Optimal Recommendation of Individual Dose Intervals

by

Xiaomao Li

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Statistics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2018

Date of final oral examination: 06/04/2018

The dissertation is approved by the following members of the Final Oral Committee:

Guanhua Chen, Assistant Professor, Department of Biostatistics & Medical Informatics

Shao Jun, Professor, Department of Statistics

Sundus Keles, Professor, Department of Statistics

Menggang Yu, Professor, Department of Biostatistics and Medical Informatics

Jun Zhu, Professor, Department of Statistics and Department of Entomology

Abstract

The recent years witnessed a surge in the intellectual interest of statistical methods for personalized dosing. It is important to investigate the Individual Dose Interval (IDI) both for medical applications and statistical researches. In this dissertation, a new method based on Outcome-Weighted-Learning (OWL) is proposed to determine two types of IDIs with different properties: Probability Dose Interval (PDI) and Expectation Dose Interval (EDI). The former identifies a range of doses where the potential outcome is greater than a threshold with a certain probability and the latter identifies a range of doses where the expectation of the potential outcome is greater than a threshold. The method adopts Reproducing Kernel Hilbert Space (RKHS) estimators from non-convex loss functions, which are solved by the Difference-of-Convex (DC) algorithm. The Fish consistency and the convergence rates of the estimators are discussed. Numerical simulations show an advantage of the proposed method over methods based on traditional outcome modeling. In addition, this new method is applied to determine personalized Hemoglobin A1c (HbA1c) control intervals for diabetic patients.

Contents

List of Tables

List of Figures

Acknowledgments

My last five years at the University of Wisconsin-Madison have been very rewarding. This accomplishment would not have been possible without the help of many people.

Chapter 1

Introduction

1.1 The Need for Personalized Interval Dosing in Practice

Personalized medicine is playing an important role in clinical researches and applications. A common belief in many medical fields is that personalized prescription based on patients' health condition and medical history will result in more effective treatment. Despite the abundance of research on personalization of categorical treatments, which aims to choose the optimal treatment out of several candidates based on patients' features, personalization of dosing has not been studied sufficiently. This is related to the inadequacy of clinical experiments targeting this area, as randomized clinical trials, the gold standard in medical research, typically engages only one dosage level. A typical dose-finding trial is often constructed in the double-blind Phase II trial to identify the no-effect, the main effect, and the maximal effective doses (Chevret, 2006). Such practices often ignore the heterogeneity of patients' reaction to the doses, and in fact, the FDA-approved doses may not work well for all the patients.

In practice, physicians may not always want to give prescriptions of fixed doses. Due to the complex nature of human's reaction to chemicals, which might vary dramatically across patients, it is difficult or undesirable to specify a fixed dose for all the situations. Instead, the dosage is sometimes prescribed to several candidate levels or to a range where physicians

can make decisions depending on the severity of the disease and patients' tolerance to side effects. Such practices are mainly driven by the demands for the maximum-tolerated doses. For example, cancer patients who survive 5 years after treatment are usually viewed as successful cases. It would be useful to know about the least amount of doses that can achieve the comparable anticipated lifetime, especially given the fact that cancer therapies generally entail high costs and cause damage to human organs. As another example, beta blockers are usually adopted to alleviate the symptoms of heart premature contraction, but overdosing may incur potential side effects. For almost any sorts of prescriptions, it would be helpful to learn about the maximum-tolerated dose with minimum side effects for patients.

In some other situations, it is impossible to enforce an exact dose of medication due to partial compliance. The HbA1c, referred to as glycated hemoglobin, is formed when hemoglobin, a protein within the red blood cells that carries oxygen throughout the body, combines with glucose in the blood and thus becomes glycated. By measuring HbA1c, clinicians are able to get an overall picture of what average blood sugar levels of the patients have been over the past three months. Diabetes patients usually try to keep their HbA1c level within a certain range, usually 4 to 7 mmol/l before meals and 4 to 8.5 mmol/l after the meals (Weykamp, 2013). But it is not very likely to be maintained at an exact value because the blood sugar cannot be accurately controlled only by medications, diets, and exercises. Moreover, patients with different physiological conditions need to control HbA1c in different ways (Radin, 2014). For instance, over-controlling blood sugar for older patients will lead to a compromised quality of life. In this case, how to provide personalized recommendations of HbA1c control intervals for patients with different conditions becomes an important question. In Chapter 7, we will tackle this problem with our proposed approaches using a large dataset.

1.2 Personalization of Discrete Treatment

The field of personalized medicine has received considerable attention over time. Unlike the personalized dosing problem, there have been well-developed lines of research on personalizing a binary treatment using the information on patients' characteristics. The most popular and intuitive statistical methods involve a two-step procedure based on outcome modeling (Zhao et al., 2009, Moodie et al., 2012, Chakraborty and Moodie, 2013). This framework sometimes referred to as the indirect methods, can be described as a two-step procedure. In the first step, an outcome model is developed to model the conditional expectation of outcome given treatment, covariates, and their interactions. The model could be parametric with Lasso-type regularization (Tibshirani, 1996), or nonparametric, such as tree-based methods (Breiman, 2017, 2001) and Support Vector Machine (Vapnik, 2013, Steinwart and Christmann, 2008). In the second step, the treatment option which maximizes the expected outcome given a patient's covariates is set to be the optimal treatment for the patient. Alternative methods based on direct modeling of regret or contrast have also been developed, which are believed to have better performance facing low signal-to-noise ratio (Moodie et al., 2009, Henderson et al., 2010, Wallace and Moodie, 2015, Liang and Yu, 2018). Although the current practices of the indirect methods mainly focus on categorical treatment, the same procedure can be utilized to handle continuous treatment (Chen et al., 2016).

However, the outcome model is susceptible to overfitting when the sample size is small. In such cases, there might be some advantages to directly estimate the Individual Treatment Rule (ITR), which often has a simpler form than the outcome model (Zhao et al., 2012). This framework is referred to as the Outcome-Weighted-Learning (OWL), which has been extended to handle multiple stages (Zhao et al., 2015), ordinal treatment (Chen et al., 2017), and multi-category treatment (Lou et al., 2017). From here, we denote Y as the outcome, $A \in \{0, 1\}$ as the binary treatment for simplicity, and $X \in \mathcal{X}$ as the d-dimensional covariates, while y, a and x are specific values of Y, A and X. Also, denote the ITR as a function

that maps patients' features to the binary treatment treatments, i.e., $f(x) \in \{0, 1\}$ for all $x \in \mathcal{X}$. The classic OWL model aims at solving a weighted binary classification problem as below.

$$
\begin{aligned}
f_{ITR} &= \arg \min_{f(x)\in\{0,1\}} E\Big[\frac{Y1\{f(X) \neq A\}}{P(A \,|\, X)}\Big] \\
&= \arg \max_{f(x)\in\{0,1\}} E\Big[\frac{Y1\{f(X) = A\}}{P(A \,|\, X)}\Big] \\
&= \arg \max_{f(x)\in\{0,1\}} E_X\Big[\int E[Y|A = a, X]1\{f(x) = a\}da\Big] \\
&= \arg \max_{f(x)\in\{0,1\}} E_X\Big[1\{f(X) = 1\}E[Y|A = 1, X] + 1\{f(X) = 0\}E[Y|A = 0, X]\Big] \\
&= \arg \max_{f(x)\in\{0,1\}} E_X\Big[1\{f(X) = 1\}\Big(E[Y|A = 1, X] - E[Y|A = 0, X]\Big) + E[Y|A = 0, X]\Big]
\end{aligned}
\tag{1.1}
$$

where $P(A \,|\, X)$ is the probability of observing A given the covariates X, often referred to as the propensity score. For a fixed $X = x$, this is equivalent to maximizing

$$1\{f(x) = 1\}\big(E[Y \mid A = 1, X = x] - E[Y \mid A = 0, X = x]\big) + E[Y \mid A = 0, X = x]$$

Therefore, the optimal treatment for patient x is

$$a^*(x) = \arg \max_{a\in\{0,1\}} E[Y|A = a, X = x]$$

1.3 Personalization of Continuous Treatment

As introduced in the previous section, there are well-developed statistical methods for personalized treatment assignment when the treatment is from two or more categories. But less attention has been paid to the continuous treatment, i.e., the dosage of a certain medication. The pioneer work of Chen et al. (2016) adopts the OWL framework to estimate the Individual Dosing Rule (IDR) that maximizes the conditional expectation of the response given covariates. In addition, Chen et al. (2016) discusses the usage of various indirect methods for problems of personalized dosing and demonstrates the superiority of OWL-based

approaches through extensive simulation and real data analysis.

The IDR from Chen et al. (2016) recommends a dose for each patient so that the patient's expected outcome given the recommended dose and covariates is maximized. Assuming now $A \in [a_L, a_U]$ is a bounded continuous dose variable and $Y(a, x)$ is the potential outcome (Rubin, 1974) given dose a and covariates x, we denote the $f : \mathcal{X} \mapsto [a_L, a_U]$ as the IDR Chen et al. (2016) which solves the following problem:

$$\begin{aligned} f_{IDR} &= \arg \max_{f(x)\in[a_L,a_U]} E\big[Y(f(X), X)\big] \\ &= \arg \max_{f(x)\in[a_L,a_U]} E_X\big[E[Y|A = f(X), X]\big] \\ &= \arg \max_{f(x)\in[a_L,a_U]} \lim_{\phi\to 0} E_X\left[\frac{E\big[Y1\big\{A \in \big(f(X) - \phi, f(X) + \phi\big)\big\}/P(A \,|\, X)\big|X\big]}{2\phi}\right] \end{aligned} \tag{1.2}$$

where ϕ is a hyperparameter in finite sample and is required to be specified to balance the bias and variance of the estimated IDR. The inverse weighting probability $P(A \,|\, X)$ is essentially a conditional density of A given X which is assumed to be known in experimental settings and has to be estimated in the observational settings. The equality holds under the assumption of ignorability, that the potential outcome is independent with the dose assignment given the covariates, i.e., $Y(\cdot, x) \perp\!\!\!\perp A \,\big|\, X = x$.

In the finite-sample setting, the indicator function in (1.2) is relaxed by a non-convex but continuous surrogate loss, which can be decomposed to the difference of two convex losses. Then the Difference-of-Convex (DC) (Le Thi Hoai and Tao, 1997) algorithm is adopted to solve the problem. Chen et al. (2016) demonstrates the small sample advantage of their method over the indirect methods based on the traditional outcome modeling when the received dose is relatively close to the optimal dose. The method that we are going to propose in Chapter 2 is partially inspired by Chen et al. (2016) and they will complement each other when applied together.

There are two main limitations of Chen et al. (2016) from the application aspects. First, it is unknown how to capture the sampling error of the estimated IDR as the statistical

inference is difficult due to the non-convexity of the objective function and the iterative nature of the DC algorithm. Second, there are cases where the estimated IDR might not be the best option, if other factors are taken into consideration, such as expenses and side effects. Especially when the outcome is insensitive to the dose, a non-optimal dose might nonetheless achieve comparable outcomes and thus it may be meaningful to look for the least-sufficient or maximum-tolerated dose. As in Fig 1.1, a hypothetical relationship between the response and the dose is graphed. The red dot indicates the optimal dose, while the orange dots signify potential estimated peaks. If the threshold is considered to be the yellow dotted line, then all potential estimated peaks would be contained within the lower bound and the upper bound marked by two green lines. Also, the lower bound itself might be a rational choice as it achieves a comparable outcome with a lower dose.

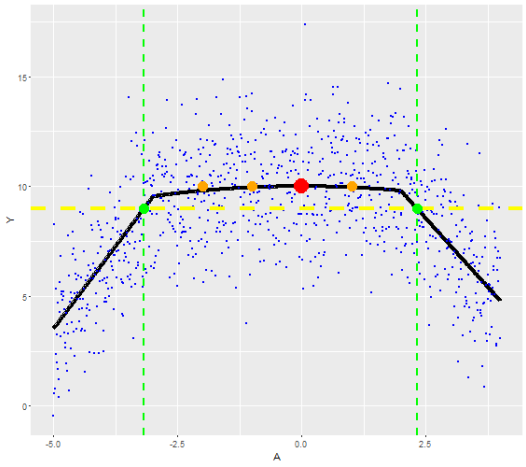

Figure 1.1: A Hypothetical Dose-response Curve

To capture the dose-response relationship in a holistic manner, we define an Individual Dose Interval (IDI) and we anticipate that the doses inside the interval are desirable whereas doses outside the interval are undesirable. Each IDI is connected with a threshold S separating good outcomes from bad ones. And a confidence probability α reflects how strongly the desirable interval is associated with the good outcomes. In the remaining of this dissertation, let A be a continuous dose with bounded support $a \in [a_L, a_U]$ and $Y(a, x)$ be the potential outcome given the dose a and the patient features x. Also, let $f_L(x) \in [a_L, a_U]$

be the lower bound of IDI and $f_U(x) \in [a_L, a_U]$ be the upper bound of IDI for patient x such that $f_L(x) < f_U(x)\ \forall x \in \mathcal{X}$. This dissertation aims to determine the following two types of IDIs.

Probability Dose Interval (PDI): (1.3)

$$\big[f_L(x), f_U(x)\big] : \forall a^* \in [f_L(x), f_U(x)],\ \ Y(a^*, x) > S \text{ with probability at least } \alpha$$

Expectation Dose Interval (EDI): (1.4)

$$\big[f_L(x), f_U(x)\big] : \forall a^* \in [f_L(x), f_U(x)],\ \ E[Y(a^*, x)] > S$$

If a patient follows the PDI recommendation, the patient is guaranteed to have the outcome above S with a probability of at least α. Similarly, a patient who follows the EDI recommendation will have the expected outcome above S. As in Rubin (1974), the potential outcome is only observed for a single treatment level, but we will discuss in Chapter 3 the assumptions required for linking the potential outcomes to estimable quantities. In the course of this dissertation, we denote the Individual Dose Interval Rule (IDIR) as the two bound functions of patients' features, which includes Probability Dose Interval Rule (PDIR) and Expectation Dose Interval Rule (EDIR), functions that map $x \in \mathcal{X}$ to (1.3) and (1.4).

1.4 Indirect Methods Based on Outcome Modeling

Classification and regression methods model the conditional distributions of discrete and continuous outcomes given covariates. In the following section, we will show that they can be adapted to estimate the IDI with the properties described in the previous section. Generally, the PDI can be estimated using classification methods whereas the EDI can be estimated using regression methods. This adaptation of classification and regression methods shares the same spirits with the other indirect methods that are used to estimate the the ITR (see Section 1.2) and the IDR (see Section 1.3). The intuition behind these approach is that since in the definition of (1.3) and (1.4), the potential outcome is never

observed, we can replace it by sample estimates, under the ignoriability assumption which we will discuss in Section 3.1. The indirect methods essentially look for following interval for patient $x \in \mathcal{X}$.

Indirect PDI: (1.5)

$$\big[f_L(x), f_U(x)\big] : \forall a^* \in [f_L(x), f_U(x)], \;\; \hat{P}\big(Y > S \,|\, A = a^*, X = x\big) \geq \alpha$$

Indirect EDI: (1.6)

$$\big[f_L(x), f_U(x)\big] : \forall a^* \in [f_L(x), f_U(x)], \;\; \hat{E}[Y|A = a^*, X = x] > S$$

This approach involves two steps. In the first step, an outcome model is constructed to attain the estimated conditional probability $\hat{P}\big(Y > S \,|\, A = a, X = x\big)$ or the estimated conditional expectation $\hat{E}[Y|A = a, X = x]$ for any $a \in [a_L, a_U]$. In the second step, a grid search is done to find $[f_L(x), f_U(x)]$. We will illustrate the detailed procedures as follows.

The estimation procedure of the PDI involves developing a classification model using $I\{Y > S\}$ as the binary label against the dose and the covariates. Subsequently, a grid search is done for each patient in order to identify a range of doses where the predicted probability of having $Y > S$ is greater than the confidence probability α, i.e., $\hat{P}(Y > S \,|\, a^*, x) \geq \alpha$. To enable the grid search, the classification model should be able to deliver a fitted probability for each patient on each dose. For instance, if the classification model is a logistic regression model, the fitted probability can be easily calculated through the link function. Another example is the Random Forest (Breiman, 2001), where the probability is the proportion of positive at the majority voting step. For the Support vector machine (SVM) (Vapnik, 2013), it more involved as a refitting step, suggested in Platt (1999), must be taken so that the SVM results can be interpreted in probabilistic language.

Similarly, EDI is estimated through regression methods. The majority of the regression methods target the conditional mean, $E[Y|a, x]$ in our case, and hence they can be utilized as the first step outcome model. For each individual, the grid search is done to find a range of doses where the predicted outcome given the dose and the covariates is above a threshold

S. In the case of the Ordinary Least Square (OLS), this implies finding a range of a such that $\hat{E}[Y|a,x] = X\hat{\beta} > S$. Other nonparametric methods or machine learning techniques, including the Random Forest regression (Breiman, 2001), can also serve the same purpose. Support vector regression (Drucker et al., 1997) essentially targets the conditional median, but it should be able to approximate the conditional mean in most of the cases.

A problem naturally arises since the predicted probabilities might be completely above or below the confidence probability α for classification methods within a reasonable range of doses. Similarly, the predicted conditional expectations can be entirely above or below the threshold S. Strategies should be prepared in order to handle such cases. For instance, if the dose has a natural lower bound a_L and a natural upper bound a_U, the PDI for a patient is $[a_L, a_U]$ when all predicted probabilities are above α, i.e., $\hat{P}(Y > S \,|\, a, x) > \alpha \;\; \forall a \in [a_L, a_U]$ and the EDI is also $[a_L, a_U]$ when all predicted conditional expectations are above S, i.e., $\hat{E}[Y|a,x] > \alpha \;\; \forall a \in [a_L, a_U]$. Conversely, when $\hat{P}(Y > S \,|\, a, x) < \alpha$ or $\hat{E}[Y|a,x] < S$ for all $a \in [a_L, a_U]$, we have the PDI or the EDI being the empty set. This procedure does not preclude alternative strategies. For example, a pre-specified conservative dose interval could be used when the grid search suggests that all doses are desirable or all doses are undesirable.

The following of this dissertation is organized as follows: In Chapter 2, we will introduce our new research method based on the OWL, where the **Quadrant Loss**, the basis of our method, is derived in Section 2.3. In Chapter 3, we will clarify the necessary assumptions in addition to several causal issues associated with our method. In Chapter 4, we discuss the estimation of the IDI based on non-convex risk functions using relaxations. In Section 4.2, the DC algorithm will be applied to solve the non-convex objective functions. In Chapter 5, theorems and corollaries will explain the reason why the proposed method is Fisher consistent, which shall be followed by the analysis of the convergence rate of the estimators. Simulation studies for different scenarios can be found in Chapter 6. Finally, in Chapter 7 we will provide an estimation of the personalized HbA1c control interval using a large EHR dataset of type II diabetes patients.

Chapter 2

Method

In Section 1.4, we have illustrated the estimation of IDI using classification and regression methods based on traditional outcome modeling. However, the success of such indirect methods necessitates the quality of the outcome model. In the cases where the outcome model can be correctly specified, the grid search might be able to produce IDIs which are beneficial to the patients. When the outcome model is hard to specify but the sample size is large, one can still rely on nonparametric methods in order to attain the predicted probabilities or conditional expectations. But when the outcome model is believed to be complicated or the sample size is small, using the indirect method could lead to biased or inefficient estimates of IDI. In this Chapter, we will provide an alternative approach based on the OWL framework without modeling the outcome.

2.1 From Binary OWL to Individual Dose Interval

Zhao et al. (2012) established the original OWL framework for personalized treatment assignment without modeling the outcome. However, their framework has to be extended as now we need to assign a personalized interval of doses. In the case of the IDI, the continuous dose A takes values in a range $[a_L, a_U]$ instead of $\{0, 1\}$, as is in the ITR case. Let $\mathcal{A}_1(x)$ denote the desirable subset of $\mathcal{A}$ for $x \in \mathcal{X}$, while $\mathcal{A}_0(x)$ denotes the undesirable

subset, such that $\mathcal{A}(x) = \mathcal{A}_0(x) \cup \mathcal{A}_1(x)$ and $\mathcal{A}_0(x) \cap \mathcal{A}_1(x) = \emptyset$. Unlike the objective function 1.1 where the ITR $f(x)$ selects from $\{0, 1\}$, the IDI selects $\mathcal{A}_1(x) \subset \mathcal{A}(x)$.

Also, in equation (1.1), the weights in the classification problem is proportional to Y, and therefore, assuming $A = 1$ is the better treatment for x, the regret for not being able to select the optimal treatment is $\big|Y(1, x) - Y(0, x)\big|$. In the IDI case, the regret for a patient x who followed the IDI but had outcome below S should be $S - Y$. On the contrary, if the patient failed to follow the IDI but received outcome above S, there would be a regret of $Y - S$. Hence, we can construct the objective function to be minimized similarly to (1.1) as below.

$$E_{(Y,A,X)}\left[\frac{1}{P(A \,|\, X)}\Big(\big(S - Y\big)_+ 1\{A \in \mathcal{A}_1(X)\} + \big(Y - S\big)_+ 1\{A \notin \mathcal{A}_1(X)\}\Big)\right]$$

where the inverse weighting probability $P(A \,|\, X)$ is the conditional density function of A given X which is assumed to be known in experimental settings and has to be estimated in the observational settings. In Section 3.2, we will discuss the estimation of $P(A \,|\, X)$. The above minimization problem is equivalent to maximize

$$\begin{aligned}
&E\left[\frac{1}{P(A \,|\, X)}\Big(\big(Y - S\big)_+ 1\{A \in \mathcal{A}_1(X)\} + \big(S - Y\big)_+ 1\{A \notin \mathcal{A}_1(X)\}\Big)\right] \\
&\quad = E_X\Bigg[\int \Bigg(E\Big[\big(Y - S\big)_+ | A = a, X\Big] 1\{a \in \mathcal{A}_1(X)\} \\
&\qquad\qquad + E\Big[\big(S - Y\big)_+ | A = a, X\Big] 1\{a \notin \mathcal{A}_1(X)\}\Bigg) da\Bigg] \\
&\quad = E_X\Bigg[\int \Bigg(1\{a \in \mathcal{A}_1(X)\}\Big(E\Big[\big(Y - S\big)_+ | A = a, X\Big] - E\Big[\big(S - Y\big)_+ | A = a, X\Big]\Big) \\
&\qquad\qquad + E\Big[\big(S - Y\big)_+ | A = a, X\Big]\Bigg) da\Bigg]
\end{aligned}$$

for a fixed $X = x$, this boils down to maximize

$$\begin{aligned}
\int \Bigg(&1\{a \in \mathcal{A}_1(x)\}\Big(E\Big[\big(Y - S\big)_+ | A = a, X = x\Big] - E\Big[\big(S - Y\big)_+ | A = a, X = x\Big]\Big) \\
&+ E\Big[\big(S - Y\big)_+ | A = a, X = x\Big]\Bigg) da
\end{aligned}$$

Therefore, the optimal division of $\mathcal{A}$ into $\mathcal{A}_1(x)$ and $\mathcal{A}_0(x)$ should satisfy

$$\begin{cases} a \in \mathcal{A}_1(x), & if \quad E\Big[Y|A=a, X=x\Big] > S \\ a \in \mathcal{A}_0(x), & if \quad E\Big[Y|A=a, X=x\Big] < S \end{cases} \tag{2.1}$$

In Section 2.3, we will show that the above objective function leads to the EDI defined in Section 1.4 with certain structural assumptions which will be covered in Section 3.1.

2.2 From Classification to Individual Dose Interval

In Section 1.4, the two-step procedure based on the classification methods and grid search is introduced to estimate the PDI with threshold S and the confidence probability α. A PDI should be both sufficient and necessary. A patient is supposed to have an outcome greater than the pre-specified constant S if the dose is inside the PDI. On the other hand, the patient's outcome is supposed to be less than the constant, if the dose is outside the IDI. Due to the randomness that cannot be ruled out even in the most rigorous experiment, violations will exist in a proportion of the observations.

A new patient $x \in \mathcal{X}$ requiring a PDI recommendation presents the risk of two possible errors. First, when the patient receives a dose inside the interval, i.e., $A \in \mathcal{A}_1(x)$, but has an outcome below the threshold, i.e., $Y < S$, it presents a false-positive error with weight α. Second, when the patient receives a dose outside the interval, i.e., $A \in \mathcal{A}_0(x)$, but has an outcome above the threshold, i.e., $Y > S$, it presents a false-negative error with weight $1-\alpha$. The weights α and $1-\alpha$ balance the magnitude of the two errors and eventually determine the PDI. Essentially, the two-step procedure minimizes the below weighted misclassification by grid searching $\mathcal{A}_1(x)$ and $\mathcal{A}_0(x)$ for each $x \in \mathcal{X}$.

$$E\Big[\alpha I(Y < S^*) \times I\big(A \in \mathcal{A}_1(X)\big) + (1-\alpha) I(Y > S^*) \times I\big(A \notin \mathcal{A}_0(X)\big)\Big] \tag{2.2}$$

The mathematical reasoning is implied in the proof of theorem 5.3.1. However, this two-step

procedure relies on the validity of the classification model. In practice, it is hard to guarantee the quality of the model as the data generating mechanism of the outcome variable is often complicated. Sometimes, it is desirable to circulate the modeling of outcome and instead directly minimize (2.2) in order to estimate $\mathcal{A}_1(x)$ and $\mathcal{A}_0(x)$. We will show in Section 2.3 that this idea results in the OWL-based PDI method.

For the classification based methods, the interactions between A and X are captured in the outcome model. The direct estimation of $\mathcal{A}_1(x)$ and $\mathcal{A}_0(x)$ from (2.2) requires a pseudo-population weighted by the inverse density $1/P(A \mid X)$ where the interactions between A and X will not bias the estimation under the ignorability assumption (see Section 3.1). In theorem 5.3.1 and corollary 5.3.1, the necessity of such adjustment is demonstrated. In addition, postulating structural assumptions of $Y(a, x)$ can facilitate the direct estimation of IDI. For example, if the potential outcome $Y(a, x)$ is a monotonically increasing function of a given x, then only the lower bound is needed to separate the desirable doses from the undesirable doses, i.e., $\mathcal{A}_1(x) = [f_L(X), a_U]$ and $\mathcal{A}_0(x) = [a_L, f_L(X)]$. Therefore, the updated risk function becomes

$$E\Bigg[\frac{1}{P(A \mid X)}\Big(\alpha I(Y < S^*) \times I\big(A \in [f_L(X), a_U]\big) + (1-\alpha) I(Y > S^*) \times I\big(A \in [a_L, f_L(X)]\big)\Big)\Bigg] \tag{2.3}$$

Similarly, if $Y(a, x)$ is a unimodal function of a and monotonic on each side of the mode, the updated risk function becomes

$$E\Bigg[\frac{1}{P(A \mid X)}\Big(\alpha I(Y < S^*) \times I\big(A \in \big[f_L(X), f_U(X)\big]\big) + (1-\alpha) I(Y > S^*) \times I\big(A \notin \big[f_L(X), f_U(X)\big]\big)\Big)\Bigg] \tag{2.4}$$

Again, the inverse density weighting $1/P(A \mid X)$ is the conditional density function of A given X which is assumed to be known in the experimental settings and has to be estimated in the observational settings. In Section 3.1, weaker structural assumptions are documented.

In Section 3.2, we will discuss the estimation of $P(A \mid X)$. In the next section, we will specifically show that both (2.1) and (2.4) are special cases of a general framework.

2.3 Quadrant Loss

The approaches derived in the previous two sections are special cases of applying the Quadrant Loss. In this section, we will demonstrate the intuition of the generalized framework of the Quadrant Loss. We start with a hypothetical case where the potential outcome $Y(a, x)$ monotonically increases with the treatment a for a specific patient with features $x \in \mathcal{X}$. For simplicity, we assume $Y(a, x) = m(a, x) + \epsilon$ and $\epsilon \sim \mathcal{N}(0, 1)$ where the subpopulation of $X = x$ is shown in Figure 2.1. Then we attempt to provide an estimate $\hat{f}_L(x)$ of the optimal lower bound of the IDI for x given a threshold S and a confidence probability $\alpha = 0.5$.

Now we consider two options of $\hat{f}_L(x)$, a_0 and a_1, where a_0 is the optimal lower bound for $x \in \mathcal{X}$ and a_1 a sub-optimal one larger than a_0. We will demonstrate why $\hat{f}_L(x)$ reaches a_0 instead of a_1 by minimizing the Quadrant Loss. The optimal lower bound a_0 takes the value where $m(a, x)$ intersects S, as is indicated by the blue dotted vertical line in Figure 2.1. In fact, the threshold S and $\hat{f}_L(x)$ would form a set of coordinates, where the Quadrant Loss only penalize the second (upper left) and the fourth (lower right) quadrants. Only the red dots in the second quadrant and the green dots in the fourth quadrant would contribute to the loss. So by letting $\hat{f}_L(x) = a_0$, the Quadrant Loss is minimized as the sum of the numbers of the red and green dots is the smallest. Alternatively, if we let $\hat{f}_L(x)$ equal to the other candidate lower bound a_1 which is slightly larger than a_0, then there will be more dots in the second Quadrant and fewer in the fourth Quadrant, as is illustrated in Fig 2.2. Since the increase of the dots in the second Quadrant (as is shown by the orange dots) would always be larger than the decrease of dots in the fourth Quadrant (as is shown by the blue dots) (see Fig 2.2), to move $\hat{f}_L(x)$ from a_0 to a_1 would increase Loss. Similarly, there is no gain moving $\hat{f}_L(x)$ from a_0 to the left. So a_0 is the optimal $\hat{f}_L(x)$ which minimizes

the Quadrant Loss. We will provide a formal proof in section 5.2.

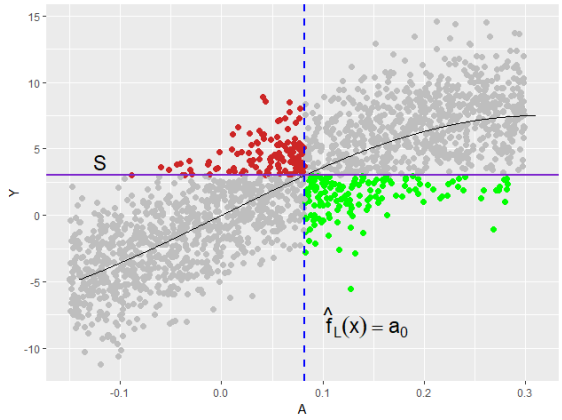

Figure 2.1: The Quadrant Loss for the Estimated Optimal Lower Bound

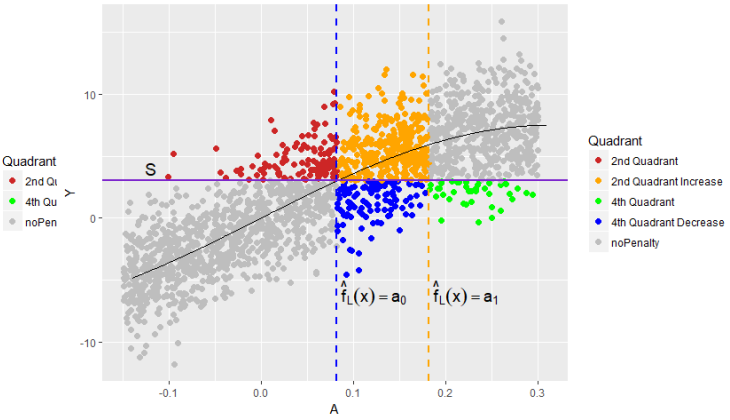

Figure 2.2: The Change of Quadrant Loss for an Alternative Lower Bound

Another implication of Figure 2.1 is that the weights assigned to the losses of the second and the fourth quadrant do not have to be the same. Naturally, the patients might care more if they followed the IDI suggestion and nonetheless had bad outcomes compared to if they failed to follow the recommendation and nonetheless had good outcomes. In that case, we are essentially looking for a conservative dose interval by assigning more weights to the loss of $Y < S$ to encourage a narrower interval. We will demonstrate in the theorem 5.3.1 that this corresponds to choosing confidence probability $\alpha \neq \frac{1}{2}$ in the PDI case. Moreover, if we let the ratio of the losses incurred by a single observation in the fourth quadrant and in the second quadrant to be $\frac{\alpha}{1-\alpha}$, the confidence probability α in PDI is guaranteed. In another word, the patients receiving doses inside the interval are guaranteed to have the response $Y > S$ with a probability of at least α. As for EDI, it requires $\alpha = \frac{1}{2}$ and it is guaranteed the patients receiving doses inside the interval to have expected response greater than S. The detailed proofs are established in Section 5.2.

Instead of identifying the IDI for a single type of patients $X = x$, we allow information to be shared between patients and the IDIR can, therefore, be learned as a function of patients' features x. Here, we denote $f_L(x)$ as the IDI lower bound function and $f_U(x)$ as

the IDI upper bound function. Formally, the risk functions that determines the one-sided and two-sided IDI are shown as below.

