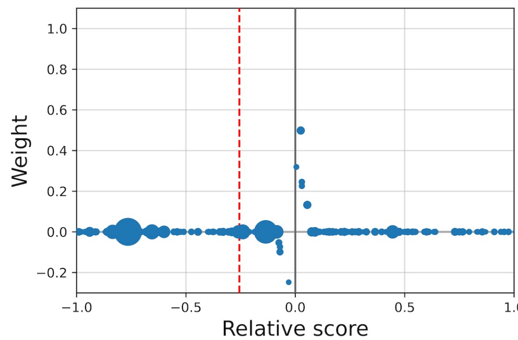


(b) Minimax MSE Under Constant Effect, $C = 0.5$

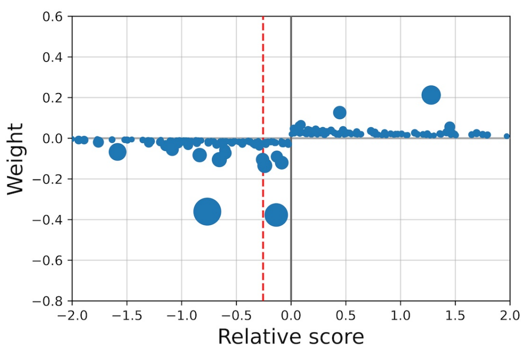


(c) Polynomial of Degree 2

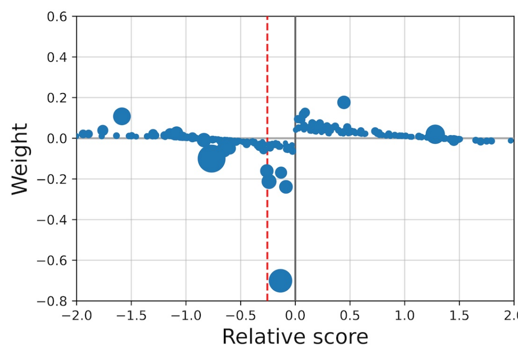


(d) Polynomial of Degree 4

*Notes*: This figure shows the weight $w_i$ attached to each village by the plug-in decision rules of the form $\delta(\boldsymbol{Y}) = \mathbf{1}\{\sum_{i=1}^n w_i Y_i \geq 0\}$. The weights are normalized so that $\sum_{i=1}^n w_i^2 = 1$. The horizontal axis indicates the relative score of each village. Each circle corresponds to each village. The size of circles is proportional to the inverse of the standard error of the enrollment rate $Y_i$. The vertical dashed line corresponds to the new cutoff $-0.256$. Panels (a) and (b) show the results for the plug-in rules based on the linear minimax MSE estimators with or without the assumption of constant conditional treatment effects when the Lipschitz constant $C$ is 0.5. Panels (c) and (d) show the results for the plug-in rules based on the polynomial regression estimators of degrees 2 and 4, respectively.

# References


Armstrong, T. B. and Kolesár, M. (2018). Optimal Inference in a Class of Regression Models. *Econometrica*, **86** (2), 655–683.

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis; 2nd ed.* Springer Series in Statistics, New York: Springer.

Donoho, D. L. (1994). Statistical Estimation and Optimal Recovery. *The Annals of Statistics*, **22** (1), 238–270.

Fan, J. (1993). Local Linear Regression Smoothers and Their Minimax Efficiencies. *The Annals of Statistics*, **21** (1), 196–216.

Imbens, G. and Wager, S. (2019). Optimized Regression Discontinuity Designs. *The Review of Economics and Statistics*, **101** (2), 264–278.

Karlin, S. and Rubin, H. (1956). The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio. *The Annals of Mathematical Statistics*, **27** (2), 272–299.

Stoye, J. (2009). Minimax Regret Treatment Choice with Finite Samples. *Journal of Econometrics*, **151** (1), 70–81.

— (2012). Minimax Regret Treatment Choice with Covariates or with Limited Validity of Experiments. *Journal of Econometrics*, **166** (1), 138–156.

Tetenov, A. (2012). Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria. *Journal of Econometrics*, **166** (1), 157–165.