The Quadrant Risk for the lower bound function in the one-sided IDI is as follows.

$$R_\Phi(f_L) = E_{Y,A,X}\left[\frac{1}{P(A\,|\,X)}\Big((1-\alpha)\Phi(Y-S)I\big(A<f_L(X)\big)+\alpha\Phi(S-Y)I\big(A>f_L(X)\big)\Big)\right] \tag{2.5}$$

The first term of the risk function above comes from the second quadrant while the second term comes from the fourth quadrant. Correspondingly, the Quadrant Risk for the IDI upper bound function is as follows.

$$R_\Phi(f_U) = E_{Y,A,X}\left[\frac{1}{P(A\,|\,X)}\Big((1-\alpha)\Phi(Y-S)I\big(A>f_U(X)\big)+\alpha\Phi(S-Y)I\big(A<f_U(X)\big)\Big)\right] \tag{2.6}$$

While the Quadrant Risk for the two-sided IDI is as follows.

$$R_\Phi(f_L, f_U) = \\ E_{Y,A,X}\left[\frac{1}{P(A\,|\,X)}\Big((1-\alpha)\Phi(Y-S)I(A\notin[f_L(X),f_U(X)])+\alpha\Phi(S-Y)I(A\in[f_L(X),f_U(X)]\Big)\right] \tag{2.7}$$

The weight function $\Phi(\cdot)$ shows the adjustment of loss based on a misclassified case. There are two choices, the indicator loss $I(\cdot>0)$ is engaged to learn the PDIR, and the hinge loss $(\cdot)_+$ is used in learning the EDIR. The two corresponding risk functions are denoted as $R_I(f_L)$ and $R_h(f_L)$ for one-sided PIDR and EDIR, as well as $R_I(f_L, f_U)$ and $R_h(f_L, f_U)$ for two-sided PDIR and EDIR. The notations of losses and risks, among relaxed losses which we will introduce in Chapter 4, are summarized comprehensively in Table 4.1. The inverse density weighting $1/P(A\,|\,X)$ is the conditional density function of A given X which is assumed to be known in experimental settings and has been estimated in observational settings. In Section 3.2, we will discuss the estimation of $P(A\,|\,X)$.

In order for the IDIs to be identifiable, certain assumptions are needed. For one-sided IDIs, we assume that $Y(a,x)$ is a Partial Monotonic function of a, and for two-sided IDIs,

we assume that $Y(a, x)$ is a Partial Unimodal function of a. These assumptions are slightly weaker than the monotonicity assumption and unimodality assumption. We will discuss the details of these assumptions in Section 3.1.

2.4 Personalization of Threshold and Confidence Probability

So far, the confidence probability α and the threshold S have been regarded as constants for both indirect methods in Section 1.4 and our OWL-based method in Section 2.3. This could potentially limit the applications of these IDI methods.

In fact, there is a large variation in patients' baseline conditions. Patients may have very different conditions prior to receiving the doses. Admittedly, in cases where the patients' response is mainly decided by the treatment dose, this will not cause any problems. However, in reality, patients' outcome may be affected by these baseline factors extensively, like the medication history, the disease progression, and demographics. These factors may overwhelmingly contribute to the patients' response as prognostic effects. Consider a case where patients' outcome is dominated by the prognostic effects, wherein, the identified confidence probability α and the threshold S would only be meaningful for a small proportion of patients. For the rest cases dominated by prognostic effects, those who have good baseline conditions may nonetheless show desirable outcomes given any reasonable doses, and those who have bad baseline conditions may nonetheless show undesirable outcomes given any reasonable dosing. This scenario may invalidate the applicability of a constant threshold of S. Therefore, we shall consider the threshold as a function of patients' baseline differences, i.e., $S = S(x)$. The same restriction may apply to the confidence probability α; while there may naturally be a higher probability for some patients to have desirable outcomes, the opposite might be true for other patients, in spite of taking optimal doses. Therefore, it is necessary to personalize the confidence probability as a function of patients' characteristics as well, i.e., $\alpha = \alpha(x)$. Grantedly, how to personalize $S(x)$ and $\alpha(x)$ is an open question which should be decided on a case-by-case basis.

As long as both of the two functions are pre-determined, the estimation procedures as in Section 4 can still be used for such personalization of patients' baseline conditions, given that the interpretations and the assumptions modified accordingly. Now the PDI is an IDI where the patient x will have outcome $Y > S(x)$ with probability at least $\alpha(x)$. We suggest that only one of $S(x)$ and $\alpha(x)$ is personalized for the PDI to avoid the difficulty of interpretation. As for the EDI, α is required to equal $\frac{1}{2}$ and the EDI provides a range of doses where the expected outcome of patient x is above $S(x)$. Further discussion of the assumptions can be found in theorems from 5.2.1 to corollary 5.2.2.

It is straightforward to use the indirect methods to solve IDI. The first step will not change as the prognostic effects can be captured by either classification methods or regression methods. The only difference is in the second step. If using grid search methods to find the IDIs for a particular patient x, we may simply replace the constant S and α by the values of the functions $S(x)$ and $\alpha(x)$.

Correspondingly, the new risk functions for the one-sided IDIs are as following.

$$R_\Phi(f_L) = E_{Y,A,X}\Big[\frac{1}{P(A\,|\,X)}\Big(\big(1-\alpha(X)\big)\Phi\big(Y-S(X)\big)I\big(A<f_L(X)\big)] \\ +\alpha(X)\Phi\big(S(X)-Y\big)I\big(A>f_L(X)\big)\Big)\Big] \tag{2.8}$$

and

$$R_\Phi(f_U) = E_{Y,A,X}\Big[\frac{1}{P(A\,|\,X)}\Big(\big(1-\alpha(X)\big)\Phi\big(Y-S(X)\big)I\big(A>f_U(X)\big)] \\ +\alpha(X)\Phi\big(S(X)-Y\big)I\big(A<f_U(X)\big)\Big)\Big] \tag{2.9}$$

The new risk function for the two-sided IDI is as following.

$$R_\Phi(f_L, f_U) =$$
$$E_{Y,A,X}\Big[\frac{1}{P(A \mid X)}\Big(\big(1-\alpha(X)\big)\Phi\big(Y - S(X)\big)I(A \notin \Big[f_L(X), f_U(X)\Big]\big) + \alpha(X)\Phi\big(S(X) - Y\big)I(A \in \Big[f_L(X), f_U(X)\Big]\Big)\Big] \tag{2.10}$$

We use the example below to illustrate the interpretation of PDI with personalized thresholds and confidence probabilities. $\text{PDI}\big(\alpha = \hat{P}(Y > 10 \mid x), S = 10, x\big)$ represents a dose interval where we have $Y(a, x) > 10$ with a probability of at least $\hat{P}(Y > 10 \mid x)$ for each patient x, where the estimated probability can be obtained from an independent model and is determined before estimating PDI. Similarly, $\text{PDI}\big(\alpha = 0.95, S = \hat{E}\left[Y \mid x\right], x\big)$ represents a dose interval where $Y(a, x)$ is guaranteed to be greater than its conditional mean with probability 95%. Besides, an $\text{EDI}\big(\alpha = \frac{1}{2}, S = \hat{E}\left[Y|x^*\right], x\big)$ will imply a dose interval where $E[Y(a, x^*)] > \hat{E}\left[Y \mid x\right]$. This suggests that the proper choice of a treatment dose is expected to results in a better outcome compared to the expected outcome assuming the treatment dose is chosen using the same strategy as in the dataset where the estimate $\hat{E}\left[Y \mid x\right]$ is obtained.

When dealing with a binary outcome $Y \in \{0, 1\}$, two special properties of our proposed methods should be noted. First, the same estimates of PDI can be yielded with any $S(x) \in (0, 1)$. Therefore, in this case, personalization should only be applied to $\alpha(x)$ but not $S(x)$. Second, the PDI with $S \in (0, 1)$ and $\alpha(x) = P(Y = 1 \mid x)$ is equivalent to the EDI with $\alpha = \frac{1}{2}$ and $S(x) = P(Y = 1 \mid x)$. This is due to the fact that such personalization implies the same weights in (2.8) and (2.10).

Chapter 3

A Causal Framework of Dose Interval Recommendation

In Chapter 2, we proposed our approach based on the OWL framework for determining IDIs. Before diving into the estimation of IDIs, we will discuss the identification problems in order for the estimated IDIs to have the causal implications. In Section 3.1, we discuss the main causal assumptions that we impose. In Section 3.2, we will discuss the estimation of the inverse density weights. In Section 3.3, we define the average treatment effect of IDI recommendation for evaluation purpose.

3.1 Assumptions

In the previous Chapters, we have defined and discussed the IDIs. Specifically, the definitions of the PDI and the EDI are stated using the potential outcome language introduced in Rubin (1974). In order for the PDI and the EDI to be identifiable from observational data, the ignorability assumption must be imposed for the linkage between the counterfactual world and real world.

In our case, the ignorability assumption implies that the received dose is independent of the potential outcomes given the covariates, i.e., $A \perp\!\!\!\perp Y(a,x)|x$ where $a \in [a_L, a_U]$. This

assumption is essentially not testable but might be guaranteed by a careful incorporation of baseline covariates. Therefore, in Chapter 7, we attempt to collect as large a set of relevant upstream variables prior to the dose as possible, so that even when this assumption is not perfectly satisfied, the impact of its violation may be alleviated.

Besides the ignorability assumption, for the IDIs to be identifiable, we need additional assumptions, which are referred to as the structural assumptions. For the one-sided interval and the two-sided interval, there are four different structural assumptions listed below for both PDI and EDI.

The partial monotonicity assumption for one-sided PDI (3.1)

- $P\Big(Y(a_L, x) > S(x)\Big) \leq \alpha(x)$ and $P\Big(Y(a_U, x) > S(x)\Big) \geq \alpha(x)$
- $P(Y(a, x) > S(x))$ across $\alpha(x)$ at most once as a goes from a_L to a_U

The partial unimodality assumption for two-sided PDI (3.2)

- $P\Big(Y(a_L, x) > S(x)\Big) \leq \alpha(x)$ and $P\Big(Y(a_U, x) > S(x)\Big) \leq \alpha(x)$
- $\exists a^* \in [a_L, a_U],\ P(Y(a^*, x) > S(x)) \geq \alpha(x)$
- $P(Y(a, x) > S(x))$ cross $\alpha(x)$ at most twice as a goes from a_L to a_U

The partial monotonicity assumption for one-sided EDI (3.3)

- $E\left[Y(a_L, x)\right] \leq S(x)$ and $E\left[Y(a_U, x)\right] \geq S(x)$
- $E(Y(a, x))$ across $S(x)$ at most once as a goes from a_L to a_U

The partial unimodality assumption for two-sided EDI (3.4)

- $E\left[Y(a_L, x)\right] \leq S(x)$ and $E\left[Y(a_U, x)\right] \leq S(x)$
- $\exists a^* \in [a_L, a_U],\ E\left[Y(a^*, x)\right] \geq S(x)$
- $E(Y(a, x))$ cross $S(x)$ at most twice as a goes from a_L to a_U

In Section 5.2, we will demonstrate that the mechanisms via which these assumptions enables the Fisher consistency. For one-sided intervals, we only present the assumptions for the lower bounds and the assumptions for the upper bounds can be adapted accordingly.

3.2 Propensity Score of Dose Assignment

In the previous sections, we refer to the inverse density weights $1/P(a \,|\, x)$ as the known function of patients' features x. This might be the case in clinical trials where the doses or even the features are determined or randomized by design. However, there is only limited number of clinical trials for which dosing is the target. What is more, there may not always exist enough treatment levels so that the treatment could be viewed as a continuous variable. The inadequacy of randomized clinical trials forces us to consider observational studies where the conditional density $P(a \,|\, x)$ has to be estimated.

Unlike the binary ITRs where the probability of receiving one of the treatments is used as the propensity score (Zhao et al., 2012, Zhou et al., 2017), the estimation of IDIR requires an inverse density weights $1/P(a \,|\, x)$ where $P(a \,|\, x)$ is the conditional density of the dose given the covariates (Chen et al., 2016). Multiple approaches are available for estimating the conditional density. First, a regression model, either parametric or nonparametric, can be constructed to model the conditional expectation of the dose given covariates. And the residuals could be used to estimate the conditional density either from a parametric distribution, i.e., the Gaussian distribution or from Kernel density estimation. The same technique has been used in Chen et al. (2016). This approach offers the best performance when the distributional assumptions are reasonable or the distribution of error does not show much variation across the range of doses. Of course, the quality of the conditional

mean model has to be guaranteed. Second, the range of doses can be divided into multiple intervals where the probability of falling into one interval is modeled using a multinomial regression. This approach introduces much more parameters for the benefit of increasing flexibility. The third approach is similar to the second approach, but the probability of falling into one interval is instead modeled using a cumulative ordinal logistic regression (Harrell, 2001). This approach has fewer parameters and thus is less flexible, however, it is indeed more robust compared to the second approach. In the real data analysis of HbA1c measure in Chapter 7, we applied the third approach, yet we believe each of these three approaches has its own suitable scenarios. Another potential candidate is Covariates-Balancing Propensity Score for continuous variable (Fong et al., 2018), which forces the balancing of covariates across the range of continuous treatment. From the experiences with the dataset in Chapter 7, this procedure leads to less efficient estimates, which can probably be attributed to the reduced emphasis on conditional density. Therefore, we do not recommend using this approach for general OWL-based IDI although we do not preclude the possibility that it may have good performance in certain cases.

3.3 Average Causal Effect of A Dose Interval Recommendation

The causal inference for IDIR has to be established in order to answer such a question: how large is the effect of following the IDI recommendation versus not following? As illustrated in the Chapter 2.4, IDI may be estimated with both the OWL-based methods and the indirect methods. However, the evaluation of these IDI estimates is not as straightforward. For an IDIR $\big[f_L(x), f_U(x)\big]$, we define T_i as the indicator of whether the patient i received a dose are inside the recommended IDI, i.e., $T_i = 1\big\{A_i \in [f_L(X_i), f_U(X_i)]\big\}$. Correspondingly, patients with $T_i = 1$ constitute the taker group and patients with $T_i = 0$ the non-taker group. Ideally, we shall compare the distribution of outcome between the takers and non-takers. A large gap in the proportions of the good outcome, $Y_i > S(X_i)$, indicates the

potential benefit of the IDIR. However, this procedure can introduce large bias because of the following facts. First, the distributions of the estimated IDI lower/upper bounds vary for different methods. When $S(x)$ and $\alpha(x)$ are not perfectly personalized, the indirect methods may distinct observations with $Y_i > S(X_i)$ from observations with $Y_i < S(X_i)$ simply by modeling the prognostic effect. Consequently, the baseline conditions of takers might be much better than those of non-takers. Therefore, a direct comparison might overestimate the effectiveness of the recommendation for the indirect methods.

To evaluate different IDI methods, we formally define the average causal effect of IDI recommendation as below. To eliminate possible confounding factors, it is necessary to make the taker group and the non-taker group comparable in terms of covariates distribution. Only in that case, the average causal effect of IDIR can be identified under th ignorability assumption in Section 3.1. Therefore, we develop a propensity score model based on the CBPS approach introduced in Imai and Ratkovic (2014), using T_i as the outcome against all the potential confounders X_i. The model delivers two estimated probability $\hat{P}(T_i = 1|X_i)$ and $\hat{P}(T_i = 0|X_i)$. The average treatment effect for the PDI is defined as the gap of proportions of $Y_i > S(X_i)$ between weighted taker group and weighted non-taker group. For the EDI, the gap of $Y_i - S(X_i)$ is used.

$$ATE_{PDI} = \sum_{i=1}^{n} \frac{1\{Y_i > S(X_i)\}T_i}{P(T_i|X_i)} - \sum_{i=1}^{n} \frac{1\{Y_i > S(X_i)\}(1-T_i)}{P(T_i|X_i)}$$

$$ATE_{EDI} = \sum_{i=1}^{n} \frac{\big(Y_i - S(X_i)\big)T_i}{P(T_i|X_i)} - \sum_{i=1}^{n} \frac{\big(Y_i - S(X_i)\big)(1-T_i)}{P(T_i|X_i)}$$

The CBPS method is believed to alleviate the imbalance more efficiently, because it incorporates the goal of balancing of covariates into the objective function. After imposing the weights, the weighted population approximates the hypothetical randomized experiment where patients are randomly assigned to the taker group and the non-taker group.

Chapter 4

Estimation

In this Chapter, we introduce the estimation methods for solving objective function (2.8) and (2.10). In section 4.1, we relax the indicator function to be a non-convex but continuous surrogate loss which is then solved by DC-algorithm. The technical details of this approach are discussed in details in section 4.2. This framework can incorporate both linear estimation and Kernel estimation, whose results are compared in Chapter 6 and Chapter 7.

4.1 Non-Convex Relaxation

With the indicator function in (2.9) and (2.10), it is difficult to optimize the empirical risk function. Convex relaxations are widely used in Statistics and Machine Learning where the indicator function is often relaxed by hinge loss, like in various Support Vector Machine studies. However, it can be shown that the indicator functions cannot be replaced by a convex loss in the optimal dose problem (Chen et al., 2016). This is also true for the IDI problem as the convex relaxation will incur extremely large losses on observations far from the optimal bounds which bias the estimation. Instead, we propose a truncated hinge loss for one-sided IDI and an integrated truncated hinge loss for the two-sided IDI. These surrogate losses are non-convex but continuous and can be viewed as natural extensions of the non-convex loss from Chen et al. (2016). Since non-convex losses are incorporated, the

DC-algorithm could be applied to solve the resulting non-convex objective functions, as is shown in the next section.

The risk function of one-sided IDI after relaxation is as follows.

$$R_{\Phi,\epsilon}(f_L) = E_{Y,A,X}\Big[\frac{1}{P(A\,|\,X)}\Big(\big(1-\alpha(X)\big)\Phi\big(Y-S(X)\big))\Psi_\epsilon(f_L(X),A) \tag{4.1}$$

$$+\,\alpha(X)\Phi\big(S(X)-Y\big))\Psi_\epsilon(A,f_L(X))\Big)\Big] \tag{4.2}$$

The risk function of two-sided IDI after relaxation is as follows.

$$R_{\Phi,\epsilon}(f_L,f_U) = \tag{4.3}$$

$$E_{Y,A,X}\Big[\frac{1}{P(A\,|\,X)}\Big(\big(1-\alpha(X)\big)\Phi\big(Y-S(X)\big)\Psi_\epsilon^{out}\Big(f_L(X),A,f_U(X)\Big) \tag{4.4}$$

$$+\,\alpha(X)\Phi\big(S(X)-Y\big)\Psi_\epsilon^{in}\Big(f_L(X),A,f_U(X)\Big)\Big)\Big] \tag{4.5}$$

where

$$\Psi_\epsilon(a,b) = min\{\frac{(a-b)_+}{\epsilon},1\}$$

$$\Psi_\epsilon^{in}\Big(a,b,c\Big) = \begin{cases} b<a<c, & 0 \\ a<b<(a+\epsilon)<c, & (b-a)/\epsilon \\ (a+\epsilon)<b<(c-\epsilon), & 1 \\ a<(c-\epsilon)<b<c, & (c-b)/\epsilon \\ a<c<b, & 0 \end{cases}$$

$$\Psi_\epsilon^{out}\Big(a,b,c\Big) = 1-\Psi_\epsilon^{in}\Big(a,b,c\Big)$$

The plots of these three surrogate losses can be found in Fig 4.1, Figure 4.2 and Figure

4.3. To summarize, we listed the combinations of the loss, the risk and the corresponding interval type that they imply in Table 4.1.

Table 4.1: Notations

Weight	Loss	Risk	Relaxed Risk	Methods
$\Phi = I(\cdot > 0)$	Ψ_ϵ	$R_I(f_L)$	$R_{I,\epsilon}(f_L)$	one-sided PDI
$\Phi = I(\cdot > 0)$	$\Psi_\epsilon^{in},\ \Psi_\epsilon^{out}$	$R_I(f_L, f_U)$	$R_{I,\epsilon}(f_L, f_U)$	two-sided PDI
$\Phi = (\cdot)_+$	Ψ_ϵ	$R_h(f_L)$	$R_{h,\epsilon}(f_L)$	one-sided EDI
$\Phi = (\cdot)_+$	$\Psi_\epsilon^{in},\ \Psi_\epsilon^{out}$	$R_h(f_L, f_U)$	$R_{h,\epsilon}(f_L, f_U)$	two-sided EDI

4.2 DC Algorithm for One-sided Dose Interval

In the previous section, the risk function is relaxed by replacing the indicator functions with surrogate losses. In this section, we apply DC algorithm which estimates the IDIs by minimizing the empirical risk. To start with one-sided IDI, let $\mathcal{H}$ being a Reproducing Kernel Hilbert Space (RKHS). The optimization procedure for one-sided IDI attempts to solve the following problem:

$$\hat{f}_L = \arg\min_{f\in\mathcal{H}} \sum_{i=1}^{n} \Big[\frac{1}{P(A_i|X_i)}\Big(\big(1-\alpha(X)\big)\Phi(Y_i - S(X))\Psi(f(X_i), A_i) \tag{4.6}$$

$$+ \alpha(X)\Phi(S(X) - Y_i)\Psi(A_i, f(X_i))\Big)\Big] + \frac{\lambda_n}{2}||f_L||_2^2 \tag{4.7}$$

The representer theorem is applicable in this case, but since the convexity no longer holds, it is difficult to solve the optimization problem numerically. Chen et al. (2016) applied the DC algorithm to solve their optimization problem. It turns out that the DC algorithm are also applicable in our case, since the truncated hinge loss $\Psi(a,b) = min\{\frac{(a-b)_+}{\epsilon}, 1\}$ can also be decomposed as the difference of two convex loss function, i.e., $\Psi_1(a,b) = (\frac{a-b}{\epsilon})_+$, $\Psi_2(a,b) = (\frac{a-b}{\epsilon} - 1)_+$ and $\Psi(a,b) = \Psi_1(a,b) - \Psi_2(a,b)$.

According to Le Thi Hoai and Tao (1997), the DC algorithm can be used to solve the non-convex minimization problem by iteratively solving a series of convex minimization problem, where the non-convex objective function are transformed to the difference of two

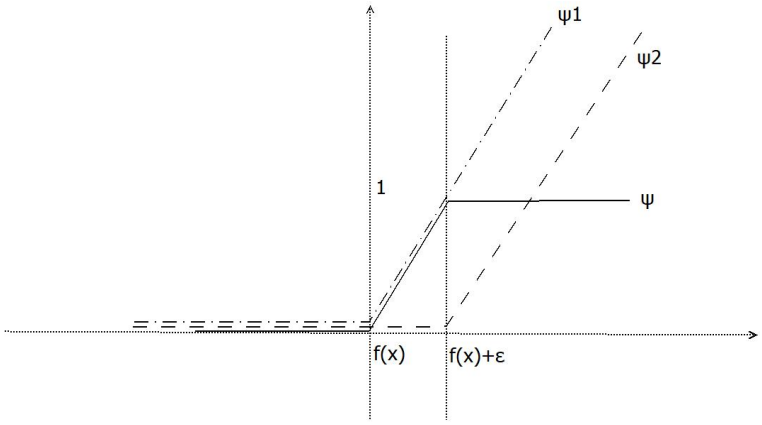

Figure 4.1: Loss Functions Ψ_ϵ,Ψ_1 and Ψ_2

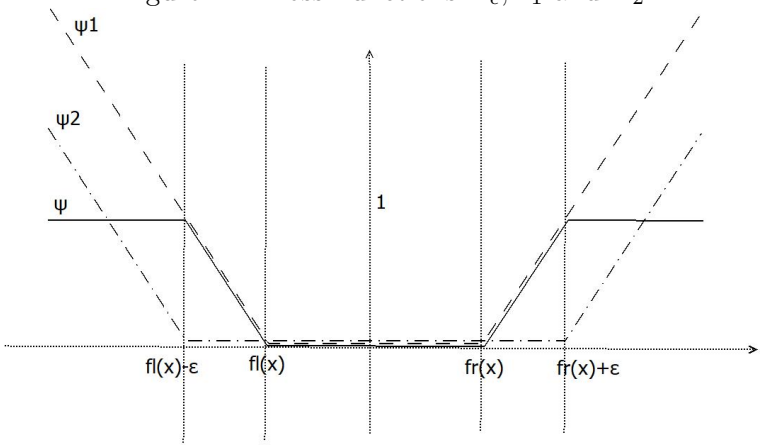

Figure 4.2: Loss Functions Ψ_ϵ^{out},Ψ_1^{out} and Ψ_2^{out}

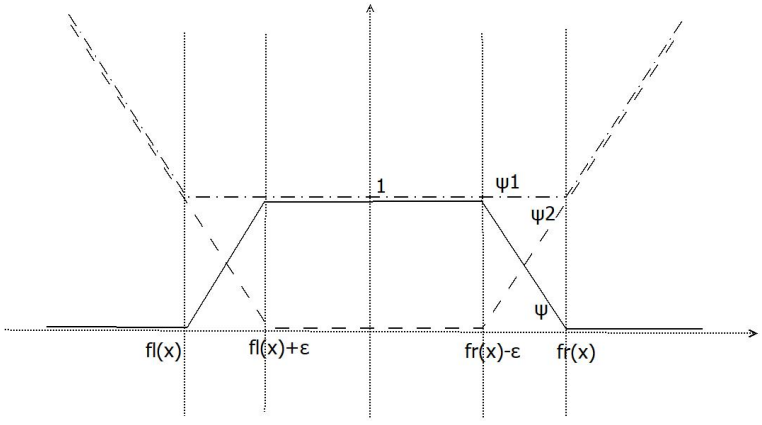

Figure 4.3: Loss Functions Ψ_ϵ^{in},Ψ_1^{in} and Ψ_2^{in}

convex functions. Denote that $f(x) = \phi(x)\beta$, where $\phi(\cdot)$ is a transformation of the original patient's features. If $\phi(x) = x$ then we are essentially looking for a linear rule, otherwise, the non-linearity is incorporated through the Kernel trick. The non-convex objective function $S(\beta)$ is the sum of the empirical risk function plus an ℓ_2 regularization term. We can construct two convex functions $S_1(\beta)$ and $S_2(\beta)$ such that $S(\beta) = S_1(\beta) - S_2(\beta)$. We define $S(\beta)$, $S_1(\beta)$ and $S_2(\beta)$ as below.

$$S(\beta) = \sum_{i=1}^{n} \Big[\frac{1}{P(A_i|X_i)} \Big(\big(1 - \alpha(X_i)\big) \Phi\big(Y_i - S(X_i)\big) \Psi\big(\phi(X_i)\beta, A_i\big) + \alpha(X_i) \Phi\big(S(X_i) - Y_i\big) \Psi\big(A_i, \phi(X_i)\beta\big) \Big) \Big] + \frac{\lambda_n}{2} \beta^T K(X_i, X_i) \beta$$

$$S_1(\beta) = \sum_{i=1}^{n} \Big[\frac{1}{P(A_i|X_i)} \Big(\big(1 - \alpha(X_i)\big) \Phi\big(Y_i - S(X_i)\big) \Psi_1\big(\phi(X_i)\beta, A_i\big) + \alpha(X_i) \Phi\big(S(X_i) - Y_i\big) \Psi_2\big(A_i, \phi(X_i)\beta\big) \Big) \Big] + \frac{\lambda_n}{2} \beta^T K(X_i, X_i) \beta$$

$$S_2(\beta) = \sum_{i=1}^{n} \Big[\frac{1}{P(A_i|X_i)} \Big(\big(1 - \alpha(X_i)\big) \Phi\big(Y_i - S(X_i)\big) \Psi_2\big(\phi(X_i)\beta, A_i\big) + \alpha(X_i) \Phi\big(S(X_i) - Y_i\big) \Psi_2\big(A_i, \phi(X_i)\beta\big) \Big) \Big]$$

According to the DC algorithm, β can be solved by repeatedly solving the following opti-

mization problem

$$
\begin{aligned}
\beta^{t+1} &= \arg\min_{\beta} \quad S_1(\beta) - [\nabla S_2(\beta^t)]^T(\beta - \beta^t) \\
&= \arg\min_{\beta} \quad S_1(\beta) - [\nabla S_2(\beta^t)]^T\beta \\
&= \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)\Big[W_i(\beta^t)\Psi_1\big(\phi(X_i)\beta, A_i\big) + Q_i(\beta^t)\Psi_1\big(A_i, \phi(X_i)\beta\big)\Big] \\
&+ \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\alpha(X_i)\Phi\big(S(X_i) - Y_i\big)\Big[\tilde{W}_i(\beta^t)\Psi_1\big(A_i, \phi(X_i)\beta\big) + \tilde{Q}_i(\beta^t)\Psi_1\big(\phi(X_i)\beta, A_i\big)\Big] \\
&+ \frac{\lambda_n}{2}\beta^T K(X_i, X_i)\beta
\end{aligned} \tag{4.8}
$$

where

$$
\begin{aligned}
W_i(\beta) &= I\big(A_i - \phi(X_i)\beta > -\epsilon\big) & \tilde{W}_i(\beta) &= I\big(A_i - \phi(X_i)\beta < \epsilon\big) \\
Q_i(\beta) &= I\big(A_i - \phi(X_i)\beta < -\epsilon\big) & \tilde{Q}_i(\beta) &= I\big(A_i - \phi(X_i)\beta > \epsilon\big)
\end{aligned}
$$

The optimization problem above boils down to the following quadratic programming problem which can be easily solved using standard packages. The details of the derivations are documented in Appendix A.1.

$$
\begin{aligned}
\arg\min_{\beta,b} \quad & \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}(\tilde{\alpha}_i - \alpha_i) < X_i, X_j > (\tilde{\alpha}_j - \alpha_j) - \sum_{i=1}^{n}(\tilde{\alpha}_i - \alpha_i)A_i \\
s.t. \quad & \sum_{i=1}^{n}\alpha_i - \sum_{i=1}^{n}\tilde{\alpha}_i = 0 \\
& 0 \le \alpha_i \le H_i W_i^t + \tilde{H}_i \tilde{Q}_i^t \\
& 0 \le \tilde{\alpha}_i \le H_i Q_i^t + \tilde{H}_i \tilde{W}_i^t
\end{aligned} \tag{4.9}
$$

where $H_i = \frac{I\left(Y_i > S(X_i)\right)\left(1-\alpha(X_i)\right)}{P(A_i|X_i)\lambda_n}$ corresponds to the penalty for patients who are outside

their IDIs but receive good outcomes and $\tilde{H}_i = \frac{I\big(S(X_i)>Y_i\big)\alpha(X_i)}{P(A_i|X_i)\lambda_n}$ represents the penalty for patients who are inside their IDIs but receive bad outcomes.

The solution of β are attained by a linear combination of X_i

$$\hat{\beta} = \sum_{i=1}^{n}(\tilde{\alpha}_i - \alpha_i)X_i$$

As for the intercept b, although there exists an analytic solution by applying K.K.T. conditions, we solve it using a line search for simplicity. Technically, any Kernel can be used for the Kernel matrix $K(X_i, X_i)$. In Section 6 and Section 7, the results of the linear Kernel and Gaussian Kernel are compared. The convergence rates in Section 5.3 are established for Gaussian Kernel.

The above procedures can be simplified if we are willing to make further assumptions. To see this, the optimization problem in each iteration can be transformed to:

$$\hat{f}_L^{t+1} = \arg\min_{f\in\mathcal{H}} \frac{1}{n}\sum_{i=1}^{n} = \frac{1}{P(A_i|X_i)}\mathcal{J}_i\,|f(X_i) - A_i| + \frac{\lambda_n}{2}||f||_2^2 \tag{4.10}$$

where

$$\mathcal{J}_i(\beta,\beta^t) = \big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)\Big[W_i(\beta^t)H_i(\beta) + Q_i(\beta^t)\tilde{H}_i(\beta)\Big] \tag{4.11}$$

$$+\,\alpha(X_i)\Phi\big(S(X_i) - Y_i\big)\Big[\tilde{W}_i(\beta^t)\tilde{H}_i(\beta) + \tilde{Q}_i(\beta^t)H_i(\beta)\Big] \tag{4.12}$$

Once we assume that $\hat{f}_{\beta^{t+1}}(X_i) - A_i$ and $\hat{f}_{\beta^t}(X_i) - A_i$ have the same sign, then the objective function (4.10) can be simplified as:

$$\begin{aligned} f_L^{t+1} = \arg\min_{f\in\mathcal{H}} \quad & \frac{\lambda_n}{2}||f||_2^2 + \\ & \frac{1}{n}\sum_{i=1}^{n}\frac{1}{P(A_i|X_i)}\Big[\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)I_{\left(0<\hat{f}^t(X_i)-A_i<\epsilon\right)} \\ & \qquad + \alpha(X_i)\Phi\big(S(X_i) - Y_i\big)I_{\left(0<A_i-\hat{f}^t(X_i)<\epsilon\right)}\Big]\Big|f(X_i) - A_i\Big| \end{aligned} \tag{4.13}$$

This simplified algorithm provides an intuitive illustration of the estimation procedure. For each iteration, only a small proportion of the sample can contribute to the optimization, i.e., only observations $\Big\{i : \ 0 < \hat{f}^t(X_i) - A_i < \epsilon \text{ or } 0 < A_i - \hat{f}^t(X_i) < \epsilon\Big\}$ are used for estimating f^{t+1}. This means the estimator $\hat{f}^t(X_i)$ only search its current neighborhood of radius ϵ to get $\hat{f}^{t+1}(X_i)$. From extensive simulation studies, we found the simplified algorithm general provides much faster estimation with a little compromise for the accuracy. The advantage of such simplicity is that equation (4.13) is equivalent to weighted Support Vector Machine which can be solved by standard libraries.

So far, we only discussed the case of the one-sided IDIs with a lower bound function. In the case where the upper bound is needed, the derivations above can be reproduced with minor modifications. Instead of deriving from the original non-convex optimization problems to its quadratic programming problem, we show its simplified form below.

$$\begin{aligned} f_U^{t+1} = \arg\min_{f\in\mathcal{H}} \quad & \frac{\lambda_n}{2}||f||_2^2 + \\ & \frac{1}{n}\sum_{i=1}^{n}\frac{1}{P(A_i|X_i)}\Big[\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)I_{(0<A_i-f^t(X_i)<\epsilon)} \\ & \qquad + \alpha(X_i)\Phi\big(S(X_i)-Y_i\big)I_{(0<f^t(X_i)-A_i<\epsilon)}\Big]\Big|f(X_i)-A_i\Big| \end{aligned} \tag{4.14}$$

4.3 DC Algorithm for Two-sided Dose Interval

The non-convex objective function for estimating two-sided IDIs below involves the empirical version of the relaxed risk (4.3) and two ℓ_2 norm regularization terms.

$$\begin{aligned} S(f_L, f_U) = & \sum_{i=1}^{n} H_i\Psi_\epsilon^{out}\Big(f_L(X), A, f_U(X)\Big) + \tilde{H}_i\Psi_\epsilon^{in}\Big(f_L(X), A, f_U(X)\Big) \\ & + \frac{\lambda_n}{2}||f_L||_2^2 + \frac{\lambda_n}{2}||f_U||_2^2 \end{aligned}$$

where H_i and $\tilde{H}_i$ are defined as in the previous section. For simplicity, the same tunning parameter λ_n is used for both f_L and f_U. In order to apply DC algorithm, we construct

the following convex decompositions

$$\Psi^{out} = \Psi_1^{out} - \Psi_2^{out}$$

$$\Psi^{in} = \Psi_1^{in} - \Psi_2^{in}$$

where

$$\Psi_1^{in}\Big(a,b,c\Big) = \begin{cases} b < a, & (a-b)/\epsilon + 1 \\ a < b < c, & 1 \\ c < b, & (b-c)/\epsilon + 1 \end{cases}$$

$$\Psi_2^{in}\Big(a,b,c\Big) = \begin{cases} b < (a+\epsilon) & (a+\epsilon-b)/\epsilon \\ (a+\epsilon) < b < (c-\epsilon), & 0 \\ (c-\epsilon) < b, & (b-c+\epsilon)/\epsilon \end{cases}$$

$$\Psi_1^{out}\Big(a,b,c\Big) = \Psi_1^{in}\Big(a,b,c\Big) - 1 = \begin{cases} b < a, & (a-b)/\epsilon \\ a < b < c, & 0 \\ c < b, & (b-c)/\epsilon \end{cases}$$

$$\Psi_2^{out}\Big(a,b,c\Big) = \begin{cases} b < (a-\epsilon) & (a-b)/\epsilon \\ (a-\epsilon) < b < (c+\epsilon), & 0 \\ (c+\epsilon) < b, & (b-c)/\epsilon \end{cases}$$

Here we assume that $c - a \leq 2\epsilon$ holds all the time. It can be easily shown that the non-convex objective function $S(f_L, f_U)$ can be decomposed to the difference of two convex

functions $S_1(f_L, f_U)$ and $S_2(f_L, f_U)$ where

$$S_1(f_L, f_U) = \sum_{i=1}^{n} H_i \Psi_1^{out}(f_L, A_i, f_U) + \sum_{i=1}^{n} \tilde{H}_i \Psi_1^{in}(f_L, A_i, f_U) + \frac{1}{2}||f_L||_2^2 + \frac{1}{2}||f_U||_2^2$$

and

$$S_2(f_L, f_U) = \sum_{i=1}^{n} H_i \Psi_2^{out}(f_L, A_i, f_U) + \sum_{i=1}^{n} \tilde{H}_i \Psi_2^{in}(f_L, A_i, f_U) \tag{4.15}$$

The DC algorithm iteratively solves the following problem

$$\begin{aligned}
&\arg \min_{\left(\phi(X_i)\beta_L + b_L\right) < \left(\phi(X_i)\beta_U + b_U\right)} S_1(\beta_L, \beta_U) - \nabla S_2(\beta_L, \beta_U) \begin{pmatrix} \beta_L - \beta_L^t \\ \beta_U - \beta_U^t \end{pmatrix} \\
&\arg \min_{\left(\phi(X_i)\beta_L + b_L\right) < \left(\phi(X_i)\beta_U + b_U\right)} S_1(\beta_L, \beta_U) - \nabla S_2(\beta_L, \beta_U) \begin{pmatrix} \beta_L \\ \beta_U \end{pmatrix} \\
&= \sum_{i=1}^{n} H_i \Big[W_i^L \Psi_1(f_L(X_i) - A_i) + Q_i^L \Psi_1(A_i - f_L(X_i)) \Big] \\
&+ \sum_{i=1}^{n} H_i \Big[W_i^U \Psi_1(A_i - f_U(X_i)) + Q_i^U \Psi_1(f_U(X_i) - A_i) \Big] \\
&+ \sum_{i=1}^{n} \tilde{H}_i \Big[\tilde{W}_i^L \Psi_1(f_L(X_i) - A_i) + \tilde{Q}_i^L \Psi_1(A_i - f_L(X_i)) \Big] \\
&+ \sum_{i=1}^{n} \tilde{H}_i \Big[\tilde{W}_i^U \Psi_1(A_i - f_U(X_i)) + \tilde{Q}_i^U \Psi_1(f_U(X_i) - A_i) \Big] \\
&+ \frac{1}{2}||f_L||_2^2 + \frac{1}{2}||f_U||_2^2
\end{aligned}$$

where $W_i^L = 1_{A_i > (f_L^t - \epsilon)}$, $\tilde{W}_i^L = 1_{A_i > (f_L^t + \epsilon)}$, $Q_i^L = 1_{A_i < (f_L^t - \epsilon)}$, $\tilde{Q}_i^L = 1_{A_i < (f_L^t + \epsilon)}$ and $W_i^U = 1_{A_i < (f_U^t + \epsilon)}$, $\tilde{W}_i^U = 1_{A_i < (f_U^t - \epsilon)}$, $Q_i^U = 1_{A_i > (f_U^t + \epsilon)}$, $\tilde{Q}_i^U = 1_{A_i > (f_U^t - \epsilon)}$.

The first two terms in (4.15) came with penalty $H_i = \frac{I\left(Y_i > S(X_i)\right)\left(1 - \alpha(X_i)\right)}{P(A_i|X_i)\lambda_n}$ for patients who are outside the IDIs but receive good outcomes and the second two terms are associated with penalty $\tilde{H}_i = \frac{I\left(S(X_i) > Y_i\right)\alpha(X_i)}{P(A_i|X_i)\lambda_n}$ for patients who are inside the IDIs but receive bad outcomes.

As in the previous section, each single iteration of the DC algorithm is eventually transformed to the following quadratic programming problem. The details of the derivations can be found in Appendix A.2.

$$\begin{aligned}
\arg\min_{\beta,b} \quad & \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}(\tilde{\alpha}_i^L - \alpha_i^L - m_i) < X_i, X_j > (\tilde{\alpha}_j^L - \alpha_j^L - m_j) \\
& + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}(\tilde{\alpha}_i^U - \alpha_i^U + m_i) < X_i, X_j > (\tilde{\alpha}_j^U - \alpha_j^U + m_j) \\
& - \sum_{i=1}^{n}(\tilde{\alpha}_i^L - \alpha_i^L + \tilde{\alpha}_i^U - \alpha_i^U)A_i - 2\epsilon\sum_{i=1}^{n} m_i \qquad (4.16) \\
s.t. \quad & 0 \le \alpha_i^L \le H_i W_i^L + \tilde{H}_i \tilde{W}_i^L \\
& 0 \le \tilde{\alpha}_i^L \le H_i Q_i^L + \tilde{H}_i \tilde{Q}_i^L \\
& 0 \le \alpha_i^U \le H_i Q_i^U + \tilde{H}_i \tilde{Q}_i^U \\
& 0 \le \tilde{\alpha}_i^U \le H_i W_i^U + \tilde{H}_i \tilde{W}_i^U \\
& 0 \le m_i \\
& \sum_{i=1}^{n} \alpha_i^L - \tilde{\alpha}_i^L + m_i = 0 \\
& \sum_{i=1}^{n} \alpha_i^U - \tilde{\alpha}_i^U - m_i = 0 \qquad (4.17)
\end{aligned}$$

After solving the problem above, the parameters can be recovered by

$$\begin{aligned}
\hat{\beta}_L &= \sum_{i=1}^{n}(\tilde{\alpha}_i^L - \alpha_i^L - m_i)X_i \\
\hat{\beta}_U &= \sum_{i=1}^{n}(\tilde{\alpha}_i^U - \alpha_i^U + m_i)X_i
\end{aligned}$$

Until now the estimation of intercept b_L and b_U is neglected. Although these two parameters can technically be solved by applying K.K.T. conditions, it will involve rather complicated derivations. Instead, we use a two-dimensional grid search to solve the intercepts.

In (4.17), the first two constraints are for the lower bound function while the second

two constraints are for the upper bound function. The positive parameter m_i corresponds to the gap $f_U(X_i) - f_L(X_i)$ for each observation which enforces every training observation to satisfy the constraint that $f_L(X_i) < f_U(X_i)$. However, for the testing dataset, this constraint can be violated for a small proportion of the sample. Our simulation and real data analysis show that this is generally not a severe problem, as the structural assumptions assume away the potential violations asymptotically.

Similar to the one-sided case, there is a simplified version of the two-sided IDI estimation. From equation (4.16) and Appendix A.7, it can be inferred that f_L and f_U can be solved independently if we do not enforce the constraint $f_L(X_i) < f_U(X_i)$ in the training data. This is only reasonable when this constraint is believed to be enforced by the finite sample as the structural assumptions in Section 3 only ensure that it is satisfied asymptotically. For each iteration in the DC algorithm, the lower bound function $\hat{f}_L^{(t)}(x)$ and the upper bound function $\hat{f}_U^{(t)}(x)$ are optimized in a sequence, with $\hat{f}_U^{(t)}(x)$ being the upper bound for $\hat{f}_L^{(t+1)}(x)$ and $\hat{f}_L^{(t+1)}(x)$ being the lower bound for $\hat{f}_U^{(t+1)}(x)$.

To summarize the simplified algorithm, we iteratively update $\hat{f}_L$ and $\hat{f}_U$ by solving the following two optimization tasks, which are not much different than (4.13) and (4.14).

$$\hat{f}_L^{t+1} = \arg \min_{f_L \in \mathcal{H}, f_L(X_i) < \hat{f}_U^t(X_i)} \frac{\lambda_n}{2}||f_L||_2^2 \\ \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)} \Big((1-\alpha)\Phi(Y_i - S^*)W_i + \alpha\Phi(S^* - Y_i)\tilde{W}_i \Big) \Big| f_L(X_i) - A_i \Big|$$

$$\hat{f}_U^{t+1} = \arg \min_{f_U \in \mathcal{H}, \hat{f}_L^t(X_i) < f_U(X_i)} \frac{\lambda_n}{2}||f_U||_2^2 \\ \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)} \Big((1-\alpha)\Phi(Y_i - S^*)\tilde{W}_i + \alpha\Phi(S^* - Y_i)W_i \Big) \Big| f_U(X_i) - A_i \Big|$$

where $W_i = I_{(0 < f^t(X_i) - A_i < \epsilon)}$ and $\tilde{W}_i = I_{(0 < A_i - f^t(X_i) < \epsilon)}$

Chapter 5

Theoretical Results

In this Chapter, theorems and corollaries will be proved to show the theoretical properties of the proposed methods. In Section 5.2, we will prove two theorems to show that the Fisher consistency holds under reasonable assumptions. Therefore, by minimizing the Quadrant risk functions, we will attain the IDIs introduced in Section 1.3 with the desired causal interpretations. In Section 5.3, we first discuss the convergence of the relaxed risk function (4.1) and (4.3) to the original Quadrant risk functions (2.8) and (2.10). And then the convergence rate of the estimated IDI bound functions to the optimal IDI bound functions are shown. Throughout this section, we assume that $\alpha = \frac{1}{2}$ for all EDIs.

5.1 Optimal Bound Functions

In this section, we define the optimal bound functions that we intend to recover from the sample. The functions satisfying the definition might not be unique for both the PDI or the EDI unless extra uniqueness assumptions are made. These definitions will be needed in the next section to prove the Fisher consistency of the proposed methods.

1. One-sided PDI optimal lower bound function:
   - Any measurable function $f_{I,L,opt}(x)$ that $P\Big(Y\big(f_{I,L,opt}(x), x\big) > S(x)\Big) = \alpha(x)$
2. Two-sided PDI optimal lower bound function:

- Any measurable function $f_{I,L,opt}(x)$ that $P\Big(Y\big(f_{I,L,opt}(x),x\big) > S(x)\Big) = \alpha(x)$
- For any $a < f_{I,L,opt}(x)$, $P\big(Y(a,x) > S(x)\big) \leq \alpha(x)$

Two-sided PDI optimal upper bound function:

- Any measurable function $f_{I,U,opt}(x)$ that $P\Big(Y\big(f_{I,U,opt}(x),x\big) > S(x)\Big) = \alpha(x)$
- For any $a > f_{I,U,opt}(x)$, $P\big(Y(a,x) > S(x)\big) \leq \alpha(x)$

3. One-sided EDI optimal lower bound function:

- Any measurable function $f_{h,L,opt}(x)$ that $E\Big[Y(f_{h,L,opt}(x),x)\Big] = S(x)$

4. Two-sided EDI optimal lower bound function:

- Any measurable function $f_{h,L,opt}(x)$ that $E\Big[Y(f_{h,L,opt}(x),x)\Big] = S(x)$
- For any $a < f_{h,L,opt}(x)$, $E\left[Y(a,x)\right] \leq S(x)$

Two-sided EDI optimal upper bound function:

- Any measurable function $f_{h,U,opt}(x)$ that $E\Big[Y(f_{h,U,opt}(x),x)\Big] = S(x)$
- For any $a > f_{h,U,opt}(x)$, $E\left[Y(a,x)\right] \geq S(x)$

The IDI bound functions above are optimal in the sense that they are the minimizers of corresponding Quadrant risk functions, i.e.,

- $f_{I,L,opt} = \arg\min_f R_I(f)$
- $\Big[f_{I,L,opt}, f_{I,U,opt}\Big] = \arg\min_{f_L \leq f_U} R_I(f_L, f_U)$
- $f_{h,L,opt} = \arg\min_f R_h(f)$
- $\Big[f_{h,L,opt}, f_{h,U,opt}\Big] = \arg\min_{f_L \leq f_U} R_h(f_L, f_U)$

These claims will be verified in the proofs of theorem 5.2.1 to Corollary 5.2.2.

5.2 Fisher Consistency

In this section, Fisher consistency of the OWL-based approach will be discussed. Theorem 5.2.1 shows that minimizing $R_I(\cdot)$ results in the PDI lower bound, above which the doses will produce outcomes larger than $S(x)$ with probability at least $\alpha(x)$. Theorem 5.2.2 shows that minimizing $R_h(\cdot)$ implies the EDI lower bound that the doses greater than the bound are expected to generate outcomes larger than $S(x)$, while the $\alpha(x)$ is set to be a constant $\frac{1}{2}$ for interpretability. Corollary 5.2.1 and corollary 5.2.2 are the two-sided counter-parts of theorem 5.2.1 and theorem 5.2.2.

It should be emphasized that the IDIs are determined by a threshold function $S(x)$ and a probability function $\alpha(x)$. The combinations of these two functions imply different interpretations and require different assumptions. It is important that these two functions are determined coherently with the interpretations and the assumptions are supported by data.

Theorem 5.2.1. *For the one-sided PDI with optimal lower bound function* $f_{I,L,opt} = \arg\min_f R_I(f)$*, any measurable function* $a(\cdot)$ *s.t.* $f_{I,L,opt}(x) \leq a(x) \leq a_U$*,* $\forall x \in \mathcal{X}$*, the potential outcome* $Y(a(x), x)$ *satisfies* $P(Y(a(x), x) \geq S(x)) \geq \alpha(x)$ *under the partial monotonicity assumption for PDI* (3.1)*, i.e.,*

1. $P\Big(Y(a_L, x) > S(x)\Big) \leq \alpha(x)$ *and* $P\Big(Y(a_U, x) > S(x)\Big) \geq \alpha(x)$
2. $P(Y(a, x) > S(x))$ *across* $\alpha(x)$ *at most once as* a *goes from* a_L *to* a_U

Proof. A brief proof can be found in Appendix A.3. □

Corollary 5.2.1. *For the two-sided PDI with optimal bound functions* $[f_{I,L,opt}, f_{I,U,opt}] = \arg\min_{f_L \leq f_U} R_I(f_L, f_U)$*, any measurable function* $a(\cdot)$ *s.t.* $f_{I,L,opt}(x) \leq a(x) \leq f_{I,U,opt}(x)$*,* $\forall x \in \mathcal{X}$*, the potential outcome* $Y(a(x), x)$ *satisfies* $P(Y(a(x), x) \geq S(x)) \geq \alpha(x)$ *, under the partial unimodality assumptions for PDI* (3.2)*, i.e.,*

1. $P\Big(Y(a_L, x) > S(x)\Big) \leq \alpha(x)$ *and* $P\Big(Y(a_U, x) > S(x)\Big) \leq \alpha(x)$

2. $\exists a^* \in [a_L, a_U],\ P(Y(a^*, x) > S(x)) \geq \alpha(x)$

3. $P(Y(a, x) > S(x))$ *cross* $\alpha(x)$ *at most twice as* a *goes from* a_L *to* a_U

Proof. A brief proof can be found in Appendix A.4. □

Theorem 5.2.2. *For the one-sided EDI with optimal lower bound function* $f_{h,L,opt} = \arg\min_f R_h(f)$*, any measurable function* $a(\cdot)$ *s.t.* $f_{h,L,opt}(x) \leq a(x) \leq a_U$*,* $\forall x \in \mathcal{X}$*, the potential outcome* $Y(a(x), x)$ *satisfies* $E[Y(a(x), x)] \geq S(x)$ *under the partial monotonicity assumptions for EDI* (3.3)*, i.e.,*

1. $E[Y(a_L, x)] \leq S(x)$ *and* $E[Y(a_U, x)] \geq S(x)$

2. $E(Y(a, x))$ *across* $S(x)$ *at most once as* a *goes from* a_L *to* a_U

Proof. A brief proof can be found in Appendix A.5. □

Corollary 5.2.2. *For the two-sided EDI with optimal bounds functions* $[f_{h,L,opt}, f_{h,U,opt}] = \arg\min_{f_L \leq f_U} R_h(f_L, f_U)$*, any measurable function* $a(\cdot)$ *s.t.* $f_{h,L,opt}(x) \leq a(x) \leq f_{h,U,opt}(x)$*,* $\forall x \in \mathcal{X}$*, the potential outcome* $Y(a(x), x)$ *satisfies* $E[Y(a(x), x)] \geq S(x)$ *under the partial unimodality assumption for EDI* (3.4)*, i.e.,*

1. $E[Y(a_L, x)] \leq S(x)$ *and* $E[Y(a_U, x)] \leq S(x)$

2. $\exists a^* \in [a_L, a_U],\ E[Y(a^*, x)] \geq S(x)$

3. $E(Y(a, x))$ *cross* $S(x)$ *at most twice as* a *goes from* a_L *to* a_U

Proof. A brief proof can be found in Appendix A.6. □

5.3 Convergence Rate

In this section, the convergence rates of the Quadrant risk of the estimated IDIR are derived. First, theorem 5.3.1 shows that, for one-sided PDI lower bounds, the difference between risk function $R_I(\cdot)$ and its non-convex relaxation $R_{I,\epsilon_n}(\cdot)$ converges to zero with the same rate

as $\epsilon_n \to 0$. Similarly, theorem 5.3.2 shows that $R_h(\cdot)$ converges to its non-convex relaxation $R_{h,\epsilon_n}(\cdot)$ with the same rate as $\epsilon_n \to 0$. Corollary 5.3.1 and corollary 5.3.2 show that the same results hold for two-sided intervals. Theorem 5.3.3 is adapted from the Chen et al. (2016), where the convergence rate of one-sided IDIR is obtained based on general results of empirical risk minimization and approximation theorem. Finally, in theorem 5.3.3, we provide the convergence rate of the two-sided IDIR.

Theorem 5.3.1. *For any measurable function $f : \mathcal{X} \mapsto \mathbb{R}$, we have $|R_I(f) - R_{I,\epsilon_n}(f)| \leq C\epsilon_n$, where C is a constant.*

Proof. A brief proof can be found in Appendix A.7. □

Corollary 5.3.1. *For any interval of two measurable functions $[f_L, f_U]$, where $f_L, f_U : \mathcal{X} \mapsto \mathbb{R}$ and $f_L(x) \leq f_U(x)\ \forall x \in \mathcal{X}$, we have $|R_I(f_L, f_U) - R_{I,\epsilon_n}(f_L, f_U)| \leq C\epsilon_n$, where C is a constant.*

Proof. A brief proof can be found in Appendix A.8. □

Theorem 5.3.2. *For any measurable function $f : \mathcal{X} \mapsto \mathbb{R}$, then $|R_h(f) - R_{h,\epsilon_n}(f)| \leq C_0\epsilon_n$, under the assumption that $E\left[\left|Y - E[Y(a,x)]\right| \middle| A = a, X = x\right] \leq C_1$, where C_0 and C_1 are two constants.*

Proof. A brief proof can be found in Appendix A.9. □

Corollary 5.3.2. *For any interval of two measurable functions $[f_L, f_U]$ where $f_L, f_U : \mathcal{X} \mapsto \mathbb{R}$ and $f_L(x) \leq f_U(x)\ \forall x \in \mathcal{X}$, we have $|R_h(f_L, f_U) - R_{h,\epsilon_n}(f_L, f_U)| \leq C_0\epsilon_n$, under the assumption that $E\left[\left|Y - E[Y(a,x)]\right| \middle| A = a, X = x\right] \leq C_1$, where C_0 and C_1 are two constants.*

Proof. A brief proof can be found in Appendix A.10. □

Theorem 5.3.3. *Under the assumptions that $f_{\Phi,L,opt} \in B^{\delta}_{1,\infty}(\mathbb{R}^d)$, a Besov space, such that, $B^{\delta}_{1,\infty}(\mathbb{R}^d) = \{f \in L_\infty((\mathbb{R}^d)) : \sup_{t>0}(t^{-\delta} w_{(r,L_1)}(f,t)) < \infty\}$, where w is the modulus*

of continuity. Then for any $\zeta > 0$, $d/(d+\tau) < p < 1$, $\tau > 0$, *and the parameter* γ_n *for the Gaussian Kernel,*

$$R_\Phi(\hat{f}_n) - R_\Phi(f_{opt}) \leq c_1 \frac{\lambda_n}{\gamma_n^d} + c_2\gamma_n^\delta + c_3 \frac{1}{\gamma_n^{\frac{(1-p)(1+\zeta)d}{2-p}} \lambda_n^{\frac{p}{2-p}} n^{\frac{1}{2-p}}} + c_4 \frac{\tau^{\frac{1}{2}}}{n^{\frac{1}{2}}} + c_5 \frac{\tau}{n} + c_6\epsilon_n \quad (5.1)$$

with probability $1-3e^{-\tau}$. *Here* d *is the dimension of* $\mathcal{X}$ *and* Φ *can be either* $\Phi = I$ *or* $\Phi = h$, *i.e., the rate above applies for both PDI and EDI. .*

Properly choosing the constant as

$$\gamma_n \propto \left(\frac{1}{n}\right)^{\frac{1}{2\delta+d}} \qquad \lambda_n \propto \left(\frac{1}{n}\right)^{\frac{\delta+d}{2\delta+d}} \qquad \epsilon_n = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right) \quad (5.2)$$

We have the following convergence rate with with probability $1 - 3e^{-\tau}$.

$$R_\Phi(\hat{f}_L) - R_\Phi(f_{opt}) = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right) \quad (5.3)$$

Proof. A brief proof can be found in Appendix A.11. □

Corollary 5.3.3. *Under the assumptions that* $f_{\Phi,L,opt} \in B^\delta_{1,\infty}(\mathbb{R}^d)$ *and* $f_{\Phi,U,opt} \in B^\delta_{1,\infty}(\mathbb{R}^d)$, *such that,* $B^\delta_{1,\infty}(\mathbb{R}^d) = \{f \in L_\infty((\mathbb{R}^d)) : \sup_{t>0}(t^{-\delta} w_{(r,L_1)}(f,t)) < \infty\}$, *where* w *is the modulus of continuity. If it is further assumed that there exists a measurable function* $f_{M,opt}(x)$, *which is known or can be estimated consistently, such that* $f_{L,opt}(x) < f_M(x) < f_{U,opt}(x)$ *for all* $x \in \mathcal{X}$. *Then for any* $\eta > 0$, $d/(d+\tau) < p < 1$, $\tau > 0$, $0 < p' < 1$,*and the parameter* γ_n *for the Gaussian Kernel, we have*

$$R_\Phi(\hat{f}_{L,n}, \hat{f}_{U,n}) - R_\Phi(f_{L,opt}, f_{U,opt}) = \mathcal{O}_p\left(n^{-\frac{\delta}{2\delta+d}}\right) \quad (5.4)$$

with probability $p' + \left(1 - 3e^{-\tau}\right)^2 - 1$. *Here* d *is the dimension of* $\mathcal{X}$ *and* Φ *can be either* $\Phi = I$ *or* $\Phi = h$, *i.e., the rate above applies for both PDI and EDI.*

Proof. A brief proof can be found in Appendix A.12. □

The theorem 5.3.3 and corollary 5.3.3 prove that the OWL-based IDI method can potentially achieve a rate of nearly $\mathcal{O}(n^{-\frac{1}{2}})$, when the δ, the smoothness parameter of optimal bound functions f_L and f_U, is relatively large compared to the d, the dimension of the covariates. This rate is essentially faster than the convergence rate in Chen et al. (2016), which is nearly $\mathcal{O}(n^{-\frac{1}{4}})$. We believe that this difference is due to the fact that in the IDIR problem, the objective is to find a compact set, while in the IDR problem, the target is a single function, which makes the problem significantly more difficult.

Chapter 6

Simulation

In this Chapter, simulation settings are designed to examine the efficiency of our OWL-based methods and the indirect methods in various scenarios. In Section 6.1, we present settings without prognostic effects. In Section 6.2, we design alternative data generating model in presence of prognostic effects, where the personalization of $S(x)$ has to rely on independent estimation steps.

The detailed two-step procedure used by the indirect methods can be found in Section 1.4 and Section 2.4. This procedure involves an outcome modeling step based on the classification models and a grid search step to find the PDIs. For the EDI, the outcome modeling step is based on the regression models. The indicator $I\{Y > S\}$ is used as the binary outcome for each classification model for which we build a model that $\hat{P}(Y > S(X)|a, X) \coloneqq h(A, X)$, and then search through the domain of A to find a range of doses $\mathcal{A}_1$ that for $\forall a \in \mathcal{A}_1$, $h(a, X) > \alpha(X)$. Similarly, Y serves as the numeric outcome for each regression model for which we build a model that $\hat{E}[Y|A, X] \coloneqq h(A, X)$, and then conduct a grid search to find a range $\mathcal{A}_1$ such that for $\forall a \in \mathcal{A}_1$, we have $h(a, X) > S(X)$. The effectiveness of both the indirect methods and our OWL-based methods are examined by evaluating the empirical risk functions on an independent testing dataset. The empirical risk functions, as defined in Chapter 2, are $\hat{R}_{I,\epsilon}$ for the PDIs and $\hat{R}_{h,\epsilon}$ for the EDIs.

Three types of classification methods are included in this simulation. Logistic Regression

incorporates all the main effect terms of dose and covariates, as well as the quadratic term of dose and interactions of dose and covariates. Besides, ℓ_2 regularization is used but the quadratic term is exempted from the regularization, as is suggested by Cai and Tian (2016). The coefficient of regularization, λ_n, is the only tuning parameter. Support Vector Machine relies on a Gaussian Kernel to capture the non-linearity and interactions between the dose and the covariates. The cost C is the only tunning parameter while the Kernel bandwidth is fixed at $\frac{1}{\sqrt{d}}$. The Random Forest Classification (RF-C) with a forest size 1000 is the only tree-based classifier in this study. The percentage of variables used for each split is the only tuning parameter.

Three types of regression methods are selected for this simulation. Lasso uses the same model as the Logistic Regression except the link function being linear instead of the logit function. Support Vector Regression (SVR) applies the ϵ-insensitive regression method (Drucker et al., 1997) and approximates the conditional mean by the conditional median. The Gaussian Kernel with bandwidth $\frac{1}{\sqrt{d}}$ is still used for SVR. The Random Forest Regression (RF-R) with a forest size 1000 shares the same robustness as the RF-C method. Parameter settings for the regression methods are the same as the corresponding classification cases.

For our OWL-based approach, the Kernel trick allows us to incorporate many types of Kernels. However, to establish easy comparisons with SVM and SVR, we only include the linear Kernel, denoted as OWL-L, and the Gaussian Kernel denoted as OWL-K. Although the convergence rate in Section 5.3 is based on the Gaussian Kernel, we will show that linear Kernel sometimes achieves no worse performance even for nonlinear IDIRs. Throughout this Chapter, as well as Chapter 7, all the tuning parameters are selected using 5-folds cross-validation while the initializations of OWL-based methods are done by using constant lower/upper bounds through population estimations without including patients' covariates.

6.1 Basic Settings

In this section, multiple simulation settings are prepared with varying sample sizes, numbers of covariates, and data generating processes. Scenario1 and Scenario2 are designed for the Monotonic case to find one-sided IDIs, where the $Y(a, x)$ has only one intersection with $Y = S^*$ for a given x. The relationship between A and Y is depicted using a logistic function while the heterogeneity is captured by manipulating the center parameter C_i as a function of X_i. Scenario3 and Scenario4 are designed for the Unimodal case to find the two-sided IDIs, where $Y(a, x)$ has two intersections with $Y = S^*$ for a given x. Negative absolute distance is used when constructing the mean generating function, while the center parameter C_i is a function of features X_i. According to Section 3.1, only Partial Monotonicity and Partial Unimodality assumptions are needed for finding IDIs. But instead we used Strict Monotonicity and Strict Unimodality for simplicity. Independent Gaussian noises are added to all the outcome variables. The four scenarios are documented as follows. The confidence probability αs are kept to $\frac{1}{2}$ for all the four scenarios.

1. Monotonicity with Linear Bound

$$S = 2.5$$

$$C_i = 0.3 * X_{i1} + 0.3 * X_{i2} + 0.3 * X_{i3}$$

$$\mu_i = \frac{re^{rA_i}}{e^{rC_i} + e^{rA_i}} \qquad r = 10$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 1)$$

2. Monotonicity with Non-linear Bound

$$S = 2.5$$

$$C_i = 0.75 * \log(|X_{i1}| + 1) - 0.2 * cos(\pi X_{i2}) - 0.2 * I(X_{i3} > 0) - 0.3896$$

$$\mu_i = \frac{re^{rA_i}}{e^{rC_i} + e^{rA_i}} \qquad r = 10$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 1)$$

3. Unimodality with Linear Bounds

$$S = 0.5$$

$$C_i = 0.15 * X_{i1} + 0.15 * X_{i2} + 0.15 * X_{i3}$$

$$\mu_i = 1 - |A_i - C_i|$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 0.2)$$

4. Unimodality with Non-linear Bounds

$$S = 0.5$$

$$C_i = 0.75 * \log(|X_{i1}| + 1) - 0.2 * cos(\pi X_{i2}) - 0.2 * I(X_{i3} > 0) - 0.3896$$

$$\mu_i = 1 - |A_i - C_i|$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 0.2)$$

where Xs are independent random variables following uniform distribution between $[-1, 1]$ and A is generated from a normal distribution independent with Xs. Figure 6.1 provides scatter plots of the simulated datasets both with and without noise. The simulation results based on 100 repetitions are summarized in Table 6.1 and Table 6.1, where the average Quadrant Losses for all settings (average weighted misclassification for the PDI cases and average weighted hinge loss for the EDI cases) are documented with the standard deviations in the parentheses. Figure 6.2 plotted the trends of the average Quadrant Losses for the PDIs with different sample sizes and Figure 6.3 plotted the trends of the average Quadrant Losses for the EDIs. The solid lines stand for the two OWL-based methods, OWL-L and OWL-K. The dashed lines represent the indirect methods.

Notice that in most cases, OWL-L and OWL-N have advantages over the indirect meth-

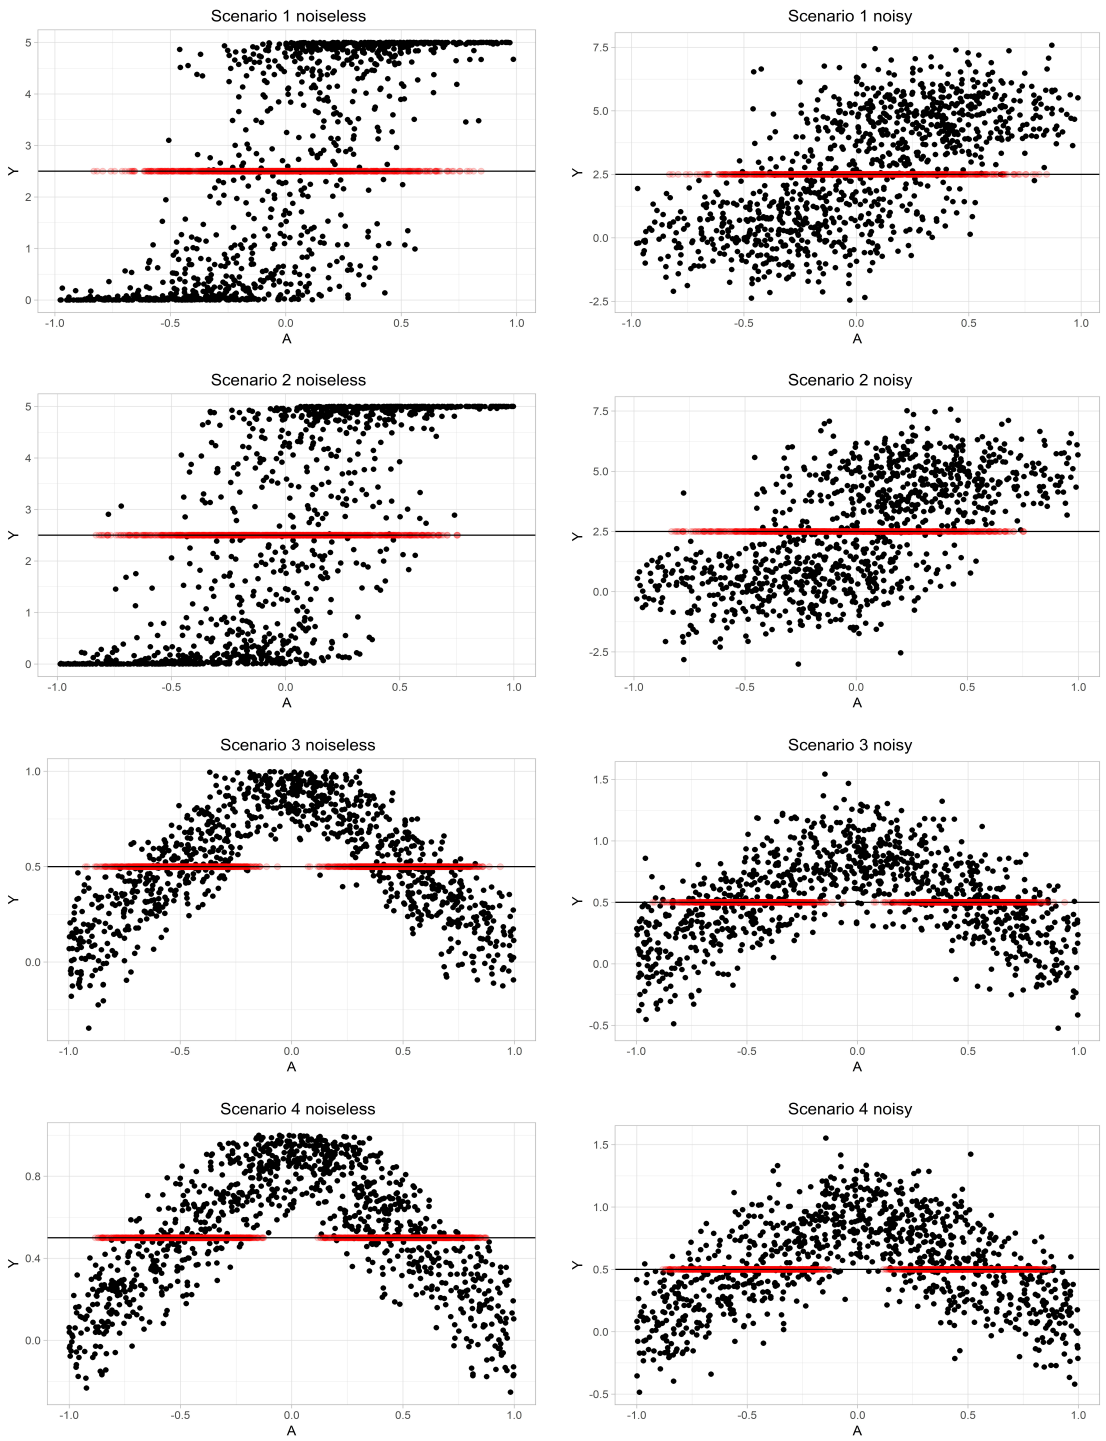

Figure 6.1: Scatter Plots of Simulation Scenarios

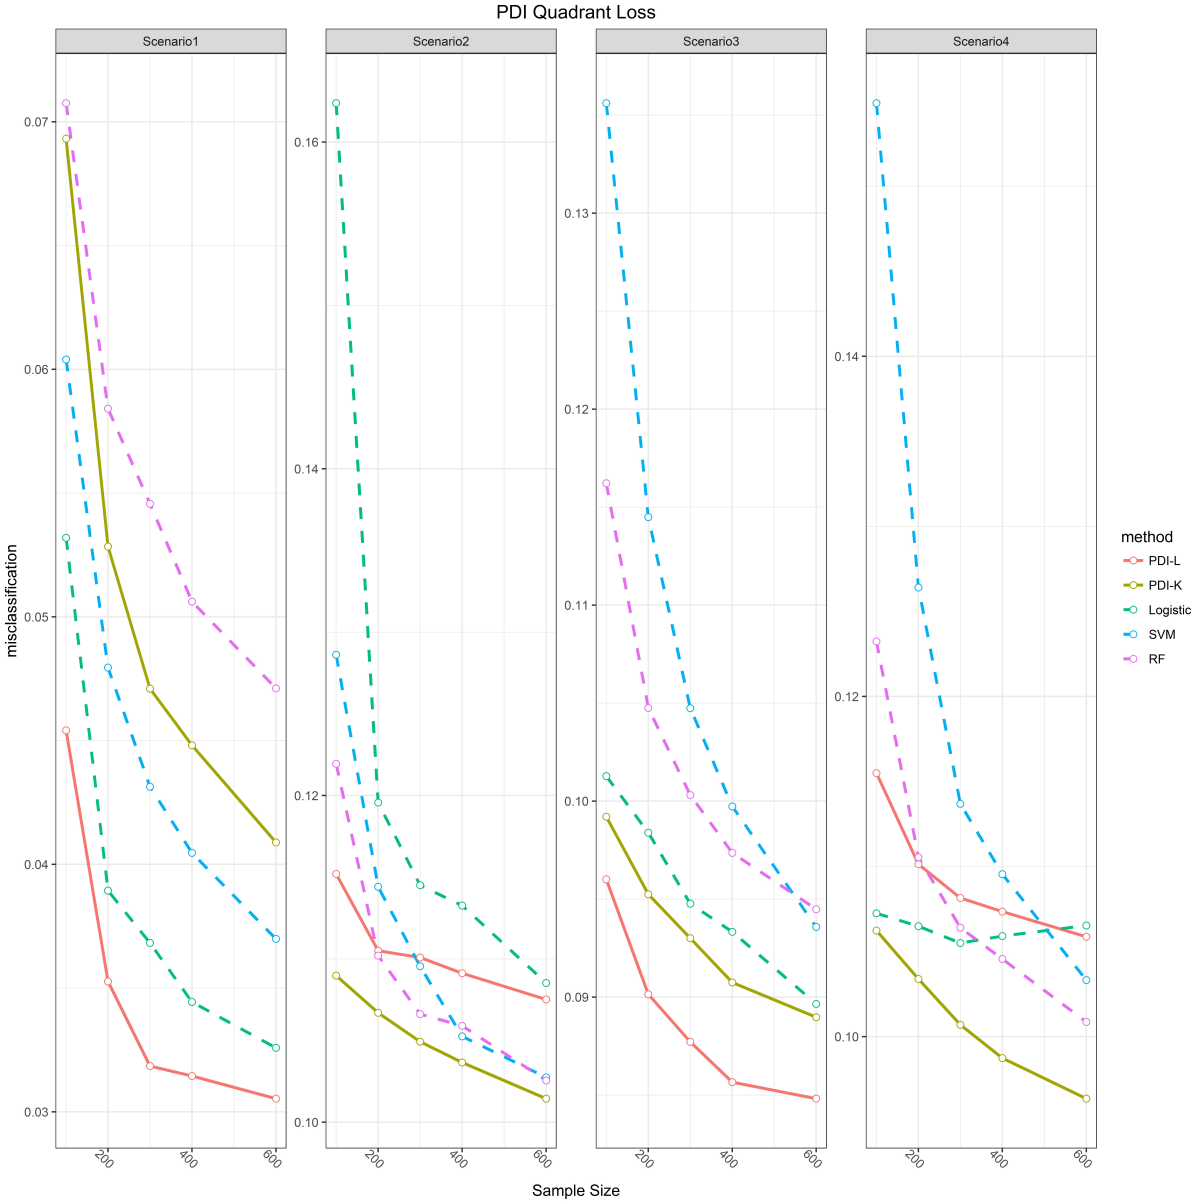

Figure 6.2: PDI Average Quadrant Losses for Simulation Settings in Section 6.1

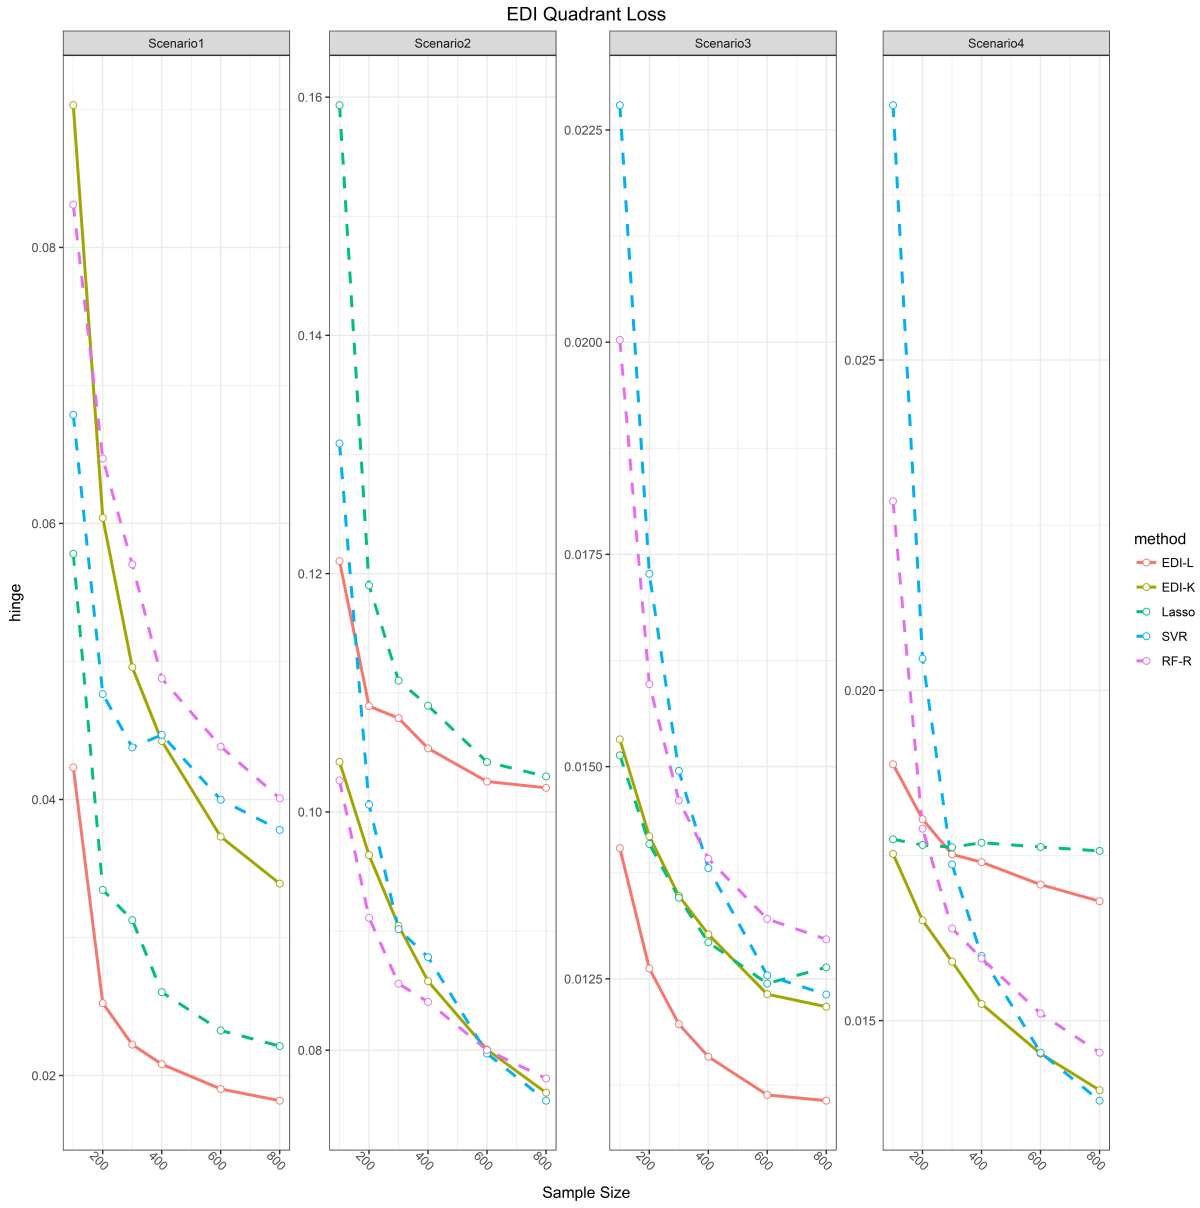

Figure 6.3: EDI Average Quadrant Losses for Simulation Settings in Section 6.1

	N	PDI-L	PDI-K	Logistic	SVM	RF-C
Scenario 1	100	**0.045 (0.010)**	0.069 (0.009)	0.053 (0.012)	0.060 (0.014)	0.071 (0.008)
	200	**0.035 (0.005)**	0.053 (0.005)	0.039 (0.005)	0.048 (0.006)	0.058 (0.007)
	300	**0.032 (0.003)**	0.047 (0.004)	0.037 (0.004)	0.043 (0.004)	0.055 (0.005)
	400	**0.031 (0.003)**	0.045 (0.004)	0.034 (0.004)	0.040 (0.003)	0.051 (0.004)
	600	**0.031 (0.002)**	0.041 (0.003)	0.033 (0.003)	0.037 (0.003)	0.047 (0.003)
Scenario 2	100	0.115 (0.007)	**0.109 (0.007)**	0.161 (0.021)	0.129 (0.008)	0.122 (0.008)
	200	0.111 (0.004)	**0.107 (0.006)**	0.119 (0.005)	0.115 (0.006)	0.111 (0.006)
	300	0.110 (0.003)	**0.105 (0.005)**	0.115 (0.004)	0.109 (0.006)	0.107 (0.004)
	400	0.109 (0.003)	**0.103 (0.003)**	0.114 (0.003)	0.105 (0.003)	0.106 (0.004)
	600	0.107 (0.003)	**0.102 (0.004)**	0.108 (0.002)	0.103 (0.003)	0.103 (0.004)
Scenario 3	100	**0.096 (0.006)**	0.099 (0.006)	0.101 (0.005)	0.136 (0.013)	0.116 (0.010)
	200	**0.090 (0.006)**	0.095 (0.004)	0.098 (0.005)	0.114 (0.008)	0.105 (0.007)
	300	**0.088 (0.003)**	0.093 (0.003)	0.095 (0.005)	0.105 (0.006)	0.100 (0.005)
	400	**0.086 (0.003)**	0.091 (0.003)	0.093 (0.005)	0.100 (0.005)	0.097 (0.004)
	600	**0.085 (0.002)**	0.089 (0.003)	0.090 (0.004)	0.094 (0.004)	0.094 (0.003)
Scenario 4	100	0.115 (0.006)	**0.106 (0.004)**	0.107 (0.003)	0.155 (0.012)	0.123 (0.011)
	200	0.110 (0.003)	**0.103 (0.003)**	0.106 (0.003)	0.126 (0.009)	0.111 (0.005)
	300	0.108 (0.003)	**0.101 (0.003)**	0.106 (0.002)	0.114 (0.006)	0.106 (0.004)
	400	0.107 (0.003)	**0.099 (0.004)**	0.106 (0.003)	0.110 (0.005)	0.105 (0.004)
	600	0.106 (0.003)	**0.096 (0.003)**	0.106 (0.002)	0.103 (0.004)	0.101 (0.003)

Table 6.1: PDI Quadrant Losses for Simulation Settings in Section 6.1

	N	EDI-L	EDI-K	Lasso	SVR	RF-R
Scenario 1	100	**0.0449 (0.0211)**	0.0913 (0.0172)	0.0580 (0.0217)	0.0649 (0.0197)	0.0841 (0.0184)
	200	**0.0252 (0.0048)**	0.0583 (0.0090)	0.0324 (0.0073)	0.0470 (0.0122)	0.0654 (0.0131)
	300	**0.0224 (0.0035)**	0.0493 (0.0080)	0.0325 (0.0096)	0.0445 (0.0133)	0.0578 (0.0095)
	400	**0.0213 (0.0028)**	0.0457 (0.0069)	0.0256 (0.0051)	0.0424 (0.0193)	0.0489 (0.0061)
	600	**0.0191 (0.0018)**	0.0375 (0.0048)	0.0234 (0.0034)	0.0406 (0.0153)	0.0439 (0.0044)
Scenario 2	100	0.1853 (0.0173)	0.1634 (0.0143)	0.2739 (0.0398)	0.2082 (0.0189)	**0.1289 (0.0176)**
	200	0.1743 (0.0086)	0.1556 (0.0143)	0.1864 (0.0136)	0.1779 (0.0096)	**0.1202 (0.0122)**
	300	0.1707 (0.0070)	0.1491 (0.0132)	0.1725 (0.0099)	0.1577 (0.0092)	**0.1175 (0.0094)**
	400	0.1657 (0.0057)	0.1404 (0.0069)	0.1691 (0.0060)	0.1505 (0.0075)	**0.1160 (0.0087)**
	600	0.1599 (0.0051)	0.1289 (0.0063)	0.1589 (0.0045)	0.1375 (0.0050)	**0.1087 (0.0094)**
Scenario 3	100	**0.0141 (0.0019)**	0.0153 (0.0016)	0.0151 (0.0015)	0.0223 (0.0028)	0.0204 (0.0035)
	200	**0.0126 (0.0014)**	0.0143 (0.0011)	0.0141 (0.0012)	0.0175 (0.0022)	0.0161 (0.0016)
	300	**0.0120 (0.0009)**	0.0135 (0.0010)	0.0132 (0.0009)	0.0149 (0.0012)	0.0147 (0.0010)
	400	**0.0116 (0.0006)**	0.0130 (0.0009)	0.0129 (0.0008)	0.0138 (0.0008)	0.0139 (0.0007)
	600	**0.0111 (0.0005)**	0.0123 (0.0007)	0.0124 (0.0006)	0.0125 (0.0007)	0.0132 (0.0006)
	100	0.0189 (0.0019)	**0.0175 (0.0018)**	0.0178 (0.0006)	0.0283 (0.0034)	0.0223 (0.0029)
	200	0.0180 (0.0009)	**0.0166 (0.0008)**	0.0177 (0.0007)	0.0205 (0.0019)	0.0178 (0.0016)
	300	0.0174 (0.0008)	**0.0158 (0.0010)**	0.0175 (0.0007)	0.0176 (0.0014)	0.0165 (0.0007)
	400	0.0174 (0.0007)	**0.0153 (0.0009)**	0.0177 (0.0007)	0.0160 (0.0010)	0.0159 (0.0007)
	600	0.0171 (0.0005)	**0.0145 (0.0007)**	0.0176 (0.0006)	**0.0145 (0.0007)**	0.0151 (0.0006)

Table 6.2: EDI Quadrant Losses for Simulation Settings in Section 6.1

ods in terms of accuracy. The standard deviations of these two OWL-based methods are at least comparable to those of the indirect methods. The EDI in Scenario 2 makes one exception, where the RF-R exhibits superior accuracy and robustness in several settings. But still, OWL-K is the second best for all the sample sizes. It is possible that the indirect methods can have average Quadrant Losses smaller than those of the OWL-based methods when the sample size is greater than 1000. This is reasonable because the qualities of the outcome models are guaranteed when the sample size is large and they might eventually attain advantages over the OWL-based methods. Therefore, with the above simulation results, we claim only the small sample advantages of the proposed OWL-based method.

6.2 In Presence of Prognostic Effects

In the previous simulations, the threshold function $S(x)$ and the confidence probability $\alpha(x)$ are chosen as constant so that the assumptions in Section 3 hold for each patient. This means that the trajectory of dose-response will always have at least one intersection with the threshold. Even though such settings are straightforward for demonstrating the proposed methods, they do not imitate the real applications except for cases where the effects of the dose are much stronger than the prognostic effects.

Therefore, the below simulation settings are designed in order to test the robustness of our approach in face of prognostic effects. When strong prognostic effects are present, it might not be reasonable to choose a constant threshold function. Instead, we build linear models with up to cubic terms to model the outcome without using the dose and then take the fitted outcome as the value of threshold function, i.e., $S(x) = \hat{E}[Y|x]$. Since it is recommended in Section 2.4 that only one of $S(x)$ and $\alpha(x)$ is personalized, we keep $\alpha(x) = \frac{1}{2}$ as in previous sections and only personalize $S(x)$. The four redesigned settings are as follows.

1. **Monotonic with Linear Bound**

$$C_i = 0.3 * X_{i1} + 0.3 * X_{i2} + 0.3 * X_{i3}$$

$$D_i = 0.6 * X_{i2} + 0.6 * X_{i3} + 0.6 * X_{i4}$$

$$A_i \sim \mathcal{N}(C_i, 0.5)$$

$$\mu_i = \frac{re^{rA_i}}{e^{rC_i} + e^{rA_i}} + D_i \quad r = 10$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 1)$$

$$S_i = fitted(lm(Y \sim X))_i$$

2. **Monotonic with Non-linear Bound**

$$C_i = 0.5 * \log(|X_{i1}| + 1) - 0.2 * I(X_{i2} > 0) + 0.2 * \sin(\pi * X_{i3})$$

$$D_i = 0.3 * \sin(\pi * X_{i2}) - 0.3 * I(X_{i3} > 0) + 0.3 * |X_{i4}|$$

$$A_i \sim \mathcal{N}(C_i, 0.5)$$

$$\mu_i = \frac{re^{rA_i}}{e^{rC_i} + e^{rA_i}} + D_i \quad r = 10$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 1)$$

$$S_i = fitted(lm(Y \sim X))_i$$

3. **Unimodal with Linear Bounds**

$$C_i = 0.3 * X_{i1} + 0.3 * X_{i2} + 0.3 * X_{i3}$$

$$D_i = 0.6 * X_{i2} + 0.6 * X_{i3} + 0.6 * X_{i4}$$

$$A_i \sim \mathcal{N}(C_i, 0.5)$$

$$\mu_i = 1 - |A_i - C_i| + D_i$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 0.2)$$

$$S_i = fitted(lm(Y \sim X))_i$$

4. **Unimodal with Non-linear Bounds**

$$C_i = 0.5 * \log(|X_{i1}| + 1) - 0.2 * I(X_{i2} > 0) + 0.2 * \sin(\pi * X_{i3})$$

$$D_i = 0.3 * \sin(\pi * X_{i3}) - 0.3 * I(X_{i4} > 0) + 0.3 * |X_{i5}|$$

$$A_i \sim \mathcal{N}(C_i, 0.5)$$

$$\mu_i = 1 - |A_i - C_i| + D_i$$

$$Y_i = \mu_i + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, 0.2)$$

$$S_i = fitted(lm(Y \sim X + X^2 + X^3))_i$$

From Figure 6.5 and Figure 6.6, the OWL-L has advantages in the two linear settings and the OWL-K has advantages in the two nonlinear settings, which is consistent with the results in Section 6.1. This fact reassures that the proposed OWL-based method will not be heavily impacted by a reasonable estimation of the threshold $S(x)$. We believe that the estimation of the confidence probability $\alpha(x)$ shares the same property. In the next section, the above IDI methods will be examined on the real dataset where the function $S(x)$ and $\alpha(x)$ are not available and thus have to be estimated from the data.

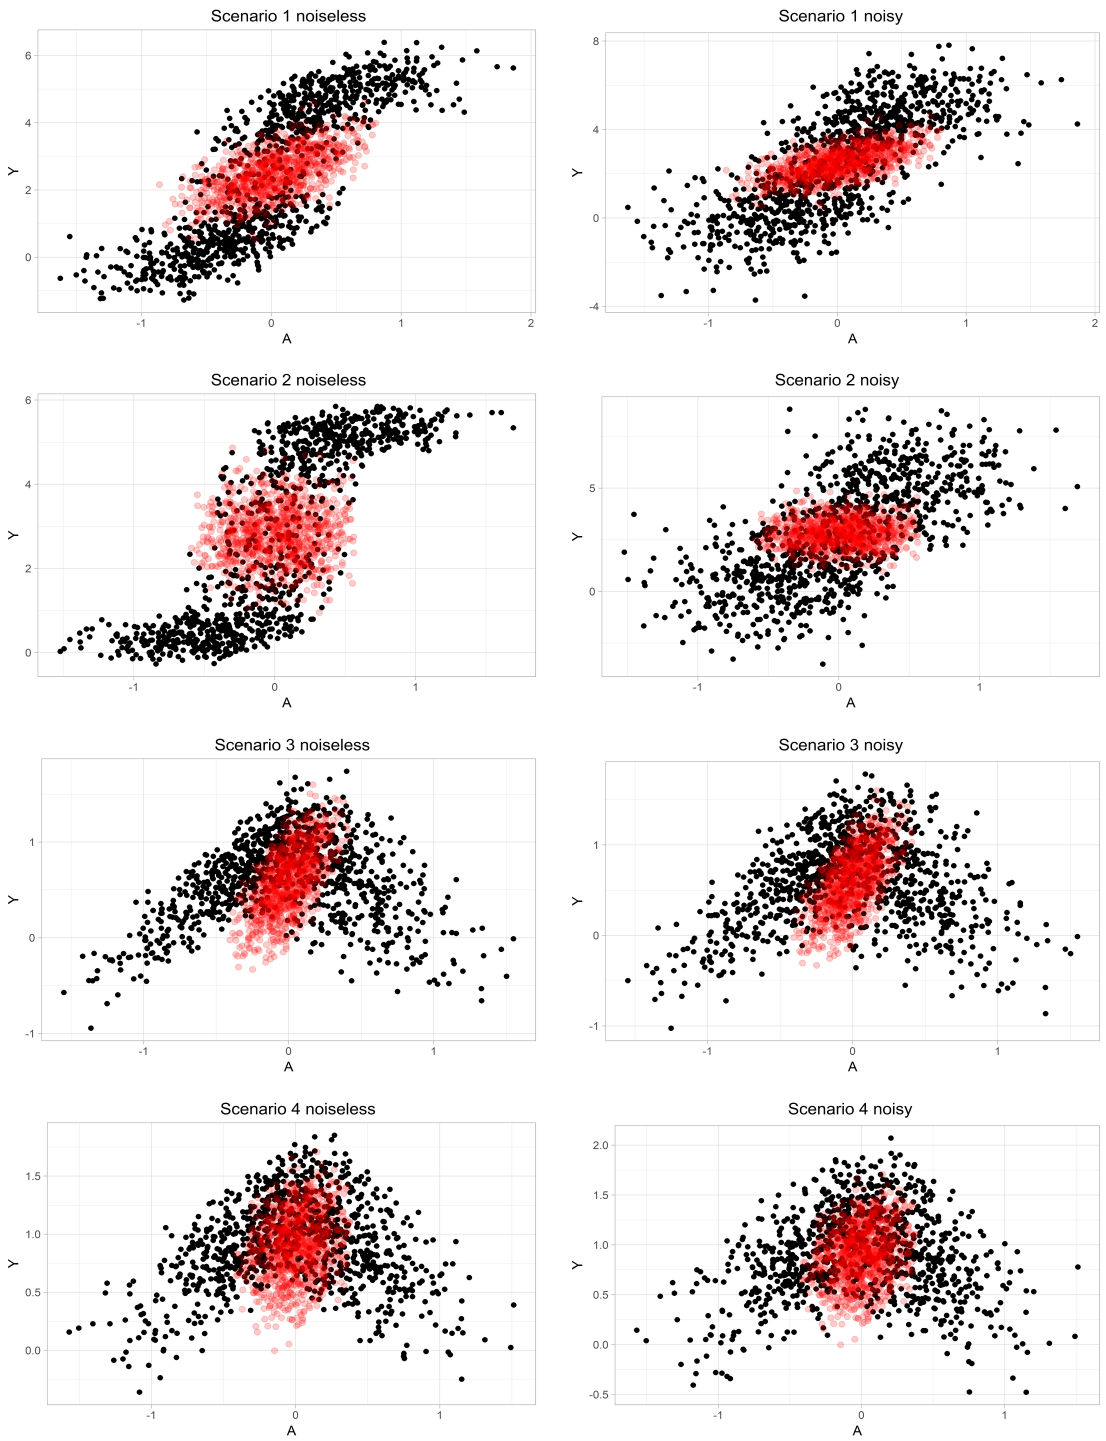

Figure 6.4: Scatter Plots of Simulation Scenarios with Prognostic Effects

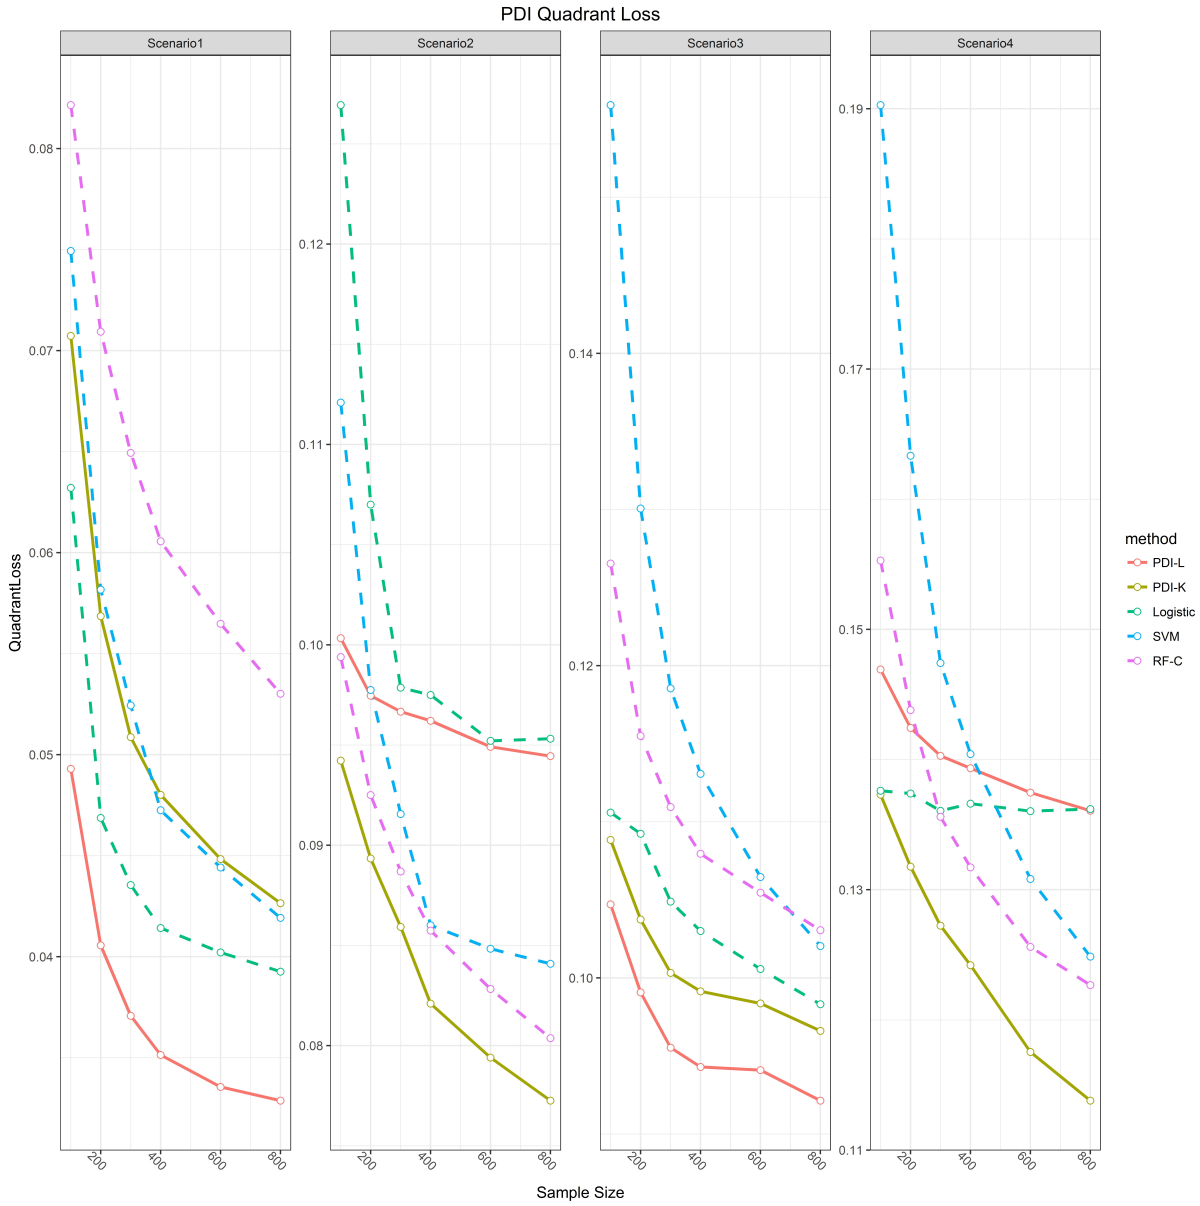

Figure 6.5: PDI Average Quadrant Losses for Simulation Settings in Section 6.2

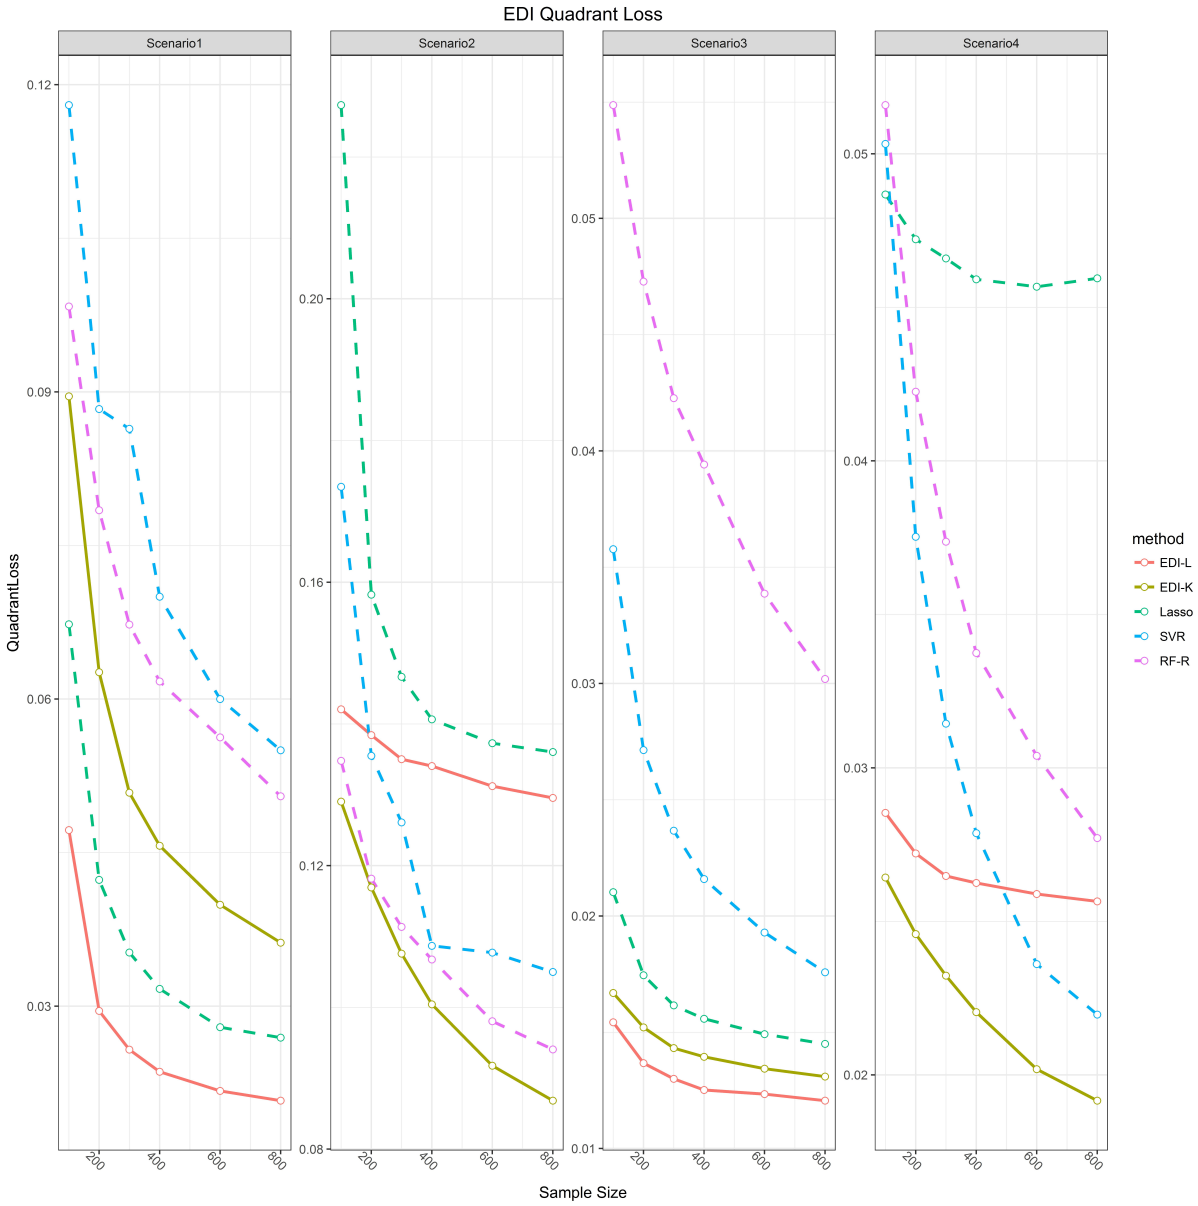

Figure 6.6: EDI Average Quadrant Losses for Simulation Settings in Section 6.2

	N	PDI-L	PDI-K	Logistic	SVM	RF-C
Scenario 1	100	**0.049 (0.010)**	0.071 (0.010)	0.063 (0.030)	0.075 (0.012)	0.082 (0.009)
	200	**0.041 (0.005)**	0.057 (0.006)	0.047 (0.007)	0.058 (0.006)	0.071 (0.006)
	300	**0.037 (0.004)**	0.051 (0.004)	0.044 (0.005)	0.052 (0.005)	0.065 (0.005)
	400	**0.035 (0.003)**	0.048 (0.004)	0.041 (0.004)	0.047 (0.004)	0.061 (0.004)
	600	**0.034 (0.002)**	0.045 (0.003)	0.040 (0.004)	0.044 (0.003)	0.056 (0.004)
	800	**0.033 (0.002)**	0.043 (0.003)	0.039 (0.003)	0.042 (0.003)	0.053 (0.003)
Scenario 2	100	0.100 (0.005)	**0.094 (0.004)**	0.127 (0.041)	0.112 (0.018)	0.099 (0.007)
	200	0.097 (0.004)	**0.089 (0.005)**	0.107 (0.026)	0.098 (0.018)	0.093 (0.005)
	300	0.097 (0.003)	**0.086 (0.004)**	0.098 (0.007)	0.092 (0.011)	0.089 (0.004)
	400	0.096 (0.003)	**0.082 (0.004)**	0.098 (0.010)	0.086 (0.008)	0.086 (0.003)
	600	0.095 (0.002)	**0.079 (0.004)**	0.095 (0.002)	0.085 (0.015)	0.083 (0.003)
	800	0.094 (0.002)	**0.077 (0.003)**	0.095 (0.004)	0.084 (0.020)	0.080 (0.003)
Scenario 3	100	**0.105 (0.007)**	0.109 (0.007)	0.111 (0.005)	0.156 (0.017)	0.127 (0.010)
	200	**0.099 (0.006)**	0.104 (0.005)	0.109 (0.005)	0.130 (0.007)	0.116 (0.006)
	300	**0.096 (0.003)**	0.100 (0.003)	0.105 (0.005)	0.119 (0.006)	0.111 (0.004)
	400	**0.094 (0.003)**	0.099 (0.003)	0.103 (0.006)	0.113 (0.005)	0.108 (0.003)
	600	**0.094 (0.003)**	0.098 (0.003)	0.101 (0.005)	0.106 (0.004)	0.105 (0.004)
	800	**0.092 (0.003)**	0.097 (0.003)	0.098 (0.004)	0.102 (0.004)	0.103 (0.003)
Scenario 4	100	0.147 (0.008)	**0.137 (0.005**)	0.138 (0.004)	0.190 (0.022)	0.155 (0.012)
	200	0.142 (0.005)	**0.132 (0.005)**	0.137 (0.003)	0.163 (0.010)	0.144 (0.007)
	300	0.140 (0.005)	**0.127 (0.006)**	0.136 (0.002)	0.147 (0.006)	0.136 (0.005)
	400	0.139 (0.004)	**0.124 (0.006)**	0.137 (0.002)	0.140 (0.006)	0.132 (0.005)
	600	0.137 (0.004)	**0.118 (0.004)**	0.136 (0.003)	0.131 (0.005)	0.126 (0.004)
	800	0.136 (0.002)	**0.114 (0.002)**	0.136 (0.002)	0.125 (0.004)	0.123 (0.004)

Table 6.3: PDI Quadrant Losses for Simulation Settings in Section 6.2

	N	EDI-L	EDI-K	Lasso	SVR	RF-R
Scenario 1	100	**0.0472 (0.0176)**	0.0896 (0.0168)	0.0673 (0.0279)	0.1180 (0.0535)	0.0983 (0.0162)
	200	**0.0295 (0.0074)**	0.0626 (0.0122)	0.0423 (0.0129)	0.0883 (0.0471)	0.0784 (0.0115)
	300	**0.0258 (0.0047)**	0.0509 (0.0075)	0.0353 (0.0090)	0.0864 (0.0636)	0.0673 (0.0092)
	400	**0.0236 (0.0036)**	0.0457 (0.0066)	0.0317 (0.0071)	0.1185 (0.0933)	0.0617 (0.0069)
	600	**0.0217 (0.0021)**	0.0399 (0.0043)	0.0280 (0.0047)	0.1769 (0.1159)	0.0563 (0.0062)
	800	**0.0208 (0.0022)**	0.0362 (0.0038)	0.0269 (0.0041)	0.2231 (0.1169)	0.0505 (0.0055)
Scenario 2	100	0.1421 (0.0097)	**0.1290 (0.0087)**	0.2273 (0.0971)	0.1735 (0.0441)	0.1348 (0.0125)
	200	0.1384 (0.0072)	**0.1169 (0.0118)**	0.1582 (0.0252)	0.1355 (0.0313)	0.1182 (0.0095)
	300	0.1350 (0.0062)	**0.1076 (0.0099)**	0.1467 (0.0126)	0.1261 (0.0384)	0.1114 (0.0090)
	400	0.1341 (0.0052)	**0.1004 (0.0105)**	0.1407 (0.0083)	0.1087 (0.0232)	0.1068 (0.0082)
	600	0.1312 (0.0042)	**0.0918 (0.0069)**	0.1373 (0.0085)	0.1077 (0.0345)	0.0981 (0.0066)
	800	0.1296 (0.0041)	**0.0868 (0.0059)**	0.1360 (0.0063)	0.1139 (0.0432)	0.0941 (0.0056)
Scenario 3	100	**0.0154 (0.0019)**	0.0167 (0.0018)	0.0210 (0.0047)	0.0358 (0.0042)	0.0549 (0.0034)
	200	**0.0137 (0.0012)**	0.0152 (0.0012)	0.0175 (0.0032)	0.0271 (0.0026)	0.0473 (0.0028)
	300	**0.0130 (0.0010)**	0.0143 (0.0010)	0.0162 (0.0018)	0.0237 (0.0022)	0.0423 (0.0032)
	400	**0.0125 (0.0006)**	0.0139 (0.0009)	0.0156 (0.0016)	0.0216 (0.0022)	0.0394 (0.0026)
	600	**0.0123 (0.0006)**	0.0134 (0.0007)	0.0149 (0.0009)	0.0193 (0.0015)	0.0339 (0.0021)
	800	**0.0121 (0.0005)**	0.0131 (0.0007)	0.0145 (0.0009)	0.0176 (0.0015)	0.0302 (0.0024)
Scenario 4	100	0.0285 (0.0023)	**0.0264 (0.0017)**	0.0487 (0.0034)	0.0503 (0.0061)	0.0516 (0.0037)
	200	0.0272 (0.0016)	**0.0246 (0.0016)**	0.0472 (0.0024)	0.0375 (0.0043)	0.0423 (0.0032)
	300	0.0265 (0.0012)	**0.0232 (0.0017)**	0.0466 (0.0021)	0.0314 (0.0033)	0.0374 (0.0026)
	400	0.0263 (0.0011)	**0.0220 (0.0017)**	0.0459 (0.0020)	0.0279 (0.0020)	0.0337 (0.0022)
	600	0.0259 (0.0010)	**0.0202 (0.0013)**	0.0457 (0.0015)	0.0236 (0.0015)	0.0304 (0.0017)
	800	0.0257 (0.0005)	**0.0192 (0.0005)**	0.0459 (0.0014)	0.0220 (0.0015)	0.0277 (0.0009)

Table 6.4: EDI Quadrant Losses for Simulation Settings in Section 6.2

Chapter 7

Study of Hemoglobin A1c Control and Diabetic Associated Events

Hemoglobin is an iron-containing oxygen which transports protein in the red blood cells. Hemoglobin A1c (HbA1c), the most abundant minor hemoglobin in the human body, is formed when glucose accumulates in red blood cells and binds to the hemoglobin. This process occurs slowly and continuously over the lifespan of red blood cells, which is 120 days on average. This makes HbA1c an ideal biomarker of long-term glycemic control. Patients who are susceptible to high blood glucose are usually recommended to have their HbA1c levels regularly measured and recorded.

HbA1c test, since became commercially available in 1978, has been one of the standard tools for monitoring diabetes progression. American Diabetes Association (ADA) recommended using A1c measurement in 1988. The Diabetes Control and Complications Trial (DCCT) demonstrated its importance as a predictor of diabetes-related outcomes, and the ADA started recommending specific A1c targets in 1994 (Little et al., 2011). In 2010, ADA added HbA1c$\geq$ 6.5% as a criterion for diabetes diagnosis. The prevalence of HbA1c test can be partially explained by its clinic convenience. As HbA1c is unaffected by acute perturbations in glucose levels, there is no need for fasting or timed samples. For people without diabetes, the normal range for the HbA1c level is between 4% and 5.6%. Hemoglobin A1c

levels between 5.7% and 6.4% imply a higher chance of getting diabetes. Levels of 6.5% or higher indicate the presence of diabetes.

Despite the popularity of HbA1c tests, HbA1c may not reflect the real progression of diabetes given various conditions that patients are predisposed to. Interfering factors may lead to false results of HbA1c tests (Radin, 2014). In some cases, the direction of bias is predictable while in other cases not. For instance, HbA1c can be falsely elevated by iron deficiency anemias, vitamin B-12 anemias, folate deficiency anemias, asplenia and other conditions associated with decreased red cell turnover (Nitin, 2010). HbA1c measures may be lower for patients with conditions that shorten the life of the red blood cell or increase its turnover rate (Freedman et al., 2010), including acute and chronic blood loss, hemolytic anemia, splenomegaly and end-stage renal disease. While in some other cases, the direction of bias is not always easily predictable. For example, the complex interplay of glycemic control and other treatments, such as erythropoietin therapy and the treatment of uremia may influence A1c level in a more case-by-case way (Radin, 2014).

Due to the susceptibility of HbA1c measures to various medical and physiological factors, HbA1c measures reflects glycemia differently for different patients. Therefore, the recommendation of a single HbA1c control level for all the patients may not always be desirable. As mentioned in Section 1.1, a lower glycemic level is not always better, and over-controlling glycemia can lead to faint or compromised life quality. Especially for patients who are older, with severe complications, or in the phase of early post-surgery rehabilitation, HbA1c is not supposed to be controlled for them as low as for younger and healthier patients.

As a response to the progress in medical research, the idea of glucose control has been evolving over time. In the past, an A1C of 7 percent was considered the golden rule of health for everyone. In recent years, however, the importance of a patient-centered approach to managing A1C levels has been recognized, which may better correspond to the patient's needs for diabetes management and their personal conditions and preferences.

The evolving needs of medical practices call for a personalized HbA1c control strategy with individualized intervals based on patients characteristics. In this Chapter, we demon-

strate how the methods that we proposed in the dissertation are ideal for the application of HbA1c control. Specifically, we apply our OWL-based approach, as well as the indirect methods based on classification, to find the IDIs for HbA1c control. The following of the Chapter is organized as follows. In Section 7.1, we introduce the EHR dataset and illustrate the data construction in details. The preparation of the estimation, including stratification of the sample, the personalization of confidence probability, and the estimation of inverse density weighting, are discussed in Section 7.2. In Section 7.3, we show that the estimated HbA1c upper bounds and compare it with the observed HbA1c levels in the dataset. In Section 7.4, we discuss the effect of IDI recommendation if patients follow the suggested HbA1c control intervals.

7.1 Data Source

In this section, we introduce the HbA1c dataset to study the relationship between HbA1c control and the events of diabetic patients. Specifically, the sample is a subset of a type II diabetes dataset which is drawn from a large Midwestern multi-specialty physician and patient group. Electronic Health Record (EHR) datasets are merged with Medicare claim data to obtain laboratory records, information on primary care visits, and medication history. All patients have been tracked for a specific time period from the first quarter of 2003 to the fourth quarter of 2011. The full observational records are not available for the majority of patients. All the patients have been medically homed to receive the plurality of primary care and/or endocrinology care at primary care clinics managed by UW Medical Foundation, UW Hospital and Clinics and UW Family Medicine.

Patients identified with diabetes must have at least 1 inpatient or SNF Medicare claim or at least two carrier claims for diabetes, which should be less than 2 years apart. The diagnosis should be defined by one or more of the following categories: the ICD-9 codes, Diabetes mellitus (250.xx), Polyneuropathy in diabetes (357.2), Diabetes retinopathy (362.0x), and Diabetic cataract (366.41). There are several types of patients who are excluded from

the dataset, who are without continuous coverage with Medicare Parts A & B for a baseline year or for at least one subsequent quarter, without Medicare railroad benefits or were not enrolled in a Medicare HMO. Each patient needs to have at least 5 quarter time periods in one of the 3 insurers from 2003 to 2011 including 4 baseline quarters and a measurement period of at least 1 subsequent quarter. Also, the patients should have A1c records over the 5 quarter period which are available in the UW health system. In the sample of analysis, 51.3% of patients are female and 91.8% of the patients are White. The ages of participants at the cohort start time range from 18 to 102 years old, with an average of 65.7 and a standard deviation of 13.7. There are 8645 patients in total whose 95130 HbA1c measurements are recorded.

Multiple patients' observations need to be aggregated to learn a single IDIR, so it is generally preferable that the patients are in similar conditions. In this data, there are 1884 patients whose maximal HbA1c are above 10%. We believe that these patients have been in severe conditions for a prolonged duration, and their events might not reflect the progress of diabetes and may be confounded by other long-term comorbidities. Therefore, only patients with maximal HbA1c below 10% are included.

We regard the HbA1c measurements as the treatment doses in this study. Observations are centered around the availability of HbA1c measurements. Each observation is associated with a unique HbA1c measurement from an individual patient proceeding the baseline period. The outcome variable is the number of negative events, which include hospitalizations and Emergency Department (ED) visits that cover 90 days following the HbA1c measurement, 90 days after the HbA1c measurement. For the purpose of notation consistency, we take the negative of the number of events such that for any $S \in (-1, 0)$, $1\{Y_i > S\}$ represents a good outcome and $1\{Y_i < S\}$ stands for a bad outcome. In addition, we also include a list of short-term covariates that we average from measurements of conditions 90 days prior to the treatment HbA1c measurement as well as a list of long-term covariates that are averaged from conditions up to 900 days prior to the treatment HbA1c measurement. After removing the ones with missing baseline covariates, 47432 observations

remain. An illustration of the data structure can be found in Figure 7.1. Each observation is generated from an outcome period (90 days after HbA1c measure), and short-term covariates period (90 days before the HbA1c measure) and a long-term covariates period (including the baseline information). Potential overlapping is allowed.

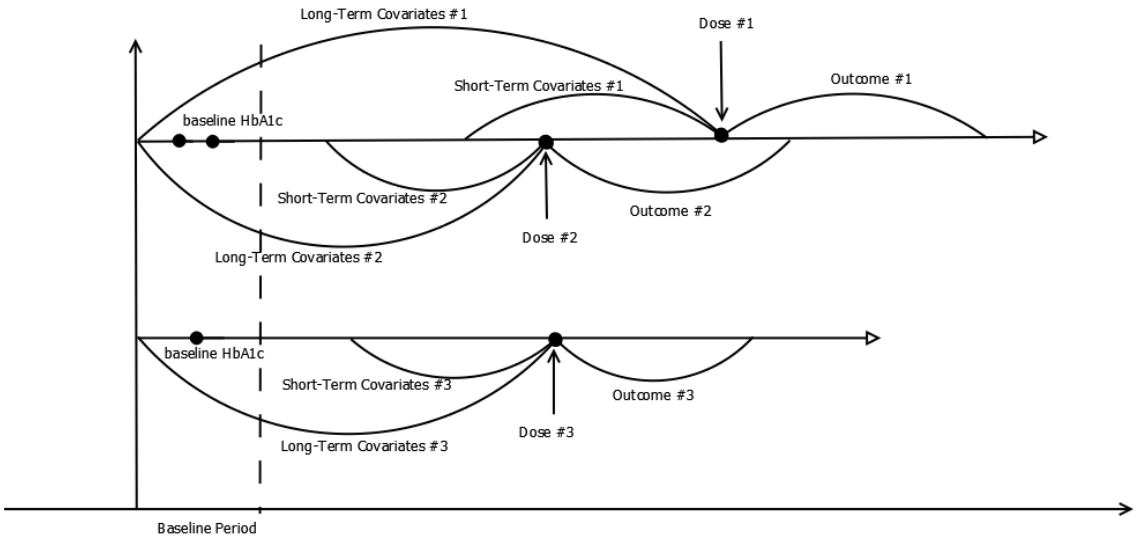

Figure 7.1: Illustration of Data Structure

There are several reasons why we choose to select the observations based on treatment HbA1c measurements rather than on patients or patient-quarters. First, this study aims to provide dynamic Individual Dose Interval (IDI) recommendations for patients. The IDI shall change over time in order to capture the effect of aging and disease progression along the study period. Therefore, it is not applicable to have observations constructed at the patient level. Second, although the HbA1c measurement generally reflects the level of blood sugar in the past 3 months, it may be impacted by recent events (Radin, 2014). For instance, a recent hospitalization or ED-visit tend to increase the following HbA1c measurements to a large extent. A patient-quarter data structure, with aggregation of treatment and events throughout a quarter, cannot really exclude the outcomes prior to the treatment and eventually may overestimate the impact of blood sugar. However, with the HbA1c-centered data structure, we are able to condition on all the events occurred prior to the HbA1c measurement as covariates, thus adjusting for these confounders. Such observations are

referred to as the new 'patient's for convenience.

The dataset also contains a large range of variables, among which are variables delineating diabetes progress, those defining the IDIR and confounders of the treatment. Specifically, the covariates may be collapsed to the following three categories:

1. Patient Demographics: age, gender, race

2. Patient Long-term Measurement: height, the indicator of disability at baseline, length of enrollment, the average of HbA1c in past 91-900 days, the indicator of Medicaid, number of measurements recorded in past 91-900 days, Sulfonylureas intake, Insulin intake, and other medications intake

3. Patient Short-term Measurement: number of HbA1c measurements in the past 90 days; the average of HbA1c measurement in the past 1-90 days, number of hospitalization, number of ED-visits; low-density lipoprotein cholesterol, Systolic blood pressure, Diastolic blood pressure, BMI, weight, cardiovascular disease, ischemic heart disease, Congestive Heart failure, Hepatocellular Carcinoma, Hypoglycemia and injury, infections; and 18 other comorbidities;

7.2 Data Preparation

It is important for defining an IDI that the threshold function $S(x)$ and the probability function $\alpha(x)$ are specified reasonably so that the interpretations are meaningful and the assumptions are realistic. First, a proper choice of $S(x)$ might be any real value between 0 and -1. This is because the 95% of patients have less than 2 events, which means the majority of the patients have outcomes are greater than -2. If we choose any value smaller than -1, it might induce severe imbalance that $1\{Y_i < S\}$ and $(Y_i < S)_+$ will be zero. Besides, the presence of events, corresponding to the outcome being zero, is a good indicator of patients' health condition. Therefore, we recode our integer outcome as a binary variable, $Y = 0$ when the patient did not have any event and $Y = -1$ when the patient had at least one

event.

Second, it might be undesirable to treat all the patients in the same manner as one patient's conditions tend to be dramatically different from the other's. As mentioned in Chapter 2.4, patients with a long diabetic history and multiple comorbidities may go to the hospital frequently even when their blood sugar level is well-controlled. In contrast, younger and healthier patients might be able to endure much higher blood sugar without hospitalizations and emergency visits. In addition, the variation of patients' conditions across time has to be considered. It will not be a good choice for a patient in the early post-surgery rehabilitation stage to control blood sugar as low as before. As is discussed in Chapter 2.4, for the binary outcome $Y \in \{-1, 0\}$, personalizing $\alpha(x)$ a the probability of having no events, $P(Y_i = 0|X_i)$, and having $S(x) \in (-1, 0)$ for the PDI is equivalent to personalizing $S(x)$ as $P(Y_i = 0|X_i)$ and letting $\alpha(x) = 0.5$ for the EDI. To simplify notations, we use the PDI notations and let $\alpha(x) = \hat{P}(Y_i = 0|X_i)$ and $S(x) = -0.5$.

In this study, we use a linear logistic regression to model $P(Y_i = 0 \,|\, X_i)$ without the information of HbA1c measure, i.e., $logit(P(Y_i = 0 \,|\, X_i)) = X_i\beta$. For each new patient with features x^*, the predicted probability of having no event is $1/\left(1 + e^{-x^*\hat{\beta}}\right)$ assuming that the dosing practice remains unchanged for the new patient, which means the dose received by the new patient is obtained from the same conditional distribution $P(a \,|\, x)$ as in the training dataset. The model contains the following variables as covariates: age, gender, BMI, short-term average HbA1c, long-term average HbA1c, the indicator of ischemic heart disease or congestive heart failure, the indicator of hepatocellular carcinoma, short-term number of hospitalization, short-term number of ED-visit, and the indicator of disability. These variables are believed to capture the most of prognostic effects. Adding extra variables is likely to reduce the stability of the estimates.

As patients at the different risk levels are not likely to have the same IDIR, separate analyses for different risk groups might be necessary. Hence, we divide the entire sample into three equal stratas according to the estimated probability of no event $\hat{P}(Y_i = 0 \,|\, X_i)$ with the top 5% and bottom 5% extremes excluded. Figure 7.2 shows a density plot

of $\hat{P}(Y_i = 0 \,|\, X_i)$ for three stratas. Analyses are conducted for each stratas. For each individual X_i, we let $\alpha(x)$ be the 75% quantile of the estimated probabilities of no event in the strata where $\hat{P}(Y_i = 0|X_i)$ falls in. For each strata, 1000 patients are randomly sampled to training dataset while the remaining around 2000 patients are left for the testing dataset. We average the results from 100 times of repetitions of the whole procedure.

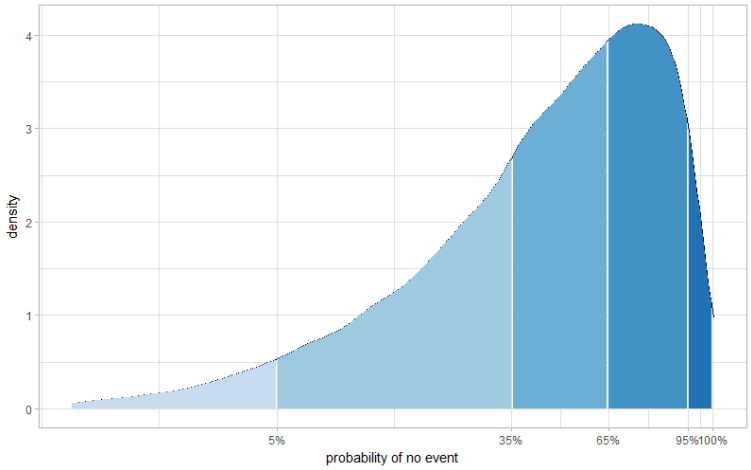

Figure 7.2: Stratification Based on the Estimated Probabilities of No Event

For each repetition, an estimate of the conditional density $P(A_i|X_i)$ is required to construct the inverse density weighting in our OWL-based method. We first uniformly split the range between 4% and 10% into 20 equal levels and use cumulative logistic regression (Harrell, 2001) to model the ordinal HbA1c levels. The probability of observing one level over the other is denoted as $P(A_i|X_i)$. We did not model the density of residuals from a regression fitting because the conditional distribution of HbA1c is likely to be different for patients with different conditions, and the noisiness of HbA1c measurement may lead to large residuals resulting in extreme probabilities. Also, multinomial regression approach may not be suitable in this case which may include too many parameters and may result in unstable estimation. However, the OWL framework is robust to minor misspecification of the propensity score model (Zhao et al., 2012, Chen et al., 2016), so we choose it as the most conservative available strategy. Our numerical results in the next two section also validate this argument.

7.3 Distribution of the Estimated HbA1c Control Upper Bounds

The observed HbA1c measures are capped between 4% and 10%, after having the patients with maximal HbA1c above 10 excluded. The distribution of our estimated upper bounds using linear Kernel for the 3 stratas can be found in Figure 7.3, Figure 7.4 and Figure 7.5. For the first strata, while the majority of patients have observed HbA1c between 5.5% and 7.5%, their estimated upper bounds are generally between 6.5% and 8%, which suggests that most of the patients are over-controlling their blood sugar. If some of them lose the control by 0.5% or 1%, their event risk is still below 18.5%. In fact, around 80% of the patients have HbA1c below the suggested upper bounds while the around 20% of them failed to control their blood sugar. Besides, 75% of patients have interval width below 1% and the median width is 0.6%. The second strata and the third strata generally have the similar distribution of the upper bounds except for a slight shift to the right.

The indirect methods give IDI recommendations that are dramatically different from those from the OWL-based method. First, as we have mentioned in Section 1.4 and Section 2.4, the indirect methods can predict if patients will have desirable outcomes by simply modeling the prognostic effects. Consequently, only a small proportion of patients, usually smaller than 50%, have informative HbA1c upper bounds. Figure 7.3 shows in the right column the distribution of the informative upper bounds for the Logistic Regression, the SVM and the Random Forest in sequence. In the left column, sample dose-response curves of the corresponding methods are presented by plotting the predicted probability of no event against the range of HbA1c measurement. It can be inferred from the curves and the density plots that the Logistic Regression has the least flexibility and assigns approximately the same upper bound to all the patients. As a contrast, the Random Forest produces the most flexible curves but results in over-spread upper bounds. The SVM is generally between the other two approaches in terms of the flexibility, but its upper bounds are highly tilted towards the right, which might not be realistic for most the diabetic patients.

The 2015 American Diabetes Association (ADA) Standards of Medical Care in Diabetes

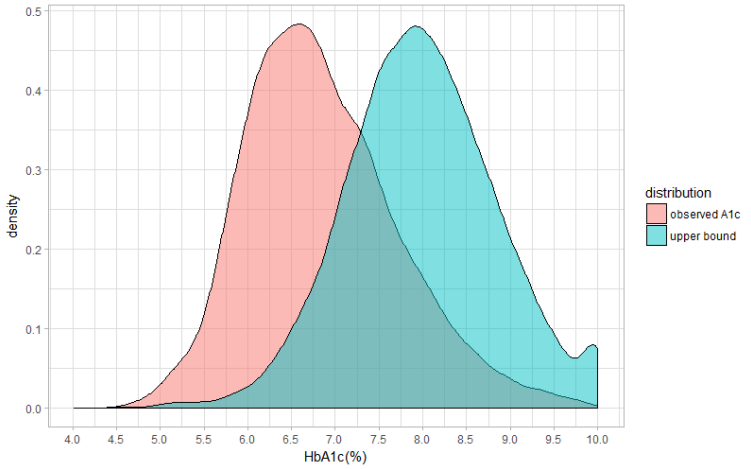

Figure 7.3: Histogram of Observed HbA1c and Estimated Upper Bounds for the 1st Strata

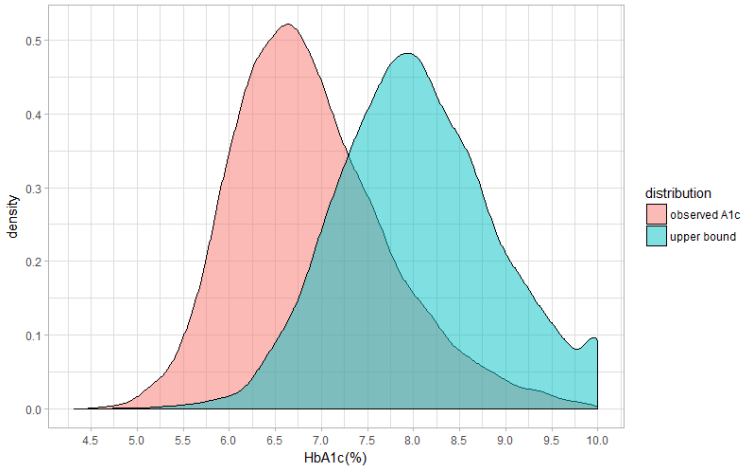

Figure 7.4: Histogram of Observed HbA1c and Estimated Upper Bounds for the 2nd Strata

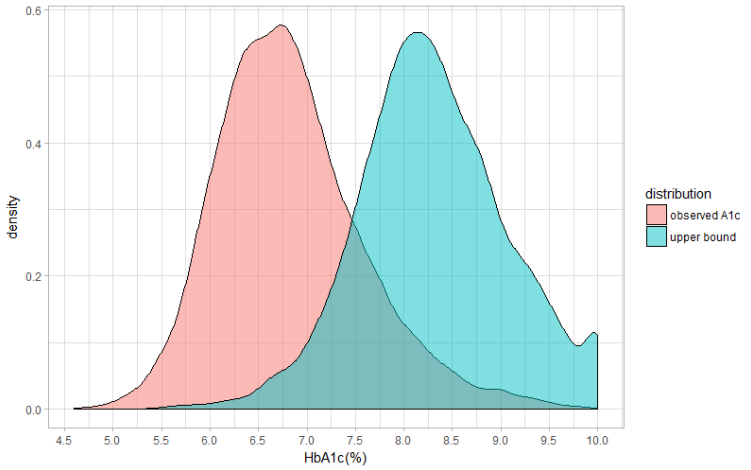

Figure 7.5: Histogram of Observed HbA1c and Estimated Upper Bounds for the 3rd Strata

advise the following A1C levels. 6.5 % is the stringent goal for people who can achieve this goal without experiencing a lot of hypoglycemia episodes or other negative effects of having lower blood glucose levels. 7 % is a reasonable A1C goal for many adults with diabetes who are not pregnant. 7.5 % is the goal recommended for all children with diabetes. 8 % is the less stringent goal for people with a history of severe hypoglycemia, people who have had diabetes for many years or have difficulty achieving tighter control, and people with limited remaining life expectancy.

The upper bounds we provided might look too high at the first glance compared with the ADA standard. However, they are actually coherent with the standard given the age distribution. In our dataset, 88% of the patients are above 60 years old, 50% of the patients are over 73 years old and 20% are aged above 80 years old. Most of the patients have diabetes for many years. Generally, if taking prescriptions from a practicing physician, most of the patients will be targeting the 8% goal in general. Our personalized HbA1c control intervals are expected to capture more heterogeneity among patients and therefore may provide more informative guidance for physicians.

7.4 Average Causal Effects of the Upper Bound Recommendations

As illustrated in the previous section, both our OWL-based methods and the indirect methods can provide estimates for the upper bounds of HbA1c for observations in the testing dataset. Due to the reasons stated in Section 3.3, it is not reasonable to interpret the gap between the shares of events of patients with doses inside IDIs and patients with doses outside IDI as the effect of IDI recommendation. First, the distribution of estimated upper bounds varies across different methods. The estimated bounds from OWL-based methods are generally concentrated between 7.5% and 8.5%. For the indirect methods, the estimated upper boundaries are more spread-out. Second, with the indirect methods, we can in addition model the prognostic effect within each strata. This enables the identification

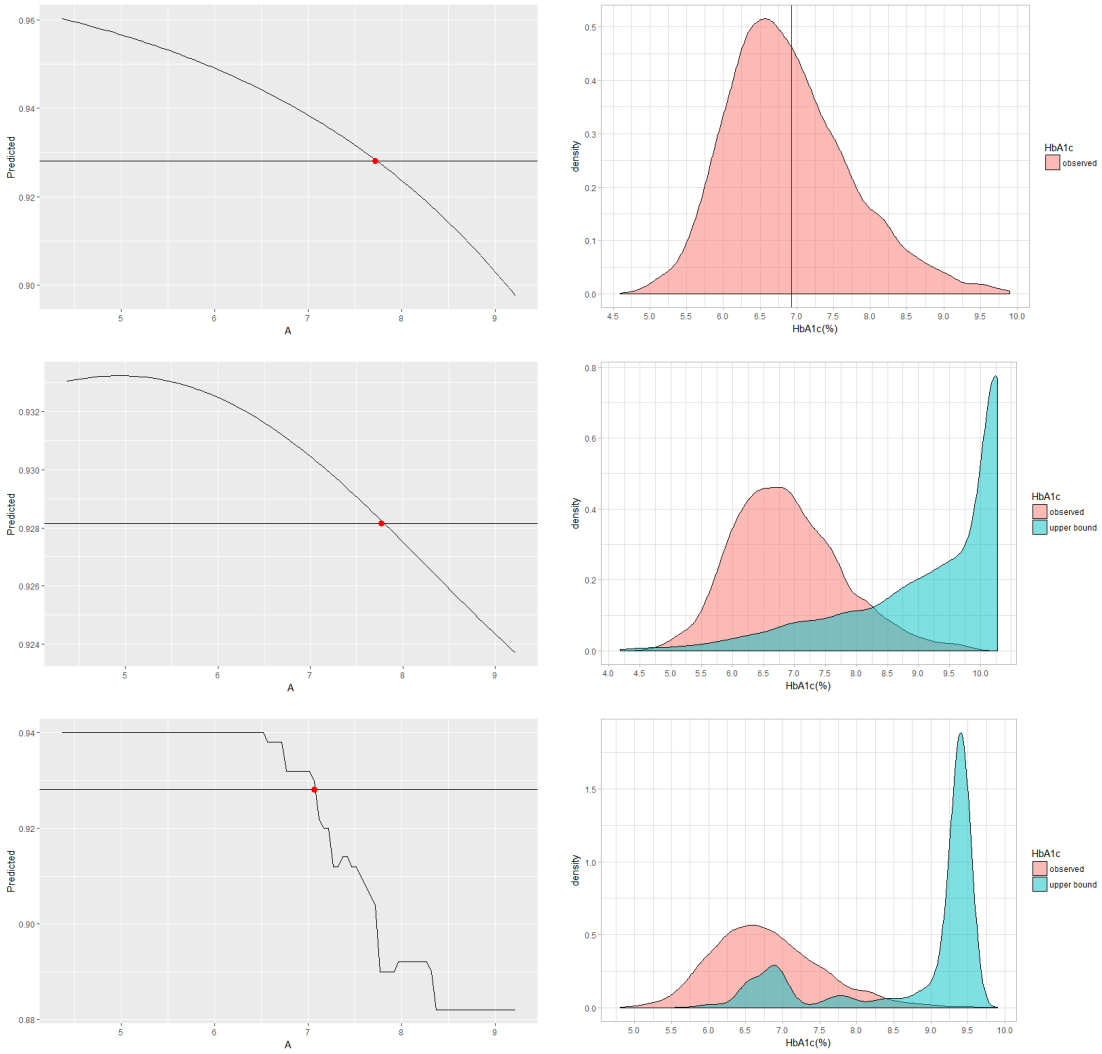

Figure 7.6: Fitted Outcome Trajectories and Histograms of the Estimated HbA1c Upper Bounds from Logistic Regression, SVM, and Random Forest

of the proportion of patients who can endure any reasonable HbA1c and the proportion of patients whose events cannot be avoided. As a consequence, patients who are located inside the IDIs estimated from indirect methods have much lower risk compared to their counterpart patients outside the IDIs. Therefore, a direct comparison might overestimate the effectiveness of the IDI recommendations estimated from the indirect methods.

To evaluate the efficacy of all the methods, we evaluate the ATE defined in Section 3.3. Patients who had HbA1c measurement below the estimated IDIs are regarded as takers. Patients failed to control the HbA1c below the estimated IDIs are regarded as non-takers. We build a CBPS model using the indicator $T_i = 1\{A_i \in [a_L, f_U(X_i)]\}$ as outcome against all the other covariates X_i. In the pseudo-population weighted by $\hat{P}\big(T_i\big)/\hat{P}\big(T_i\big|X_i\big)$, the standardized differences of covariates between the taker group and the non-taker group are generally below 0.1, which indicates an acceptable balance of covariates. The ATE in our case is defined as the gap in the share of those having events between the weighted taker group and weighted non-taker group.

$$ATE = \sum_{i=1}^{n} \frac{1\{Y_i < 0\}(1 - T_i)}{P(T_i|X_i)} - \sum_{i=1}^{n} \frac{1\{Y_i < 0\}T_i}{P(T_i|X_i)}$$

In Table 7.4, the estimated average causal effects with the OWL-based methods and indirect methods are summarized. It is clear that the OWL-based approach with linear Kernel yields the largest ATEs among all these cases. Besides, it is apparent in the within-strata analysis that the difficulty of achieving the threshold decrease from the 1st strata to the 3rd strata. Because most of the methods include a larger proportion of patients into the interval, while the corresponding ATEs decrease as the number of stratas goes up.

Combined with the results from Section 7.3, the causal effects in Table 7.4 confirm the validity of using our OWL-based methods as a new approach to study the HbA1c control interval because the distribution of the estimated HbA1c upper bounds is reasonable given the current researches and the causal effect of OWL-L recommendations is larger compared to that of the indirect methods.

	Strata 1		Strata 2		Strata 3	
	%taker	ATE	%taker	ATE	%inside	ATE
Logistic	20.97 (19.45)	3.96 (0.53)	35.98 (22.64)	1.16 (0.12)	22.70 (15.89)	0.78 (0.14)
SVM	17.94 (14.63)	1.44 (0.07)	30.34 (22.00)	1.74 (0.07)	60.93 (21.00)	1.76 (0.09)
RF	33.03 (5.36)	2.86 (0.06)	45.68 (5.70)	1.67 (0.04)	58.79 (7.10)	1.31 (0.04)
OWL-L	90.00 (3.29)	**7.42** (0.34)	91.67 (2.38)	**3.28** (0.40)	93.13 (4.00)	**2.73** (0.17)
OWL-K	88.80 (4.96)	4.54 (0.21)	90.52 (3.23)	2.51 (0.19)	89.32 (5.46)	1.63 (0.10)

Table 7.1: The Percentage of Takers and the ATE of Recommended HbA1c Upper Bounds

Chapter 8

Future Work

In this dissertation, we consider the problem of Individual Dose Interval (IDI) recommendation. Through the discussion of the indirect methods based on traditional outcome modeling and grid search, we develop a new framework based on Outcome-Weighted-Learning (OWL). We believe our method has several advantages. First, it is a unified framework for both Probability Dose Interval (PDI) and Expectation Dose Interval (EDI). Second, the modeling of the outcome, known to be susceptible to misspecification, is avoided in our framework. Third, in contrast to the indirect methods, our method provides an informative IDI for each patient since it is less affected by the prognostic effect. Simulation studies show that our method has a small sample advantage over the indirect methods. In a study of the HbA1c control upper bound, our method is proven to have better robustness and produces recommendations of upper bounds that are more beneficial to the patients compared to the indirect methods.

We hope that more attention will be drawn to the optimal dose interval problem. There are several interesting open topics that deserve more discussion in the future. First, the statistical inference of IDI estimation has not been established. There are still major obstacles due to the non-convex objective function. Second, faster convergence rates for Kernel estimation can be derived, especially for the two-sided IDIs. Third, high dimensional techniques might be integrated with this method to broaden the applications. Finally, an extension of

this framework to dynamic treatment regime with multiple stages might be a challenging and rewarding direction.

Reference

Breiman, L. (2001), 'Random forests', *Machine learning* **45**(1), 5–32.

Breiman, L. (2017), *Classification and regression trees*, Routledge.

Cai, T. and Tian, L. (2016), 'Comment on: 'personalized dose finding using outcome weighted learning".

Chakraborty, B. and Moodie, E. (2013), *Statistical methods for dynamic treatment regimes*, Springer.

Chen, G., Zeng, D. and Kosorok, M. R. (2016), 'Personalized dose finding using outcome weighted learning', *Journal of the American Statistical Association* **111**(516), 1509–1521.

Chen, J., Fu, H., He, X., Kosorok, M. R. and Liu, Y. (2017), 'Estimating individualized treatment rules for ordinal treatments', *arXiv preprint arXiv:1702.04755* .

Chevret, S. (2006), *Statistical methods for dose-finding experiments*, Vol. 24, Wiley Online Library.

Drucker, H., Burges, C. J., Kaufman, L., Smola, A. J. and Vapnik, V. (1997), Support vector regression machines, *in* 'Advances in neural information processing systems', pp. 155–161.

Eberts, M., Steinwart, I. et al. (2013), 'Optimal regression rates for svms using gaussian kernels', *Electronic Journal of Statistics* **7**, 1–42.

Fong, C., Hazlett, C., Imai, K. et al. (2018), 'Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements', *The Annals of Applied Statistics* **12**(1), 156–177.

Freedman, B. I., Shenoy, R. N., Planer, J. A., Clay, K. D., Shihabi, Z. K., Burkart, J. M., Cardona, C. Y., Andries, L., Peacock, T. P., Sabio, H. et al. (2010), 'Comparison of glycated albumin and hemoglobin a1c concentrations in diabetic subjects on peritoneal and hemodialysis', *Peritoneal Dialysis International* **30**(1), 72–79.

Harrell, F. E. (2001), Ordinal logistic regression, *in* 'Regression modeling strategies', Springer, pp. 331–343.

Henderson, R., Ansell, P. and Alshibani, D. (2010), 'Regret-regression for optimal dynamic treatment regimes', *Biometrics* **66**(4), 1192–1201.

Imai, K. and Ratkovic, M. (2014), 'Covariate balancing propensity score', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(1), 243–263.

Le Thi Hoai, A. and Tao, P. D. (1997), 'Solving a class of linearly constrained indefinite quadratic problems by dc algorithms', *Journal of global optimization* **11**(3), 253–285.

Liang, M. and Yu, M. (2018), 'A semiparametric approach to model effect modification', *arXiv preprint arXiv:1804.05373* .

Little, R. R., Rohlfing, C. L. and Sacks, D. B. (2011), 'Status of hemoglobin a1c measurement and goals for improvement: from chaos to order for improving diabetes care', *Clinical chemistry* **57**(2), 205–214.

Lou, Z., Shao, J. and Yu, M. (2017), 'Optimal treatment assignment to maximize expected outcome with multiple treatments', *Biometrics* .

Luedtke, A. R. and van der Laan, M. J. (2016), 'Comment on: 'personalized dose finding using outcome weighted learning", *Journal of the American Statistical Association* **111**(516), 1526–1530.

Moodie, E. E., Chakraborty, B. and Kramer, M. S. (2012), 'Q-learning for estimating optimal dynamic treatment rules from observational data', *Canadian Journal of Statistics* **40**(4), 629–645.

Moodie, E. E., Platt, R. W. and Kramer, M. S. (2009), 'Estimating response-maximized decision rules with applications to breastfeeding', *Journal of the American Statistical Association* **104**(485), 155–165.

Nitin, S. (2010), 'Hba1c and factors other than diabetes mellitus affecting it', *Singapore Med J* **51**(8), 616–622.

Platt, J. (1999), 'Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods', *Advances in large margin classifiers* **10**(3), 61–74.

Radin, M. S. (2014), 'Pitfalls in hemoglobin a1c measurement: when results may be misleading', *Journal of general internal medicine* **29**(2), 388–394.

Rubin, D. B. (1974), 'Estimating causal effects of treatments in randomized and nonrandomized studies.', *Journal of educational Psychology* **66**(5), 688.

Steinwart, I. and Christmann, A. (2008), *Support vector machines*, Springer Science & Business Media.

Tibshirani, R. (1996), 'Regression shrinkage and selection via the lasso', *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 267–288.

Vapnik, V. (1998), *Statistical learning theory. 1998*, Wiley, New York.

Vapnik, V. (2013), *The nature of statistical learning theory*, Springer science & business media.

Wallace, M. P. and Moodie, E. E. (2015), 'Doubly-robust dynamic treatment regimen estimation via weighted least squares', *Biometrics* **71**(3), 636–644.

Weykamp, C. (2013), 'Hba1c: a review of analytical and clinical aspects', *Annals of laboratory medicine* **33**(6), 393–400.

Zhao, Y., Kosorok, M. R. and Zeng, D. (2009), 'Reinforcement learning design for cancer clinical trials', *Statistics in medicine* **28**(26), 3294–3315.

Zhao, Y.-Q., Zeng, D., Laber, E. B. and Kosorok, M. R. (2015), 'New statistical learning methods for estimating optimal dynamic treatment regimes', *Journal of the American Statistical Association* **110**(510), 583–598.

Zhao, Y., Zeng, D., Rush, A. J. and Kosorok, M. R. (2012), 'Estimating individualized treatment rules using outcome weighted learning', *Journal of the American Statistical Association* **107**(499), 1106–1118.

Zhou, X., Mayer-Hamblett, N., Khan, U. and Kosorok, M. R. (2017), 'Residual weighted learning for estimating individualized treatment rules', *Journal of the American Statistical Association* **112**(517), 169–187.

Chapter 9

Appendix

A.1 Derivation of Algorithm for One-sided Dose Interval

As is claimed in Section 4.2, the DC algorithm will repeatedly update

$$\beta^{t+1} = \arg\min_{\beta} \big(S_1(\beta) - [\nabla S_2(\beta^t)]^T(\beta - \beta^t)\big) \tag{A.1}$$

We will derive the form of quadratic programming by first showing that

$$\begin{aligned}\nabla S_2(\beta) = \nabla_\beta \Bigg(& \sum_{i=1}^{n} \Big(\frac{1}{P(A_i|X_i)} \big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)\Psi_2\big(\phi(X_i)^T\beta - A_i\big)\Big) \\ & + \sum_{i=1}^{n} \Big(\frac{1}{P(A_i|X_i)} \alpha(X_i)\Phi\big(S(X_i) - Y_i\big)\Psi_2\big(A_i - \phi(X_i)^T\beta\big)\Big)\Bigg) \\ = & \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)} \big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big) Q_i(\beta) \frac{\phi(X_i)}{\epsilon} \\ & - \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)} \alpha(X_i)\Phi\big(S(X_i) - Y_i\big)\tilde{Q}_i(\beta)\frac{\phi(X_i)}{\epsilon}\end{aligned}$$

where $Q_i(\beta) = I_{\left(A_i - \phi(X_i)\beta < -\epsilon\right)}$ and $\tilde{Q}_i(\beta) = I_{\left(A_i - \phi(X_i)\beta > \epsilon\right)}$

Following that, we have

$$\nabla S_2(\beta^t)\beta = \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)Q_i(\beta^t)\frac{\phi(X_i)\beta}{\epsilon} + \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\alpha(X_i)\Phi\big(S(X_i)-Y_i\big)\tilde{Q}_i(\beta^t)(\frac{-\phi(X_i)\beta}{\epsilon})$$

Therefore,

$$\begin{aligned}
&\nabla S_2(\beta^t)\beta + \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)Q_i(\beta^t)(-\frac{A_i}{\epsilon}) \\
&\qquad + \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\alpha(X_i)\Phi\big(S(X_i)-Y_i\big)\tilde{Q}_i(\beta^t)\frac{A_i}{\epsilon} \\
&= \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)Q_i(\beta^t)\Big(\frac{\phi(X_i)\beta - A_i}{\epsilon}\Big) \\
&\qquad + \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\alpha(X_i)\Phi\big(S(X_i)-Y_i\big)\tilde{Q}_i(\beta^t)\Big(\frac{A_i - \phi(X_i)\beta}{\epsilon}\Big) \qquad \text{(A.2)}
\end{aligned}$$

By plugging in the previous results into the DC procedure,

$$
\begin{aligned}
\beta^{t+1} &= \arg\min_{\beta} \quad S_1(\beta) - [\nabla S_2(\beta^t)]^T(\beta - \beta^t) \\
&= \arg\min_{\beta} \quad S_1(\beta) - [\nabla S_2(\beta^t)]^T\beta \\
&= \arg\min_{\beta} \quad S_1(\beta) - \Big(A.2\Big) \\
&= \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)\Big[\Psi_1\big(\phi(X_i)\beta, A_i\big) - Q_i(\beta^t)\Big(\frac{\phi(X_i)\beta - A_i}{\epsilon}\Big)\Big] \\
&+ \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\alpha(X_i)\Phi\big(S(X_i) - Y_i\big)\Big[\Psi_1\big(A_i, \phi(X_i)\beta\big) - \tilde{Q}_i(\beta^t)\Big(\frac{A_i - \phi(X_i)\beta}{\epsilon}\Big)\Big] \\
&+ \frac{\lambda_n}{2}\beta^T K(X_i, X_i)\beta \\
&= \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\big(1-\alpha(X_i)\big)\Phi\big(Y_i - S(X_i)\big)\Big[W_i(\beta^t)\Psi_1\big(\phi(X_i)\beta, A_i\big) + Q_i(\beta^t)\Psi_1\Big(A_i, \phi(X_i)\beta\Big)\Big] \\
&+ \sum_{i=1}^{n} \frac{1}{P(A_i|X_i)}\alpha(X_i)\Phi\big(S(X_i) - Y_i\big)\Big[\tilde{W}_i(\beta^t)\Psi_1\big(A_i, \phi(X_i)\beta\big) + \tilde{Q}_i(\beta^t)\Psi_1\Big(\phi(X_i)\beta, A_i\Big)\Big] \\
&+ \frac{\lambda_n}{2}\beta^T K(X_i, X_i)\beta \qquad \text{(A.3)}
\end{aligned}
$$

where $W_i(\beta) = I_{(A_i - \phi(X_i)\beta > -\epsilon)}$ and $\tilde{W}_i(\beta) = I_{(A_i - \phi(X_i)\beta < \epsilon)}$.

To solve the optimization above problem, we now show that it is essentially a quadratic programming problem which can be easily solved using standard packages.

Starting from equation (A.3), let $H_i = \frac{I\big(Y_i > S(X_i)\big)\big(1-\alpha(X_i)\big)}{P(A_i|X_i)\lambda_n}$ and $\tilde{H}_i = \frac{I\big(S(X_i) > Y_i\big)\alpha(X_i)}{P(A_i|X_i)\lambda_n}$, the equation (A.3) can be written as

$$
\begin{aligned}
\arg\min_{\beta} &\sum_{i=1}^{n}(H_i W_i^t + \tilde{H}_i\tilde{Q}_i^t)\Psi_1\big(\phi(X_i)\beta, A_i\big) + \sum_{i=1}^{n}(H_i Q_i^t + \tilde{H}_i\tilde{W}_i^t)\Psi_1\big(A_i, \phi(X_i)\beta\big) \\
&+ \frac{1}{2}\beta^T K(X_i, X_i)\beta
\end{aligned}
$$

which is equivalent to

$$\begin{aligned}
&\arg\min_{\beta,b} \sum_{i=1}^{n}(H_i W_i^t + \tilde{H}_i \tilde{Q}_i^t)\xi_i + \sum_{i=1}^{n}(H_i Q_i^t + \tilde{H}_i \tilde{W}_i^t)\tilde{\xi}_i + \frac{1}{2}\beta^T K(X_i, X_i)\beta \\
&\quad s.t. \quad \xi_i - \phi(X_i)\beta - b + A_i \geq 0, \qquad \xi_i \geq 0 \\
&\qquad\quad \tilde{\xi}_i + \phi(X_i)\beta + b - A_i \geq 0, \qquad \tilde{\xi}_i \geq 0
\end{aligned}$$

Applying the Lagrangian multiplier, the problem above is equivalent to

$$\begin{aligned}
&\arg\max_{\alpha_i,\tilde{\alpha}_i,\mu_i,\tilde{\mu}_i} \sum_{i=1}^{n}(H_i W_i^t + \tilde{H}_i \tilde{Q}_i^t)\xi_i + \sum_{i=1}^{n}(H_i Q_i^t + \tilde{H}_i \tilde{W}_i^t)\tilde{\xi}_i + \frac{1}{2}\beta^T K(X_i, X_i)\beta \\
&\quad - \sum_{i=1}^{n}\alpha_i(\xi_i - \phi(X_i)\beta - b + A_i) - \sum_{i=1}^{n}\tilde{\alpha}_i(\tilde{\xi}_i + \phi(X_i)\beta + b - A_i) - \sum_{i=1}^{n} u_i\xi_i - \sum_{i=1}^{n}\tilde{u}_i\tilde{\xi}_i \\
&s.t. \qquad \beta + \sum_{i=1}^{n}\alpha_i X_i - \sum_{i=1}^{n}\tilde{\alpha}_i X_i = 0 \\
&\qquad \sum_{i=1}^{n}\alpha_i - \sum_{i=1}^{n}\tilde{\alpha}_i = 0 \\
&\qquad H_i W_i^t + \tilde{H}_i \tilde{Q}_i^t - \alpha_i - \mu_i = 0 \quad \forall i = 1, \ldots, n \\
&\qquad H_i Q_i^t + \tilde{H}_i \tilde{W}_i^t - \tilde{\alpha}_i - \tilde{\mu}_i = 0 \quad \forall i = 1, \ldots, n
\end{aligned} \tag{A.4}$$

Combined with the constraint that $\alpha_i \geq 0,\ \tilde{\alpha}_i \geq 0,\ \mu_i \geq 0,\ \tilde{\mu}_i \geq 0$, we have the final quadratic programming problem as

$$\begin{aligned}
&\arg\min_{\beta,b} \qquad \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}(\tilde{\alpha}_i - \alpha_i) < X_i, X_j > (\tilde{\alpha}_j - \alpha_j) - \sum_{i=1}^{n}(\tilde{\alpha}_i - \alpha_i)A_i \\
&\quad s.t. \qquad \sum_{i=1}^{n}\alpha_i - \sum_{i=1}^{n}\tilde{\alpha}_i = 0 \\
&\qquad\quad 0 \leq \alpha_i \leq H_i W_i^t + \tilde{H}_i \tilde{Q}_i^t \\
&\qquad\quad 0 \leq \tilde{\alpha}_i \leq H_i Q_i^t + \tilde{H}_i \tilde{W}_i^t
\end{aligned} \tag{A.5}$$

where the solution of β are attained by

$$\hat{\beta} = \sum_{i=1}^{n} (\tilde{\alpha}_i - \alpha_i) X_i$$

A.2 Derivation of Algorithm for Two-sided Dose Interval

In order to derive the final quadratic programming problem from the DC minimization objective function

$$\beta^{t+1} = \arg\min_{\beta} \left(S_1(\beta) - [\nabla S_2(\beta^t)]^T (\beta - \beta^t) \right)$$

We point out that

$$S_1(f_L, f_U) = \sum_{i=1}^{n} H_i \Psi_1^{out}(f_L, A_i, f_U) + \sum_{i=1}^{n} \tilde{H}_i \Psi_1^{in}(f_L, A_i, f_U) + \frac{1}{2}||f_L||_2^2 + \frac{1}{2}||f_U||_2^2$$

and

$$S_2(f_L, f_U) = \sum_{i=1}^{n} H_i \Psi_2^{out}(f_L, A_i, f_U) + \sum_{i=1}^{n} \tilde{H}_i \Psi_2^{in}(f_L, A_i, f_U)$$

where H_i and $\tilde{H}_i$ are defined as in the previous section.

It follows that

$$\begin{aligned} \nabla S_2(f_L^t, f_U^t) \begin{pmatrix} \beta_L \\ \beta_U \end{pmatrix} = & \sum_{i=1}^{n} H_i \begin{pmatrix} \frac{X_i}{\epsilon} 1_{A_i < (f_L^t - \epsilon)} & -\frac{X_i}{\epsilon} 1_{A_i > (f_U^t + \epsilon)} \end{pmatrix} \begin{pmatrix} \beta_L \\ \beta_U \end{pmatrix} \\ & + \sum_{i=1}^{n} \tilde{H}_i \begin{pmatrix} \frac{X_i}{\epsilon} 1_{A_i < (f_L^t + \epsilon)} & \frac{-X_i}{\epsilon} 1_{A_i > (f_U^t - \epsilon)} \end{pmatrix} \begin{pmatrix} \beta_L \\ \beta_U \end{pmatrix} \end{aligned}$$

Denote $W_i^L = 1_{A_i > (f_L^t - \epsilon)}$, $\tilde{W}_i^L = 1_{A_i > (f_L^t + \epsilon)}$, $Q_i^L = 1_{A_i < (f_L^t - \epsilon)}$, $\tilde{Q}_i^L = 1_{A_i < (f_L^t + \epsilon)}$ and $W_i^U = 1_{A_i < (f_U^t + \epsilon)}$, $\tilde{W}_i^U = 1_{A_i < (f_U^t - \epsilon)}$, $Q_i^U = 1_{A_i > (f_U^t + \epsilon)}$, $\tilde{Q}_i^U = 1_{A_i > (f_U^t - \epsilon)}$.

Then we have

$$
\begin{aligned}
&\arg \min_{(X_i\beta_L+b_L)<(X_i\beta_U+b_U)} S_1(\beta_L,\beta_U) - \nabla S_2(\beta_L,\beta_U) \begin{pmatrix} \beta_L - \beta_L^t \\ \beta_U - \beta_U^t \end{pmatrix} \\
&\arg \min_{(X_i\beta_L+b_L)<(X_i\beta_U+b_U)} S_1(\beta_L,\beta_U) - \nabla S_2(\beta_L,\beta_U) \begin{pmatrix} \beta_L \\ \beta_U \end{pmatrix} \\
&= \sum_{i=1}^{n} H_i \Big[W_i^L \Psi_1(f_L(X_i) - A_i) + Q_i^L \Psi_1(A_i - f_L(X_i)) \Big] \\
&+ \sum_{i=1}^{n} H_i \Big[W_i^U \Psi_1(A_i - f_U(X_i)) + Q_i^U \Psi_1(f_U(X_i) - A_i) \Big] \\
&+ \sum_{i=1}^{n} \tilde{H}_i \Big[\tilde{W}_i^L \Psi_1(f_L(X_i) - A_i) + \tilde{Q}_i^L \Psi_1(A_i - f_L(X_i)) \Big] \\
&+ \sum_{i=1}^{n} \tilde{H}_i \Big[\tilde{W}_i^U \Psi_1(A_i - f_U(X_i)) + \tilde{Q}_i^U \Psi_1(f_U(X_i) - A_i) \Big] \\
&+ \frac{1}{2}||f_L||_2^2 + \frac{1}{2}||f_U||_2^2
\end{aligned} \tag{A.6}
$$

Starting from here, we use the primal-dual property to derive the final quadratic programming form. The primal problem is defined as follows.

$$
\begin{aligned}
&\arg\min \sum_{i=1}^{n} (H_i W_i^L + \tilde{H}_i \tilde{W}_i^L) \xi_i^L + \sum_{i=1}^{n} (H_i Q_i^L + \tilde{H}_i \tilde{Q}_i^L) \tilde{\xi}_i^L \\
&\qquad + \sum_{i=1}^{n} (H_i Q_i^U + \tilde{H}_i \tilde{Q}_i^U) \xi_i^U + \sum_{i=1}^{n} (H_i W_i^U + \tilde{H}_i \tilde{W}_i^U) \tilde{\xi}_i^U \\
&\qquad + \frac{1}{2}||f_L||_2^2 + \frac{1}{2}||f_U||_2^2 \\
S.t. \quad &(\xi_i^L - \phi(X_i)\beta_L - b_L + A_i) \geq 0 \;\; \xi_i^L \geq 0 \\
&(\tilde{\xi}_i^L + \phi(X_i)\beta_L + b_L - A_i) \geq 0 \;\; \tilde{\xi}_i^L \geq 0 \\
&(\xi_i^U - \phi(X_i)\beta_U - b_R + A_i) \geq 0 \;\; \xi_i^U \geq 0 \\
&(\tilde{\xi}_i^U + \phi(X_i)\beta_U + b_R - A_i) \geq 0 \;\; \tilde{\xi}_i^U \geq 0 \\
&(\phi(X_i)\beta_U + b_R - \phi(X_i)\beta_L - b_L) \geq 2\epsilon
\end{aligned} \tag{A.7}
$$

The corresponding Dual problem after applying Lagrangian multiplier is as following.

$$
\begin{aligned}
\arg\max & \sum_{i=1}^{n}(H_i W_i^L + \tilde{H}_i \tilde{W}_i^L)\xi_i^L + \sum_{i=1}^{n}(H_i Q_i^L + \tilde{H}_i \tilde{Q}_i^L)\tilde{\xi}_i^L \\
& + \sum_{i=1}^{n}(H_i Q_i^U + \tilde{H}_i \tilde{Q}_i^U)\xi_i^U + \sum_{i=1}^{n}(H_i W_i^U + \tilde{H}_i \tilde{W}_i^U)\tilde{\xi}_i^U \\
& + \frac{1}{2}||f_L||_2^2 + \frac{1}{2}||f_U||_2^2 \\
& - \sum_{i=1}^{n}\alpha_i^L(\xi_i^L - \phi(X_i)\beta_L - b_L + A_i) - \sum_{i=1}^{n}\mu_i^L\xi_i^L \\
& - \sum_{i=1}^{n}\tilde{\alpha}_i^L(\tilde{\xi}_i^L + \phi(X_i)\beta_L + b_L - A_i) - \sum_{i=1}^{n}\tilde{\mu}_i^L\tilde{\xi}_i^L \\
& - \sum_{i=1}^{n}\alpha_i^U(\xi_i^U - \phi(X_i)\beta_U - b_R + A_i) - \sum_{i=1}^{n}\mu_i^U\xi_i^U \\
& - \sum_{i=1}^{n}\tilde{\alpha}_i^U(\tilde{\xi}_i^U + \phi(X_i)\beta_U + b_R - A_i) - \sum_{i=1}^{n}\tilde{\mu}_i^U\tilde{\xi}_i^U \\
& - \sum_{i=1}^{n} m_i(\phi(X_i)\beta_U + b_R - \phi(X_i)\beta_L - b_L - 2\epsilon)
\end{aligned}
$$

$$
\begin{aligned}
s.t. \quad & \beta_L + \sum_{i=1}^{n}\alpha_i^L\phi(X_i) - \sum_{i=1}^{n}\tilde{\alpha}_i^L\phi(X_i) + \sum_{i=1}^{n}m_i^L\phi(X_i) = 0 \\
& \beta_U + \sum_{i=1}^{n}\alpha_i^U\phi(X_i) - \sum_{i=1}^{n}\tilde{\alpha}_i^U\phi(X_i) - \sum_{i=1}^{n}m_i^U\phi(X_i) = 0 \\
& \sum_{i=1}^{n}\alpha_i^L - \tilde{\alpha}_i^L + m_i = 0 \\
& \sum_{i=1}^{n}\alpha_i^U - \tilde{\alpha}_i^U - m_i = 0 \\
& H_i W_i^L + \tilde{H}_i \tilde{W}_i^L - \alpha_i^L - \mu_i^L = 0, \ \ \forall i = 1, \ldots, n \\
& H_i Q_i^L + \tilde{H}_i \tilde{Q}_i^L - \alpha_i^L - \mu_i^L = 0, \ \ \forall i = 1, \ldots, n \\
& H_i Q_i^U + \tilde{H}_i \tilde{Q}_i^U - \alpha_i^L - \mu_i^L = 0, \ \ \forall i = 1, \ldots, n \\
& H_i W_i^U + \tilde{H}_i \tilde{W}_i^U - \alpha_i^L - \mu_i^L = 0, \ \ \forall i = 1, \ldots, n
\end{aligned}
$$

Plugging in the constraint results in the following simplified form.

$$
\begin{aligned}
\arg\min_{\beta,b} \quad & \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}(\tilde{\alpha}_i^L - \alpha_i^L - m_i) < X_i, X_j > (\tilde{\alpha}_j^L - \alpha_j^L - m_j) \\
& + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}(\tilde{\alpha}_i^U - \alpha_i^U + m_i) < X_i, X_j > (\tilde{\alpha}_j^U - \alpha_j^U + m_j) \\
& - \sum_{i=1}^{n}(\tilde{\alpha}_i^L - \alpha_i^L + \tilde{\alpha}_i^U - \alpha_i^U)A_i - 2\epsilon\sum_{i=1}^{n} m_i \\
s.t. \quad & 0 \le \alpha_i^L \le H_i W_i^L + \tilde{H}_i \tilde{W}_i^L \\
& 0 \le \tilde{\alpha}_i^L \le H_i Q_i^L + \tilde{H}_i \tilde{Q}_i^L \\
& 0 \le \alpha_i^U \le H_i Q_i^U + \tilde{H}_i \tilde{Q}_i^U \\
& 0 \le \tilde{\alpha}_i^U \le H_i W_i^U + \tilde{H}_i \tilde{W}_i^U \\
& 0 \le m_i \\
& \sum_{i=1}^{n} \alpha_i^L - \tilde{\alpha}_i^L + m_i = 0 \\
& \sum_{i=1}^{n} \alpha_i^U - \tilde{\alpha}_i^U - m_i = 0
\end{aligned} \tag{A.8}
$$

After solving the problem above,

$$
\begin{aligned}
\hat{\beta}_L &= \sum_{i=1}^{n}(\tilde{\alpha}_i^L - \alpha_i^L - m_i)X_i \\
\hat{\beta}_U &= \sum_{i=1}^{n}(\tilde{\alpha}_i^U - \alpha_i^U + m_i)X_i
\end{aligned}
$$

A.3 A Proof of Theorem 5.2.1

Proof. First, let f^* be a function such that $f^* \in F \triangleq \{f : \forall x \in \mathcal{X}, P(Y(f(x), x) > S(x)) = \alpha(x)\}$, and then let f' be a function that $f' \neq f^*$.

Denote $\mathcal{X}_0 = \{x : x \in \mathcal{X}, \exists f \in F \ s.t. \ f'(x) = f(x)\}$, $\mathcal{X}_1 = \{x : x \in \mathcal{X}, \ f'(x) > \max_{f\in F} f(x)\}$, and $\mathcal{X}_2 = \{x : x \in \mathcal{X}, \ f'(x) < \min_{f\in F} f(x)\}$

Therefore, under the ignorability assumption, $P(Y > S(x)|A = f'(x), x) = P(Y(f'(x), x) > S(x)) > \alpha(x)$ for $x \in \mathcal{X}_1$, while $P(Y > S(x)|A = f'(x), x) = P(Y(f'(x), x) > S(x)) < \alpha(x)$ for $x \in \mathcal{X}_2$.

Then,

$$
\begin{aligned}
&R_I(f') - R_I(f^*) \\
&= E\Big[\frac{1}{P(A \mid X)}\Big(\big(1-\alpha(X)\big)I\big(Y > S(X)\big)\Big(I\big(A < f'(X)\big) - I\big(A < f^*(X)\big)\Big) \\
&\quad + \alpha(X)I\big(Y < S(X)\big)\Big(I\big(A > f'(X)\big) - I\big(A > f^*(X)\big)\Big)\Big)\Big] \\
&= E\Bigg[\int_{f^*(X)}^{f'(X)} \Big(\big(1-\alpha(X)\big)P\big(Y > S(X)|A = a, X\big) \\
&\qquad\qquad - \alpha(X)P\big(Y < S(X)|A = a, X\big)\Big)da \,\Big|\, X \in \mathcal{X}_1\Bigg] P(X \in \mathcal{X}_1) \qquad \text{(A1)} \\
&\quad + E\Bigg[\int_{f'(X)}^{f^*(X)} \Big(-\big(1-\alpha(X)\big)P\big(Y > S(X)|A = a, X\big) \\
&\qquad\qquad + \alpha(X)P\big(Y < S(X)|A = a, X\big)\Big)da \,\Big|\, X \in \mathcal{X}_2\Bigg] P\big(X \in \mathcal{X}_2\big) \qquad \text{(A2)} \\
&\quad + E\Bigg[\int_{f^*(X)}^{f'(X)} \Big(\big(1-\alpha(X)\big)\alpha(X) - \alpha(X)\big(1-\alpha(X)\big)\Big)da \,\Big|\, X \in \mathcal{X}_0\Bigg] P\big(X \in \mathcal{X}_0\big) \qquad \text{(A3)}
\end{aligned}
$$

Where

$$
\begin{aligned}
(A3) &= 0 \\
(A1) &= E\left[\int_{f^*(X)}^{f'(X)} \Big(P\big(Y > S(X)|A = a, X\big) - \alpha(X)\Big)da \,\Big|\, X \in \mathcal{X}_1\right] P\big(X \in \mathcal{X}_1\big) \\
&= E\left[\int_{\max_{f\in\mathcal{F}} f(X)}^{f'(X)} \Big(P\big(Y > S(X)|A = a, X\big) - \alpha(X)\Big)da \,\Big|\, X \in \mathcal{X}_1\right] P\big(X \in \mathcal{X}_1\big) \\
&\geq 0 \\
&> 0 \qquad \text{if} \quad P(X \in \mathcal{X}_1) \neq 0
\end{aligned}
$$

$$
\begin{aligned}
(A2) =& E\left[\int_{f'(X)}^{f^*(X)} \Big(- P\big(Y > S(X)|A=a, X\big) + \alpha(X)\Big) da \middle| X \in \mathcal{X}_2\right] P(X \in \mathcal{X}_2) \\
=& E\left[\int_{f'(X)}^{\min_{f\in\mathcal{F}} f(X))} \Big(- P\big(Y > S(X)|A=a, X\big) + \alpha(X)\Big) da \middle| X \in \mathcal{X}_2\right] P(X \in \mathcal{X}_2) \\
\geq& 0 \\
>& 0 \qquad \text{if} \quad P(X \in \mathcal{X}_2) \neq 0
\end{aligned}
$$

Therefore, $R_I(f') - R_I(f^*) \geq 0$. And $R_I(f') - R_I(f^*) = 0$ if and only if $P(X \in \mathcal{X}_1) = P(X \in \mathcal{X}_2) = 0$, i.e, $f' \in F$. Hence any possible optimal left bound function $f_{I,L,opt}(x) = \arg\min_f R_I(f)$ belongs to F, and according to the assumption of partial monotonicity for PDI (3.1) that $P\Big(Y(a_L, x) > S(x)\Big) \leq \alpha(x)$, $P\Big(Y(a_U, x) > S(x)\Big) \geq \alpha(x)$ and $P(Y(a, x) > S(x))$ across $\alpha(x)$ at most once, we have that that $P\left(Y\Big(f_{I,L,opt}(x), x\Big) > S(x)\right) = \alpha(x)$ and $P\left(Y\Big(a'(x), x\Big) > S(x)\right) \geq \alpha(x)$ for $\forall x \in \mathcal{X}$, where $a'(x)$ is an arbitrary measurable function s.t. $f_{I,L,opt}(x) \leq a'(x) \leq a_U$. Then theorem 5.2.1 is proved.

□

A.4 A Proof of Corollary 5.2.1

Proof. First,

$$F_L \triangleq \{f : \forall x \in \mathcal{X}, P(Y(f(x), x) > S(x)) = \alpha(x) \quad and \quad \exists a^* < f(x),\ P(Y(a^*, x) > S(x)) \leq \alpha(x)\}$$

$$F_U \triangleq \{f : \forall x \in \mathcal{X}, P(Y(f(x), x) > S(x)) = \alpha(x) \quad and \quad \exists a^* > f(x),\ P(Y(a^*, x) > S(x)) \leq \alpha(x)\}$$

Let $f_L^* \in F_L$ and $f_U^* \in F_U$, then for arbitrary f_L' and f_U', the sample space of $\mathcal{X}$ can be

decomposed as $\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_3 \cup \mathcal{X}_4$ where

$$\mathcal{X}_0 = \{x : x \in \mathcal{X}, \quad \exists f_L \in F_L \; s.t. \; f'_L(x) = f_L(x) \quad and \quad \exists f_U \in F_U \; s.t. \; f'_U(x) = f_U(x)\}$$

$$\mathcal{X}_1 = \{x : x \in \mathcal{X}, \quad \max_{f_L \in F_L} f_L(x) < f'_L(x) \leq f'_U(x) < \min_{f_U \in F_U} f_U(x)\}$$

$$\mathcal{X}_2 = \{x : x \in \mathcal{X}, \quad f'_L(x) < \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x) < f'_U(x)\}$$

$$\mathcal{X}_3 = \{x : x \in \mathcal{X}, \quad f'_L(x) \leq \min_{f_L \in F_L} f_L(x) \leq f'_U(x) \leq \max_{f_U \in F_U} f_U(x)\}$$

$$\mathcal{X}_4 = \{x : x \in \mathcal{X}, \quad \min_{f_L \in F_L} f_L(x) \leq f'_L(x) \leq \max_{f_U \in F_U} f_U(x) \leq f'_U(x)\}$$

$$\mathcal{X}_5 = \{x : x \in \mathcal{X}, \quad f'_L(x) \leq f'_U(x) < \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x)\}$$

$$\mathcal{X}_6 = \{x : x \in \mathcal{X}, \quad \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x) < f'_L(x) \leq f'_U(x)\}$$

then

$$\begin{aligned}
&R_I(f') - R_I(f^*) \\
&= E\Big[\frac{1}{P(A \,|\, X)}\Big(\big(1-\alpha(x)\big) I\big(Y > S(X)\big)\Big(I(A \notin [f'_L(X), f'_U(X)]) - I(A \notin [f^*_L(X), f^*_U(X)])\Big) \\
&\quad + \alpha(x) I\big(Y < S(X)\big)\Big(I(A \in [f'_L(X), f'_U(X)]) - I(A \in [f^*_L(X), f^*_U(X)])\Big)\Big)\Big] \\
&= E\Big[\ldots \Big| X \in \mathcal{X}_1\Big] \, P(X \in \mathcal{X}_1) && \text{(B1)} \\
&+ E\Big[\ldots \Big| X \in \mathcal{X}_2\Big] \, P(X \in \mathcal{X}_2) && \text{(B2)} \\
&+ E\Big[\ldots \Big| X \in \mathcal{X}_3\Big] \, P(X \in \mathcal{X}_3) && \text{(B3)} \\
&+ E\Big[\ldots \Big| X \in \mathcal{X}_4\Big] \, P(X \in \mathcal{X}_4) && \text{(B4)} \\
&+ E\Big[\ldots \Big| X \in \mathcal{X}_5\Big] \, P(X \in \mathcal{X}_5) && \text{(B5)} \\
&+ E\Big[\ldots \Big| X \in \mathcal{X}_6\Big] \, P(X \in \mathcal{X}_6) && \text{(B6)}
\end{aligned}$$

where $(B3) \geq 0$, $(B3) > 0$ if $P(X \in \mathcal{X}_3) > 0$, and $(B4) \geq 0$, $(B4) > 0$ if $P(X \in \mathcal{X}_3) > 0$, according to the results of theorem 5.2.1. These results can be directly applied since it is

easily to see that within $\mathcal{X}_3$ and $\mathcal{X}_4$ the assumption meets the requirement of theorem 5.2.1.

$$
\begin{aligned}
(B0) =& 0 \\
(B1) =& E\Bigg[\int_{f_L^*(X)}^{f_L'(X)} \Big(\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) - \alpha(X)P(Y < S(X)|A=a,X)\Big)da \\
&+ \int_{f_U'(X)}^{f_U^*(X)} \Big(\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) \\
&\qquad\qquad - \alpha(X)P(Y < S(X)|A=a,X)\Big)da \Bigg| X \in \mathcal{X}_1\Bigg] P(X \in \mathcal{X}_1) \\
=& E\Bigg[\int_{\max_{f_L \in F_L} f_L(X)}^{f_L'(X)} \Big(P(Y > S(X)|A=a,X) - \alpha(X)\Big)da \\
&+ \int_{f_U'(X)}^{\min_{f_U \in F_U} f_U(X)} \Big(P(Y > S(X)|A=a,X) - \alpha(X)\Big)da \Bigg| X \in \mathcal{X}_1\Bigg] P(X \in \mathcal{X}_1) \\
\geq& 0 \\
>& 0 \qquad \text{if} \quad P(X \in \mathcal{X}_1) \neq 0 \\
(B2) =& E\Bigg[\int_{f_L'(X)}^{f_L^*(X)} \Big(-\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) + \alpha(X)P(Y < S(X)|A=a,X)\Big)da \\
&+ \int_{f_U^*(X)}^{f_U'(X)} \Big(-\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) \\
&\qquad\qquad + \alpha(X)P(Y < S(X)|A=a,X)\Big)da \Bigg| X \in \mathcal{X}_2\Bigg] P(X \in \mathcal{X}_2) \\
=& E\Bigg[\int_{f_L'(X)}^{\min_{f_L \in F_L} f_L(X)} \Big(-P(Y > S(X)|A=a,X) + \alpha(X)\Big)da \\
&+ \int_{\max_{f_U \in F_U} f_U(X)}^{f_U'(X)} \Big(-P(Y > S(X)|A=a,X) + \alpha(X)\Big)da \Bigg| X \in \mathcal{X}_2\Bigg] P(X \in \mathcal{X}_2) \\
\geq& 0 \\
>& 0 \qquad \text{if} \quad P(X \in \mathcal{X}_2) \neq 0
\end{aligned}
$$

$$
\begin{aligned}
(B5) =& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(-\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) + \alpha(X)P(Y < S(X)|A=a,X)\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) \\
&\qquad\qquad - \alpha(X)P(Y < S(X)|A=a,X)\Big)da \Big| X \in \mathcal{X}_5\Bigg] P(X \in \mathcal{X}_5) \\
=& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(-P(Y > S(X)|A=a,X) + \alpha(X)\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(P(Y > S(X)|A=a,X) - \alpha(X)\Big)da \Big| X \in \mathcal{X}_5\Bigg] P(X \in \mathcal{X}_5) \\
\geq& 0 \\
>& 0 \qquad \text{if} \quad P(X \in \mathcal{X}_5) \neq 0
\end{aligned}
$$

$$
\begin{aligned}
(B6) =& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(-\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) + \alpha(X)P(Y < S(X)|A=a,X)\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X) \\
&\qquad\qquad - \alpha(X)P(Y < S(X)|A=a,X)\Big)da \Big| X \in \mathcal{X}_6\Bigg] P(X \in \mathcal{X}_6) \\
=& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(-P(Y > S(X)|A=a,X) + \alpha(X)\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(P(Y > S(X)|A=a,X) - \alpha(X)\Big)da \Big| X \in \mathcal{X}_6\Bigg] P(X \in \mathcal{X}_6) \\
\geq& 0 \\
>& 0 \qquad \text{if} \quad P(X \in \mathcal{X}_6) \neq 0
\end{aligned}
$$

Therefore, $R_I(f'_U, f'_U) - R_I(f^*_L, f^*_U) \geq 0$, and $R_I(f'_U, f'_U) - R_I(f^*_L, f^*_U) = 0$ if and only if $P(X \notin \mathcal{X}_0) = 0$. Hence for any optimal two-sided bound functions $[f_{I,L,opt}, f_{I,U,opt}] = \arg\min_{f_L \leq f_U} R_I(f_L, f_U)$, we have $f_{I,L,opt} \in F_L$ and $f_{I,L,opt} \in F_U$. According to the assump-

tion of partial unimodality for PDI (3.2) that $P\Big(Y(a_L, x) > S(x)\Big) \leq \alpha(x)$, $P\Big(Y(a_U, x) > S(x)\Big) \leq \alpha(x)$, $\exists a^* \in [a_L, a_U]$, $P(Y(a^*, x) > S(x)) \geq \alpha(x)$ and $P(Y(a, x) > S(x))$ cross $\alpha(x)$ at most twice as a goes from a_L to a_U, we have that that $P\left(Y\Big(f_{I,L,opt}(x), x\Big) > S(x)\right) = P\left(Y\Big(f_{I,U,opt}(x), x\Big) > S(x)\right) = \alpha(x)$ and $P\left(Y\Big(a'(x), x\Big) > S(x)\right) \geq \alpha(x)$ for $\forall x \in \mathcal{X}$, where $a'(x)$ is an arbitrary measurable function s.t. $f_{I,L,opt}(x) \leq a'(x) \leq f_{I,U,opt}(x)$. □

A.5 A Proof of Theorem 5.2.2

Proof. The proof of theorem 5.2.2 can be easily adapted from theorem 5.2.1.
First, let f^* be a function such that $f^* \in F \triangleq \{f : \forall x \in \mathcal{X}, E\left[Y(f(x), x)\right] = S(x)\}$, and then let f' be a function that $f' \neq f^*$.
Denote $\mathcal{X}_0 = \{x : x \in \mathcal{X}, \exists f \in F \ s.t. \ f'(x) = f(x)\}$, $\mathcal{X}_1 = \{x : x \in \mathcal{X}, \ f'(x) > \max_{f \in F} f(x)\}$, and $\mathcal{X}_2 = \{x : x \in \mathcal{X}, \ f'(x) < \min_{f \in F} f(x)\}$

Therefore, under the ignorability assumption, $E\Big[Y|A = f'(x), X = x\Big] = E\Big[Y(f'(x), x)\Big] > S(x)$ for $x \in \mathcal{X}_1$, while $E\Big[Y|A = f'(x), X = x\Big] = E\Big[Y(f'(x), x)\Big] < S(x)$ for $x \in \mathcal{X}_2$.

Then,

$$
\begin{aligned}
&R_h(f') - R_h(f^*) \\
&=E\Big[\frac{1}{P(A\,|\,X)}\Big(\frac{1}{2}\big(Y - S(X)\big)_+\Big(I\big(A < f'(X)\big) - I\big(A < f^*(X)\big)\Big) \\
&\quad + \frac{1}{2}\big(S(X) - Y\big)_+\Big(I\big(A > f'(X)\big) - I\big(A > f^*(X)\big)\Big)\Big)\Big] \\
&=E\Bigg[\int_{f^*(X)}^{f'(X)}\Big(\frac{1}{2}E\Big[\big(Y - S(X)\big)_+|A = a, X\Big] \\
&\qquad\qquad - \frac{1}{2}E\Big[\big(S(X) - Y\big)_+|A = a, X\Big]\Big)da\Big|X \in \mathcal{X}_1\Bigg]P(X \in \mathcal{X}_1) \qquad \text{(A1)} \\
&\quad + E\Bigg[\int_{f'(X)}^{f^*(X)}\Big(-\frac{1}{2}E\Big[\big(Y - S(X)\big)_+|A = a, X\Big] \\
&\qquad\qquad + \frac{1}{2}E\Big[\big(S(X) - Y\big)_+|A = a, X\Big]\Big)da\Big|X \in \mathcal{X}_2\Big]P(X \in \mathcal{X}_2) \qquad \text{(A2)} \\
&\quad + E\Bigg[\int_{f^*(X)}^{f'(X)}\Big(\frac{1}{4}E\Big[\Big|Y - E\big[Y(a, X)\big]\Big|\Big|A = a, X\Big] \\
&\qquad\qquad - \frac{1}{4}E\Big[\Big|Y - E\big[Y(a, X)\big]\Big|\Big|A = a, X\Big]\Big)da\Big|X \in \mathcal{X}_0\Bigg]P(X \in \mathcal{X}_0) \qquad \text{(A3)}
\end{aligned}
$$

Where

$$
\begin{aligned}
(A3) =& 0 \\
(A1) =& E\left[\int_{f^*(X)}^{f'(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+ - \big(S(X)-Y\big)_+ | A=a, X\Big]\Big)da \middle| X \in \mathcal{X}_1\right] P(X \in \mathcal{X}_1) \\
=& E\left[\int_{\max_{f\in\mathcal{F}} f(X)}^{f'(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+ - \big(S(X)-Y\big)_+ | A=a, X\Big]\Big)da \middle| X \in \mathcal{X}_1\right] P(X \in \mathcal{X}_1) \\
\geq& E\left[\int_{\max_{f\in\mathcal{F}} f(X)}^{f'(X)} \Big(\frac{1}{2}E\Big[(Y - E\Big[Y(a,X)\Big])_+ \right. \\
& \qquad \left. - (E\Big[Y(a,X)\Big] - Y)_+ | A=a, X\Big]\Big)da \middle| X \in \mathcal{X}_1\right] P(X \in \mathcal{X}_1) \\
=& 0 \\
>& 0 \qquad \text{if} \quad P(X \in \mathcal{X}_1) \neq 0
\end{aligned}
$$

$$
\begin{aligned}
(A2) =& E\left[\int_{f'(X)}^{f^*(X)} \Big(\frac{1}{2}E\Big[\big(S(X)-Y\big)_+ - \big(Y-S(X)\big)_+ | A=a, X\Big]\Big)da \middle| X \in \mathcal{X}_2\right] P(X \in \mathcal{X}_2) \\
=& E\left[\int_{f'(X)}^{\min_{f\in\mathcal{F}} f(X))} \Big(\frac{1}{2}E\Big[\big(S(X)-Y\big)_+ - \big(Y-S(X)\big)_+ | A=a, X\Big]\Big)da \middle| X \in \mathcal{X}_2\right] P(X \in \mathcal{X}_2) \\
\geq& E\left[\int_{f'(X)}^{\min_{f\in\mathcal{F}} f(X))} \Big(\frac{1}{2}E\Big[(E\Big[Y(a,X)\Big] - Y)_+ \right. \\
& \qquad \left. - (Y - E\big[Y(a,X)\big])_+ | A=a, X\Big]\Big)da \middle| X \in \mathcal{X}_2\right] P(X \in \mathcal{X}_2) \\
=& 0 \\
>& 0 \qquad \text{if} \quad P(X \in \mathcal{X}_2) \neq 0
\end{aligned}
$$

Therefore, $R_h(f') - R_h(f^*) \geq 0$. And $R_h(f') - R_h(f^*) = 0$ if and only if $P(X \in \mathcal{X}_1) = P(X \in \mathcal{X}_2) = 0$, i.e, $f' \in F$. Hence any possible optimal left bound function $f_{h,L,opt}(x) = \arg\min_f R_h(f)$ belongs to F, and according to the assumption of partial monotonicity for

EDI (3.3) that $E[Y(a_L, x)] \leq S(x)$, $E[Y(a_U, x)] \geq S(x)$ and $E[Y(a, x)]$ across $S(x)$ at most once, we have that that $E\left[Y\left(f_{h,L,opt}(x), x\right)\right] = S(x)$ and $E[Y(a'(x), x)] \geq S(x)$ for $\forall x \in \mathcal{X}$, where $a'(x)$ is an arbitrary measurable function s.t. $f_{I,L,opt}(x) \leq a'(x) \leq a_U$. Then theorem 5.2.1 is proved. □

A.6 A Proof of Corollary 5.2.2

Proof. Similarly as the previous theorem,

$$F_L \triangleq \{f : \forall x \in \mathcal{X}, E[Y(f(x), x)] = S(x) \quad and \quad \exists a^* < f(x),\ E[Y(a^*, x)] < S(x)\}$$

$$F_U \triangleq \{f : \forall x \in \mathcal{X}, E[Y(f(x), x)] = S(x) \quad and \quad \exists a^* > f(x),\ E[Y(a^*, x)] < S(x)\}$$

Let $f_L^* \in F_L$ and $f_U^* \in F_U$, then for arbitrary f_L' and f_U', the sample space of $\mathcal{X}$ can be decomposed as $\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_3 \cup \mathcal{X}_4$ where

$$\mathcal{X}_0 = \{x : x \in \mathcal{X}, \quad \exists f_L \in F_L \ s.t.\ f_L'(x) = f_L(x) \quad and \quad \exists f_U \in F_U \ s.t.\ f_U'(x) = f_U(x)\}$$

$$\mathcal{X}_1 = \{x : x \in \mathcal{X}, \quad \max_{f_L \in F_L} f_L(x) < f_L'(x) \leq f_U'(x) < \min_{f_U \in F_U} f_U(x)\}$$

$$\mathcal{X}_2 = \{x : x \in \mathcal{X}, \quad f_L'(x) < \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x) < f_U'(x)\}$$

$$\mathcal{X}_3 = \{x : x \in \mathcal{X}, \quad f_L'(x) \leq \min_{f_L \in F_L} f_L(x) \leq f_U'(x) \leq \max_{f_U \in F_U} f_U(x)\}$$

$$\mathcal{X}_4 = \{x : x \in \mathcal{X}, \quad \min_{f_L \in F_L} f_L(x) \leq f_L'(x) \leq \max_{f_U \in F_U} f_U(x) \leq f_U'(x)\}$$

$$\mathcal{X}_5 = \{x : x \in \mathcal{X}, \quad f_L'(x) \leq f_U'(x) < \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x)\}$$

$$\mathcal{X}_6 = \{x : x \in \mathcal{X}, \quad \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x) < f_L'(x) \leq f_U'(x)\}$$

then

$$
\begin{aligned}
&R_I(f') - R_I(f^*) \\
&= E\Big[\frac{1}{P(A \mid X)}\Big((1-\alpha)\big(Y - S(X)\big)_+\Big(I(A \notin [f'_L(X), f'_U(X)]) - I(A \notin [f^*_L(X), f^*_U(X)])\Big) \\
&\quad + \alpha\big(S(X) - Y\big)_+\Big(I(A \in [f'_L(X), f'_U(X)]) - I(A \in [f^*_L(X), f^*_U(X)])\Big)\Big)\Big] \\
&= E\Big[\dots\Big|X \in \mathcal{X}_1\Big]P(X \in \mathcal{X}_1) && \text{(B1)} \\
&+ E\Big[\dots\Big|X \in \mathcal{X}_2\Big]P(X \in \mathcal{X}_2) && \text{(B2)} \\
&+ E\Big[\dots\Big|X \in \mathcal{X}_3\Big]P(X \in \mathcal{X}_3) && \text{(B3)} \\
&+ E\Big[\dots\Big|X \in \mathcal{X}_4\Big]P(X \in \mathcal{X}_4) && \text{(B4)} \\
&+ E\Big[\dots\Big|X \in \mathcal{X}_5\Big]P(X \in \mathcal{X}_5) && \text{(B5)} \\
&+ E\Big[\dots\Big|X \in \mathcal{X}_6\Big]P(X \in \mathcal{X}_6) && \text{(B6)}
\end{aligned}
$$

where $(B3) \geq 0$, $(B3) > 0$ if $P(X \in \mathcal{X}_3) > 0$, and $(B4) \geq 0$, $(B4) > 0$ if $P(X \in \mathcal{X}_3) > 0$, according to the results of theorem 5.2.2. These results can be directly applied since it is easily to see that within $\mathcal{X}_3$ and $\mathcal{X}_4$ the assumption meets the requirement of theorem 5.2.2.

$$
\begin{aligned}
(B0) =& 0 \\
(B1) =& E\Big[\int_{f_L^*(X)}^{f_L'(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big]-\frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+\int_{f_U'(X)}^{f_U^*(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big] \\
&\qquad\qquad -\frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da\Big|X\in\mathcal{X}_1\Big]P(X\in\mathcal{X}_1) \\
=& E\Big[\int_{\max_{f_L\in F_L} f_L(X)}^{f_L'(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+-\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+\int_{f_U'(X)}^{\min_{f_U\in F_U} f_U(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+ \\
&\qquad\qquad -\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da\Big|X\in\mathcal{X}_1\Big]P(X\in\mathcal{X}_1) \\
\geq& E\Big[\int_{\max_{f_L\in F_L} f_L(X)}^{f_L'(X)} \Big(\frac{1}{2}E\Big[\big(Y-E\big[Y(a,X)\big]\big)_+-\big(E\big[Y(a,X)\big]-Y\big)_+|A=a,X\Big]\Big)da \\
&+\int_{f_U'(X)}^{\min_{f_U\in F_U} f_U(X)} \Big(\frac{1}{2}E\Big[\big(Y-E\big[Y(a,X)\big]\big)_+ \\
&\qquad\qquad -\big(E\big[Y(a,X)\big]-Y\big)_+|A=a,X\Big]\Big)da\Big|X\in\mathcal{X}_1\Big]P(X\in\mathcal{X}_1) \\
=& 0 \\
>& 0 \qquad \text{if} \quad P(X\in\mathcal{X}_1)\neq 0
\end{aligned}
$$

$$
\begin{aligned}
(B2) =& E\Bigg[\int_{f_L'(X)}^{f_L^*(X)} \Big(-\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big]+\frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+\int_{f_U^*(X)}^{f_U'(X)} \Big(-\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big] \\
&\qquad\qquad +\frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da\Bigg|X\in\mathcal{X}_2\Bigg]P(X\in\mathcal{X}_2) \\
=& E\Bigg[\int_{f_L'(X)}^{\min_{f_L\in F_L} f_L(X)} \Big(\frac{1}{2}E\Big[-\big(Y-S(X)\big)_+ +\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+\int_{\max_{f_U\in F_U} f_U(X)}^{f_U'(X)} \Big(\frac{1}{2}E\Big[-\big(Y-S(X)\big)_+ \\
&\qquad\qquad +\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da\Bigg|X\in\mathcal{X}_2\Bigg]P(X\in\mathcal{X}_2) \\
\geq& E\Bigg[\int_{f_L'(X)}^{\min_{f_L\in F_L} f_L(X)} \Big(\frac{1}{2}E\Big[-\big(Y-E[Y(a,X)]\big)_+ +\big(E[Y(a,X)]-Y\big)_+|A=a,X\Big]\Big)da \\
&+\int_{\max_{f_U\in F_U} f_U(X)}^{f_U'(X)} \Big(\frac{1}{2}E\Big[-\big(Y-E[Y(a,X)]\big)_+ \\
&\qquad\qquad +\big(E[Y(a,X)]-Y\big)_+|A=a,X\Big]\Big)da\Bigg|X\in\mathcal{X}_1\Bigg]P(X\in\mathcal{X}_1) \\
=& 0 \\
>& 0 \qquad \text{if} \quad P(X\in\mathcal{X}_2)\neq 0
\end{aligned}
$$

$$
\begin{aligned}
(B5) =& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(-\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big] + \frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big] \\
&\qquad\qquad - \frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \Big| X\in\mathcal{X}_5\Bigg] P(X\in\mathcal{X}_5) \\
=& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(\frac{1}{2}E\Big[-\big(Y-S(X)\big)_+ + \big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+ \\
&\qquad\qquad - \big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \Big| X\in\mathcal{X}_5\Bigg] P(X\in\mathcal{X}_5) \\
\geq& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(\frac{1}{2}E\Big[-\big(Y-E[Y(a,X)]\big)_+ + \big(E[Y(a,X)]-Y\big)_+|A=a,X\Big]\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\frac{1}{2}E\Big[\big(Y-E[Y(a,X)]\big)_+ \\
&\qquad\qquad - \big(E[Y(a,X)]-Y\big)_+|A=a,X\Big]\Big)da \Big| X\in\mathcal{X}_5\Bigg] P(X\in\mathcal{X}_5) \\
=& 0 \\
>& 0 \qquad \text{if} \quad P(X\in\mathcal{X}_5)\neq 0
\end{aligned}
$$

$$
\begin{aligned}
(B6) =& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(-\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big] + \frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+|A=a,X\Big] \\
&\qquad\qquad - \frac{1}{2}E\Big[\big(S(X)-Y\big)_+|A=a,X\Big]\Big)da\Big|X\in\mathcal{X}_6\Bigg]P(X\in\mathcal{X}_6) \\
=& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(\frac{1}{2}E\Big[-\big(Y-S(X)\big)_+ + \big(S(X)-Y\big)_+|A=a,X\Big]\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\frac{1}{2}E\Big[\big(Y-S(X)\big)_+ \\
&\qquad\qquad - \big(S(X)-Y\big)_+|A=a,X\Big]\Big)da\Big|X\in\mathcal{X}_6\Bigg]P(X\in\mathcal{X}_6) \\
\geq& E\Bigg[\int_{f'_L(X)}^{f'_U(X)} \Big(\frac{1}{2}E\Big[-(Y-E[Y(a,X)])_+ + (E[Y(a,X)]-Y)_+|A=a,X\Big]\Big)da \\
&+ \int_{f^*_U(X)}^{f^*_L(X)} \Big(\frac{1}{2}E\Big[(Y-E[Y(a,X)])_+ \\
&\qquad\qquad - (E[Y(a,X)]-Y)_+|A=a,X\Big]\Big)da\Big|X\in\mathcal{X}_6\Bigg]P(X\in\mathcal{X}_6) \\
=& 0 \\
>& 0 \qquad \text{if} \quad P(X\in\mathcal{X}_6)\neq 0
\end{aligned}
$$

Therefore, $R_h(f'_U,f'_U) - R_h(f^*_L,f^*_U) \geq 0$, and $R_h(f'_U,f'_U) - R_h(f^*_L,f^*_U) = 0$ if and only if $P(X\notin\mathcal{X}_0)=0$. Hence for any optimal two sided bound functions $[f_{h,L,opt}, f_{h,U,opt}] = \arg\min_{f_L\leq f_U} R_h(f_L,f_U)$, we have $f_{h,L,opt}\in F_L$ and $f_{h,L,opt}\in F_U$. According to the assumption of partial unimodality for EDI (3.4) that $E\left[Y(a_L,x)\right]\leq S(x)$, $E\left[Y(a_U,x)\right]\leq S(x)$, $\exists a^*\in[a_L,a_U]$, $E\left[Y(a^*,x)\right]\geq S(x)$ and $E[Y(a,x)]$ across $S(x)$ at most twice, we further have that $E[Y(f_{h,L,opt}(x),x)] = E[Y(f_{h,U,opt}(x),x)] = S(x)$ and $E[Y(a'(x),x)]\geq S(x)$ for $\forall x\in\mathcal{X}$, where $a'(x)$ is an arbitrary measurable function s.t. $f_{h,L,opt}(x)\leq a'(x)\leq f_{h,U,opt}(x)$. Then thereom 5.2.2 is proved. □

A.7 A Proof of Theorem 5.3.1

Proof. for any measurable function $f : \mathcal{X} \to \mathbb{R}$, we have

$$
\begin{aligned}
&|R_I(f) - R_{I,\epsilon_n}(f)| \\
=&E\Big|\frac{1}{\epsilon_n P(A \,|\, X)}\Big(\big(1-\alpha(X)\big)I(Y > S(X))\max(\epsilon_n - (f(X)-A)_+, 0) \\
&+ \alpha(X)I(Y < S(X))\max(\epsilon_n - (A - f(X))_+, 0)\Big)\Big| \\
=&E_X\Big|\frac{1}{\epsilon_n}\int_{a\in[a_L,a_U]}\Big(\big(1-\alpha(X)\big)P(Y > S(X)|A=a,X)\max(\epsilon_n - (f(X)-a)_+, 0) \\
&\qquad + \alpha(X)P(Y < S(X)|A=a,X)\max(\epsilon_n - (a - f(X))_+, 0)\Big)da\,\Big| \\
=&E_X\Big|\frac{\big(1-\alpha(X)\big)}{\epsilon_n}\int_{a\in[f(X)-\epsilon_n,f(X)]} P(Y > S(X)|A=a,X)(\epsilon_n - (f(X)-a))da \\
&\qquad + \frac{\alpha(X)}{\epsilon_n}\int_{a\in[f(X),f(X)+\epsilon_n]} P(Y < S(X)|A=a,X)(\epsilon_n - (a - f(X)))da\,\Big| \\
=&E_X\Big|\frac{\big(1-\alpha(X)\big)}{\epsilon_n}\int_{z\in[0,\epsilon_n]} P(Y > S(X)|A=f(X)-z,X)(\epsilon_n - z)dz \\
&\qquad + \frac{\alpha(X)}{\epsilon_n}\int_{z\in[0,\epsilon_n]} P(Y < S(X)|A=f(X)+z,X)(\epsilon_n - z)dz\,\Big| \\
\le&\ \epsilon
\end{aligned}
$$

then it follows the theorem 5.3.1. □

A.8 A Proof of Corollary 5.3.1

Proof. for any measurable function $f_L : \mathcal{X} \to \mathbb{R}$, $f_U : \mathcal{X} \to \mathbb{R}$ and $f_L(X) \leq f_U(X) \forall x \in \mathcal{X}$, we have

$$
\begin{aligned}
&|R_I(f) - R_{I,\epsilon_n}(f)| \\
&= E_X \Bigg| \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{a\in[f_L(X)-\epsilon_n, f_L(X)]} P(Y > S(X)|A = a, X)(\epsilon_n - (f_L(X) - a))da \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{a\in[f_L(X), f_L(X)+\epsilon_n]} P(Y < S(X)|A = a, X)(\epsilon_n - (a - f_L(X)))da \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{a\in[f_U(X)-\epsilon_n, f_U(X)]} P(Y < S(X)|A = a, X)(\epsilon_n - (f_U(X) - a))da \\
&\qquad + \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{a\in[f_U(X), f(X)_R+\epsilon_n]} P(Y > S(X)|A = a, X)(\epsilon_n - (a - f_U(X)))da \Bigg| \\
&= E_X \Bigg| \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} P(Y > S(X)|A = f_L(X) - z, X)(\epsilon_n - z)dz \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} P(Y < S(X)|A = f_L(X) + z, X)(\epsilon_n - z)dz \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} P(Y < S(X)|A = f_U(X) - z, X)(\epsilon_n - z)dz \\
&\qquad + \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} P(Y > S(X)|A = f_U(X) + z, X)(\epsilon_n - z)dz \Bigg| \\
&= 2 * \epsilon_n
\end{aligned}
$$

then it follows the corollary 5.3.1. □

A.9 A Proof of Theorem 5.3.2

Proof. for any measurable function $f : \mathcal{X} \to \mathbb{R}$, we have

$$
\begin{aligned}
&|R_h(f) - R_{h,\epsilon_n}(f)| \\
&= E\Big|\frac{1}{\epsilon_n P(A \,|\, X)}\Big(\big(1-\alpha(X)\big)(Y - S(X))_+ \max(\epsilon_n - (f(X) - A)_+, 0) \\
&\quad + \alpha(X)(S(X) - Y)_+ \max(\epsilon_n - (A - f(X))_+, 0)\Big)\Big| \\
&= E_X\Big|\frac{1}{\epsilon_n}\int_{a\in[a_L,a_U]}\Big(E(\big(1-\alpha(X)\big)(Y - S(X))_+ | A = a, X)\max(\epsilon_n - (f(X) - A)_+, 0) \\
&\qquad + E(\alpha(X)(S(X) - Y)_+ | A = a, X)\max(\epsilon_n - (a - f(X))_+, 0)\Big)da\Big| \\
&= E_X\Big|\frac{\big(1-\alpha(X)\big)}{\epsilon_n}\int_{a\in[f(X)-\epsilon_n, f(X)]} E((Y - S(X))_+ | A = a, X)(\epsilon_n - (f(X) - a))da \\
&\qquad + \frac{\alpha(X)}{\epsilon_n}\int_{a\in[f(X), f(X)+\epsilon_n]} E((S(X) - Y)_+ | A = a, X)(\epsilon_n - (a - f(X)))da\Big| \\
&= E_X\Big|\frac{\big(1-\alpha(X)\big)}{\epsilon_n}\int_{z\in[0,\epsilon_n]} E((Y - S(X))_+ | A = f(X) - z, X)(\epsilon_n - z)dz \\
&\qquad + \frac{\alpha(X)}{\epsilon_n}\int_{z\in[0,\epsilon_n]} E((S(X) - Y)_+ | A = f(X) + z, X)(\epsilon_n - z)dz\Big| \\
&\le C_1\epsilon_n
\end{aligned}
$$

Then it follows theorem 5.3.2. □

A.10 A Proof of Corollary 5.3.2

Proof. for any measurable function $f_L : \mathcal{X} \to \mathbb{R}$, $f_U : \mathcal{X} \to \mathbb{R}$ and $f_L(X) \leq f_U(X) \forall x \in \mathcal{X}$, we have

$$
\begin{aligned}
&|R_h(f) - R_{h,\epsilon_n}(f)| \\
&= E_X \Bigg| \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{a\in[f_L(X)-\epsilon_n, f_L(X)]} E((Y - S(X))_+ | A = a, X)(\epsilon_n - (f_L(X) - a))da \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{a\in[f_U(X), f_U(X)+\epsilon_n]} E((S(X) - Y)_+ | A = a, X)(\epsilon_n - (a - f_L(X)))da \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{a\in[f_U(X)-\epsilon_n, f_U(X)]} E((S(X) - Y)_+ | A = a, X)(\epsilon_n - (f_U(X) - a))da \\
&\qquad + \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{a\in[f_L(X), f_L(X)+\epsilon_n]} E((Y - S(X))_+ | A = a, X)(\epsilon_n - (a - f_U(X)))da \Bigg| \\
&= E_X \Bigg| \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} E((Y - S(X))_+ | A = f_L(X) - z, X)(\epsilon_n - z)dz \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} E((S(X) - Y)_+ | A = f_L(X) + z, X)(\epsilon_n - z)dz \\
&\qquad + \frac{\alpha(X)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} E((S(X) - Y)_+ | A = f_U(X) - z, X)(\epsilon_n - z)dz \\
&\qquad + \frac{\big(1-\alpha(X)\big)}{\epsilon_n} \int_{z\in[0,\epsilon_n]} E((Y - S(X))_+ | A = f_U(X) + z, X)(\epsilon_n - z)dz \Bigg| \\
&\leq 2 * C_1 * \epsilon_n
\end{aligned}
$$

Then it follows Corollary 5.3.2. □

Theorem A.10.1. *Under the assumptions that the optimal regime $f_{\Phi,L,opt} \in B^{\alpha}_{1,\infty}(\mathbb{R}^d)$, a Besov space, such that, $B^{\alpha}_{1,\infty}(\mathbb{R}^d) = \{f \in L_\infty((\mathbb{R}^d)) : \sup_{t>0}(t^{-\alpha} w_{(r,L_1)}(f,t)) < \infty\}$, where w is the modulus of continuity. Then for any $\eta > 0$, $d/(d+\tau) < p < 1$, $\tau > 0$, and the parameter γ_n for the Gaussian Kernel*

$$R_\Phi(f_{opt}) - R_\Phi(\hat{f}_n) = \mathcal{O}_p\bigg(\Big(\frac{1}{n}\Big)^{1/(4+3d/\alpha)}\bigg)$$

where Φ can be either $\Phi_1(\cdot) = I(\cdot > 0)$ or $\Phi_2(\cdot) = (\cdot)_+$.

A.11 A Proof of Theorem 5.3.3

Proof. According to theorem 5.3.1 and theorem 5.3.2, we have

$$\begin{aligned} R_\Phi(\hat{f}_L) - R_\Phi(f_{opt}) &\leq R_{\Phi,\epsilon_n}(\hat{f}_L) - R_{\Phi,\epsilon_n}(f_{L,opt}) + 2 * C\epsilon_n \\ &\leq R_{\Phi,\epsilon_n}(\hat{f}_L) - R_{\Phi,\epsilon_n}(f_L^*) + 2 * C\epsilon_n \\ &\leq \Big(\lambda_n ||\hat{f}_L||_k^2 + R_{\Phi,\epsilon_n}(\hat{f}_L) - R^*_{\Phi,\epsilon_L}\Big) + \Big(2 * C\epsilon_n\Big) \\ &= (I) + (II) \end{aligned} \tag{A.9}$$

where f_L^* is the minimizer of $R_{\Phi,\epsilon_n}(\cdot)$ and $R^*_{\Phi,\epsilon_n} = R_{\Phi,\epsilon_n}(f_L^*)$.

To bound (A.9), we need to first bound (I). As is in Chen et al. (2016), we refer to theorem 7.23 in Steinwart and Christmann (2008). In order to use the oracle inequality in the theorem, there are four conditions to be satisfied.

- (B1) The loss function $L(\cdot)$ has a supremum bound $L(\cdot) \leq B$ for a constant $B > 0$.
- (B2) The loss function $L(\cdot)$ is locally Lipschitz continuous loss that can be clipped at a constant $M > 0$ such that $\tilde{f} = I(|f| \leq M)f + I(|f| \geq M)M$.
- (B3) The variance bound $E_P\big(L \circ \tilde{f} - L \circ f^*_{L,P}\big)^2 \leq V\Big(E_P\big(L \circ \tilde{-} L \circ f^*_{L,P}\big)\Big)^v$ is satisfied for constant $v \in [0,1]$, $V \geq B^{2-v}$ and all $f \in \mathcal{H}$.
- (B4) For fixed $n \geq 1$, there exist constant $p \in (0,1)$ and $a \geq B$ such that the entropy number $E_{D_X \sim P_X^n} e_i(id : \mathcal{H} \to L_2(D_X)) \leq a i^{-\frac{1}{2p}}$, $i \leq 1$.

To verify the condition (B1), recall the relaxed loss function of the lower bound of the one-sided interval is as follows.

$$\begin{aligned} L_{\Phi,\epsilon}(X, A, f_L(X)) = \frac{1}{P(A \,|\, X)}\Big(\big(1 - \alpha(X)\big)\Phi\big(Y - S(X)\big))\Psi_\epsilon(f_L(X), A) \\ + \alpha(X)\Phi\big(S(X) - Y\big))\Psi_\epsilon(A, f_L(X))\Big) \end{aligned}$$

Assuming the inverse probability $\frac{1}{P(A\,|\,X)}$ is bounded by a constant B, then $L_{I,\epsilon}(X, A, f_L(X))$ is naturally bounded by B for the PDI case. For the EDI case, when Y is bounded, and $S(x)$ is properly chosen, $L_{h,\epsilon}(X, A, f_L(X))$ is bounded by $B = \max\{\frac{1}{P(A\,|\,X)}, \max|Y - S(X)|\}$. To simplify notations, we say the loss function $L_{\Phi,\epsilon}(X, A, f_L(X))$ is bounded by B.

To verify condition (B2), similar to (B1), $L_{\Phi,\epsilon}(X, A, f_L(X))$ is Lipschitz continuous with a Lipschitz constant B. Also, f_L can be clipped to have a smaller risk. Suppose that M is a upper bound of absolute value of a reasonable range of dose, say, $M = \max\{|a_L|, |a_U|\}$, and $\tilde{f}_L = I(|f_L| \leq M)f_L + I(|f_L| \geq M)M$. It naturally follows that $R_\phi(\tilde{f}_L) \leq R_\Phi(f_L)$, since any unreasonable large dose recommendations introduce larger risk.

The condition (B3) is satisfied with $v = 0$ and $V = 4B^2$, because

$$E\big[\big(L_{\Phi,\epsilon} \circ \tilde{f}_L - L_{\Phi,\epsilon} \circ f_L^*\big)^2\big] \leq 2E\big[\big(L_{\Phi,\epsilon} \circ \tilde{f}_L\big)^2 + \big(L_{\Phi,\epsilon} \circ f_L^*\big)^2\big] \leq 4B^2$$

The Gaussian Kernel that are used in Chapter 6 and Chapter 7 is one type of benign Kernel. According to theorem 7.34 of Steinwart and Christmann (2008), (B4) is satisfied with the constant $a = c_{\epsilon,p}\gamma_n^{-\frac{(1-p)(1+\zeta)d}{2p}}$, where $d/(d+\tau) < p < 1$ and $\zeta > 0$ are two constant.

Since condition (B1)-(B4) are satisfied in our case, applying Theorem 7.23 from Steinwart and Christmann (2008) yields

$$(I) \leq 9 * \left(\lambda_n||f_L^0||_k^2 + R_{\Phi,\epsilon_n}(f_L^0) - R^*_{\Phi,\epsilon_n}\right) + K_0\left[\frac{1}{\gamma_n^{(1-p)(1+\zeta)d}\lambda_n^p n}\right]^{\frac{1}{2-p}} + 36\sqrt{2}B(\frac{\tau}{n})^{\frac{1}{2}} + \frac{15B\tau}{n} \tag{A.10}$$

with probability $1 - 3e^{-\tau}$.

In order to bound $\left(\lambda_n||f_L^0||_k^2 + R_{\Phi,\epsilon_n}(f_L^0) - R^*_{\Phi,\epsilon_n}\right)$, the approximation error, we refer to Section 2 in Eberts et al. (2013). As (A.10) holds for any $f_L^0 \in \mathcal{H}_\gamma$, we could construct a specific f_L^0 to facilitate the proof. First, as (8) in Eberts et al. (2013), we define a function $\mathcal{K}(x) = \sum_{j=1}^r \binom{r}{j}(-1)^{1-j}\frac{1}{j^d}(\frac{2}{\gamma^2})^{d/2}\mathcal{K}_{j\gamma/\sqrt{2}}(x)$ where $\mathcal{K}_\gamma = exp(-\gamma^2||x||_2^2)$ for all $x \in R^d$, and subsequently define f_L^0 via convolution.

$$f_L^0 = \mathcal{K} * f_L^* = \int_{R^d} \mathcal{K}(X-t) f_L^*(t)dt, \;\; x \in R^d$$

If we assume $f_L^* \in L_2(R^d) \cap L_\infty(R^d)$, then from theorem 2.3 in Eberts et al. (2013), it can be shown that $f_L^0 \in \mathcal{H}_\gamma$, where $\mathcal{H}_\gamma$ is the RKHS of Gaussian RBF Kernel. In addition, theorem 2.2 provides the upper bound for the access risk of f_L^0, which can be incorporated to have the results below.

$$\begin{aligned}
&\lambda_n ||f_L^0||_k^2 + R_{\Phi,\epsilon_n}(f_L^0) - R_{\Phi,\epsilon_n}^* \\
&= \lambda_n ||\mathcal{K} * f_L^*||_k^2 + R_{\Phi,\epsilon_n}(\mathcal{K} * f_L^*) - R_{\Phi,\epsilon_n}^* \\
&\leq \lambda_n (\gamma_n \sqrt{\pi})^{-d} (2^r - 1)^2 ||f_L^*||_{L_2(R^d)}^2 + R_{\Phi,\epsilon_n}(\mathcal{K} * f_L^*) - R_{\Phi,\epsilon_n}^* \\
&\leq \lambda_n (\gamma_n \sqrt{\pi})^{-d} (2^r - 1)^2 ||f_L^*||_{L_2(R^d)}^2 + B|\mathcal{K} * f_L^* - f_L^*|_{L_1(P_X)} \\
&\leq \lambda_n (\gamma_n \sqrt{\pi})^{-d} (2^r - 1)^2 ||f_L^*||_{L_2(R^d)}^2 + BC_{r,1}||g||_{L_p(P_X)} \omega_{r,L_1(R^d)}(f_L^*, \gamma_n/2)
\end{aligned} \tag{A.11}$$

Given the fact that we assume $f_{opt} \in B_{1,\infty}^\delta(R^d)$, if we are willing to further assume $f_L^* \in B_{1,\infty}^\delta(R^d)$, i.e., $B_{1,\infty}^\delta(R^d) = \{f \in L_\infty(R^d) : \sup_{t>0}(t^{-\delta}\omega_{r,L_1(R^d)})(f,t) < \infty\}$, then $w_{r,L_1}(R^d)(f_L^*, \gamma_n/2) < c_0 \gamma^\delta$, where c_0 is a constant. Plugging in this into (A.11), we obtain

$$\lambda_n ||f_L^0||_k^2 + R_{\Phi,\epsilon_n}(f_L^0) - R_{\Phi,\epsilon_n}^* \leq c_1 \lambda_n \gamma_n^{-d} + c_2 \gamma_n^\delta$$

After Combining all the parts together,

$$R_\Phi(\hat{f}_L) - R_\Phi(f_{opt}) \leq c_1 \frac{\lambda_n}{\gamma_n^d} + c_2 \gamma_n^\delta + c_3 \frac{1}{\gamma_n^{\frac{(1-p)(1+\zeta)d}{2-p}} \lambda_n^{\frac{p}{2-p}} n^{\frac{1}{2-p}}} + c_4 \frac{\tau^{\frac{1}{2}}}{n^{\frac{1}{2}}} + c_5 \frac{\tau}{n} + c_6 \epsilon_n$$

By properly choosing the constant, the desired rate can be attained.

$$\gamma_n \propto \left(\frac{1}{n}\right)^{\frac{1}{2\delta+d}} \qquad \lambda_n \propto \left(\frac{1}{n}\right)^{\frac{\delta+d}{2\delta+d}} \qquad \epsilon_n = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right)$$

We have the convergence rate:

$$R_\Phi(\hat{f}_L) - R_\Phi(f_{opt}) = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right)$$

□

A.12 A Proof of Corollary 5.3.3

Proof. Proving the convergence rate of the two-sided IDI estimators can be heavily involved. To the best of our knowledge, there is no well-established convergence theory for algorithms like DC which use an iterative procedure to tackle non-convex objective functions, especially when there are two functions estimated from the empirical risk minimization (ERM). To limit the scope of this dissertation, we instead provide a convergence rate by borrowing the results from 5.3.3, and replace the convergence of the two-sided IDI estimator with two estimators of one-sided IDIs. In order to do so, several extra assumptions have to be made. It is possible that better rates under weaker the assumptions can be achieved in the future.

The proof of corollary 5.3.3 consists of two steps. In the first step, we assume that there exists a function or a consistent estimator of a function that separate the dataset into a left subset and a right subset, in each of which the Partial Monotonicity assumption in (3.1) and (3.3) is satisfied. In the second step, the integrated access risk of the two-sided IDI estimator is bounded by the sum of two access risks of two independent one-sided IDI estimators. The detailed proof will be documented below.

Suppose there is a function $f_{M,opt}(x)$, such that $f_{L,opt}(x) < f_{M,opt}(x) < f_{U,opt}(x)$ for all $x \in \mathcal{X}$. Such a function is more likely to be known when the outcome is largely determined by the dose. For instance, it might be able to separate the left subset and the right subset by a single dose value, so that the lower bounds will be contained to only the left side of the value and the upper bounds to only the right side. Our simulation in Section 6.1 is qualified as an example where the optimal lower bounds are all below zero and the optimal upper bounds are all above zero. Thus, we can let $f_{M,opt}(x) = 0$. However, this is not always the

case in practice. Both the signal-to-noise ratio and the distribution of lower/upper bounds might make it impossible for visual separation. Therefore, we will discuss the cases where $f_{M,opt}(x)$ is estimated from data.

One possible choice of $f_{M,opt}$ is $f_{\max,opt} = \arg\max_f E\big[Y|A = f(X), X\big]$, the optimal individual dose rule (IDR) in Chen et al. (2016). Theoretical properties of the estimator $\hat{f}_{\max}$ have been established in the original paper and also discussed in Luedtke and van der Laan (2016). Despite the fact that it is suggested in Luedtke and van der Laan (2016) that the OWL-based estimator might convergence in nearly $\mathcal{O}(n^{-1/2})$, the rate established so far are based on the risk (or value function) instead of the rule $f_{\max,opt}$ itself due to major obstacles imposed by the DC algorithm. In order to facilitate the proof, we make the following assumption for the estimated optimal dose function $\hat{f}_{\max}$.

$$f_{L,opt}(x) < \hat{f}_{\max}(x) < f_{U,opt}(x), \quad \forall x \in \mathcal{X} \tag{A.12}$$

In the second step, two censored datasets are defined where the lower bound function and the upper bound function can be estimated individually, assuming there is a $f_{M,opt}$, known or estimated, such that $f_{L,opt}(x) < f_{M,opt}(x) < f_{U,opt}(x)$ for all $x \in \mathcal{X}$. Denote the original dataset without censoring as $\mathcal{D} = \{\big(X_i, A_i, Y_i\big),\ 1 \le i \le n\}$ and then we can subsequently define $\mathcal{D}_{\mathcal{L}} = \{\big(X_i, A_i, Y_i\big),\ \forall i\ s.t.\ A_i \le f_{M,opt}(X_i)\}$ and $\mathcal{D}_{\mathcal{U}} = \{\big(X_i, A_i, Y_i\big),\ \forall i\ s.t.\ A_i > f_{M,opt}(X_i)\}$. These two subsets are independent as there is no shared observations.

Now define $\hat{f}'_L = \min_{f_L} \hat{R}_{\Phi,\epsilon}(f_L)$ on dataset $\mathcal{D}_{\mathcal{L}}$ and $\hat{f}'_U = \min_{f_U} \hat{R}_{\Phi,\epsilon}(f_U)$ on dataset $\mathcal{D}_{\mathcal{U}}$, while $\big\{\hat{f}_L, \hat{f}_U\big\} = \min_{f_L, f_U} \hat{R}_{\Phi,\epsilon}(f_L, f_U)$ on the dataset of $\mathcal{D}$. Similarly, define $f'_{L,*} = \min_{f_L} R_{\Phi,\epsilon}(f_L)$ on the censored population of $\mathcal{D}_{\mathcal{L}}$ and $f'_{U,*} = \min_{f_U} R_{\Phi,\epsilon}(f_U)$ on the censored population of $\mathcal{D}_{\mathcal{U}}$, while $\big\{f_{L,*}, f_{U,*}\big\} = \min_{f_L, f_U} R_{\Phi,\epsilon}(f_L, f_U)$ on the original population. To simplify the problem, we assume that $f_{L,*}(x) - f_{U,*}(x) \le -2\epsilon$, since this inequality will hold when $\epsilon \to 0$ as $n \to 0$. By the definitions of $\mathcal{D}_{\mathcal{L}}$, $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}$, we have $\big\{f'_{L,*}, f'_{U,*}\big\} = \big\{f_{L,*}, f_{U,*}\big\}$, i.e., the population estimates of the lower and upper bounds estimated jointly are the same as the population estimates of the lower and upper bounds

estimated individually while assuming the other part is known. In addition, we make the following assumption of the uniform convergence of the empirical risk functions.

$$\sup_{f_L, f_U \in \mathcal{H}} \left| \hat{R}_{\Phi,\epsilon}(f_L, f_U) - R_{\Phi,\epsilon}(f_L, f_U) \right| \leq \kappa \sqrt{\frac{\log(n)}{n}} \text{ with probability } p' \tag{A.13}$$

where κ is a constant related to p' and the complexity of the RKHS. Without proving the rate in A.13, we believe it is consistent with the general uniform convergence rate in Vapnik (1998) where only one functions is estimated through ERM.

We will show how the loss function of the two-sided IDI can be decomposed as the sum of two one-sided interval losses plus a quantity not related to the dose. Recall the loss function for the lower bound of the one-sided interval:

$$L_{\Phi,\epsilon}(X, A, f_L(X)) = \frac{1}{P(A \mid X)} \Big(\big(1 - \alpha(X)\big) \Phi\big(Y - S(X)\big)) \Psi_\epsilon(f_L(X), A) + \alpha(X) \Phi\big(S(X) - Y\big)) \Psi_\epsilon(A, f_L(X)) \Big)$$

And the loss function for the upper bound of the one-sided interval:

$$L_{\Phi,\epsilon}(X, A, f_U(X)) = \frac{1}{P(A \mid X)} \Big(\big(1 - \alpha(X)\big) \Phi\big(Y - S(X)\big)) \Psi_\epsilon(A, f_U(X)) + \alpha(X) \Phi\big(S(X) - Y\big)) \Psi_\epsilon(f_U(X), A) \Big)$$

And the loss function for the two-sided interval:

$$L_{\Phi,\epsilon}(X, A, f_L(X), f_U(X)) = \frac{1}{P(A \mid X)} \Big(\big(1 - \alpha(X)\big) \Phi\big(Y - S(X)\big)) \Psi_\epsilon^{out}(f_L(X), A, f_U(X)) + \alpha(X) \Phi\big(S(X) - Y\big)) \Psi_\epsilon^{in}(f_L(X), A, f_U(X)) \Big)$$

It follows that the difference between two-sided loss and the sum of one-sided losses is a constant which does not contain the bound functions, under the assumption that $f_{L,*}(x) -$

$f_{U,*}(x) \leq -2\epsilon$.

$$
\begin{aligned}
&L_{\Phi,\epsilon}(X, A, f_L(X), f_U(X)) - L_{\Phi,\epsilon}(X, A, f_L(X)) - L_{\Phi,\epsilon}(X, A, f_U(X)) \\
&= \frac{1}{P(A \mid X)}\Big(\big(1-\alpha(X)\big)\Phi\big(Y - S(X)\big)\Psi_\epsilon(A, f_U(X)) \\
&\qquad + \alpha(X)\Phi\big(S(X) - Y\big)\big(\Psi_\epsilon(f_U(X), A) - 1\big)\Big) - L_{\Phi,\epsilon}(X, A, f_U(X)) \\
&= \frac{-1}{P(A \mid X)}\Big(\alpha(X)\Phi\big(S(X) - Y\big)\Big)
\end{aligned}
$$

Combined with corollary 5.3.1 and corollary 5.3.2, we have

$$
\begin{aligned}
R_\Phi(\hat{f}_L, \hat{f}_U) - R_\Phi(f_{L,opt}, f_{U,opt}) &\leq R_{\Phi,\epsilon_n}(\hat{f}_L, \hat{f}_U) - R_{\Phi,\epsilon_n}(f_{L,opt}, f_{U,opt}) + 2C\epsilon_n \\
&\leq R_{\Phi,\epsilon_n}(\hat{f}_L, \hat{f}_U) - R_{\Phi,\epsilon_n}(f_{L,*}, f_{U,*}) + 2C\epsilon_n \\
&\leq R_{\Phi,\epsilon_n}(\hat{f}'_L, \hat{f}'_U) - R_{\Phi,\epsilon_n}(f_{L,*}, f_{U,*}) + 2C\epsilon_n + 2\kappa\sqrt{\frac{\log(n)}{n}} \\
&= R_{\Phi,\epsilon_n}(\hat{f}'_L) - R_{\Phi,\epsilon_n}(f'_{L,*}) + R_{\Phi,\epsilon_n}(\hat{f}'_U) - R_{\Phi,\epsilon_n}(f'_{U,*}) + 2C\epsilon_n + 2\kappa\sqrt{\frac{\log(n)}{n}}
\end{aligned}
\tag{A.14}
$$

with probability p' as $n \to \infty$. The probability p' comes from the inequality (A.13). The first inequality of A.14 is guaranteed by theorem 5.2.1 and theorem 5.2.2. The second inequality is from the definition of $f_{L,*}$ and $f_{U,*}$. The third inequality is from the assumption A.13. The coefficient is $2 \times \kappa$ because the uniform bound has to be applied twice as the definitions of $\hat{f}_L$ and $\hat{f}_U$ guarantee $\hat{R}_{\Phi,\epsilon_n}(\hat{f}_L, \hat{f}_U) < \hat{R}_{\Phi,\epsilon_n}(\hat{f}'_L, \hat{f}'_U)$ but not $R_{\Phi,\epsilon_n}(\hat{f}_L, \hat{f}_U) < R_{\Phi,\epsilon_n}(\hat{f}'_L, \hat{f}'_U)$. The last equality is from the decomposition of the two-sided loss and the equivalence of $\{f'_{L,*}, f'_{U,*}\}$ and $\{f_{L,*}, f_{U,*}\}$.

The theorem follows by applying (A.10) twice, because $\hat{f}'_L$ and $\hat{f}'_U$ are independently estimated on two non-overlapping datasets $\mathcal{D}_\mathcal{L}$ and $\mathcal{D}_\mathcal{U}$ and, therefore, $R_{\Phi,\epsilon_n}(\hat{f}'_L) - R_{\Phi,\epsilon_n}(f'_{L,*})$ is independent with $R_{\Phi,\epsilon_n}(\hat{f}'_U) - R_{\Phi,\epsilon_n}(f'_{U,*})$. The uniform bound (A.13) and theorem 5.3.3 hold for the same time with a probability of at least $p' + (1 - 3e^{-\tau})^2 - 1$. Here ϵ_n only shows up in the term $2C\epsilon_n$. Hence by choosing $\epsilon_n = \mathcal{O}\big(n^{-\frac{\delta}{2\delta+d}}\big)$, it will not impact the conver-

gence rate. The term $2\kappa\sqrt{\frac{\log(n)}{n}}$, which corresponds to the convergence rate of $\mathcal{O}\big((\frac{\log n}{n})^{\frac{1}{2}}\big)$, is dominated by the entire (A.14) is dominated by the first two risk terms. Therefore, the same convergence rate in theorem 5.3.3 can be achieved in corollary 5.3.3. $\square$