Journal of Machine Learning Research 27 (2026) 1-113 Submitted 11/23; Revised 10/26; Published 10/26

Causal Inference under Unmeasured Confounding with Negative Controls: A Minimax Learning Approach

Nathan Kallus* KALLUS@CORNELL.EDU
Cornell Tech
Cornell University
2 West Loop Rd, NY 10044, USA

Xiaojie Mao* MAOXJ@SEM.TSINGHUA.EDU.CN
School of Economics and Management
Research Center for Contemporary Management
Tsinghua University
Beijing, 100084, China

Masatoshi Uehara* MU223@CORNELL.EDU
Cornell Tech
Cornell University
2 West Loop Rd, NY 10044, USA

Zhiheng Zhang* ZHANGZHIHENG@MAIL.SHUFE.EDU.CN
School of Statistics and Data Science
Institute of Data Science and Statistics
Institute of Big Data Research
Shanghai University of Finance and Economics
Shanghai, 200433, China

Editor: Joris Mooij

Abstract

We study the estimation of causal parameters when not all confounders are observed and instead negative controls are available. Recent work has shown how these can enable identification and efficient estimation via two so-called bridge functions. In this paper, we tackle the primary challenges in causal inference with negative controls: identifying and estimating these bridge functions. Previous work has typically relied on completeness conditions for identification, imposed uniqueness assumptions for estimation, and focused on parametric models for the bridge functions. Instead, we provide an alternative identification strategy that assumes the existence of two bridge functions rather than completeness conditions. We also provide new estimators for these functions based on minimax learning formulations. These estimators accommodate general function classes such as Reproducing Kernel Hilbert Spaces and neural networks. We study finite-sample convergence results both for estimating bridge functions themselves and for the final estimation of the causal parameter under a variety of combinations of assumptions. We avoid uniqueness conditions on the bridge functions as much as possible.

Keywords: double robustness, minimax estimation, policy evaluation, treatment effect estimation, unmeasured confounding

∗. Authors are listed in alphabetical order.

©2026 Nathan Kallus, Xiaojie Mao, Masatoshi Uehara, and Zhiheng Zhang.

License: CC-BY 4.0, see https://creativecommons.org/licenses/by/4.0/. Attribution requirements are provided at http://jmlr.org/papers/v27/23-1535.html.

1. Introduction

Causal inference from observational data is a necessity in many fields where experimentation and randomization is limited. Even when experimentation is feasible, observational data can help support initial or supplementary investigations. Compared to experimental-intervention data, the key difficulty with observational data is confounding or endogeneity: correlations between observed actions A and outcomes Y that are not due to a causal relationship, as might be induced by common causes such as a healthy lifestyle leading to both selection into a pharmaceutical therapy and good health outcomes regardless of therapy. A common identification strategy is to control for many baseline covariates and assume they fully account for all such common causes, termed unconfoundedness, ignorability, or exchangeability. However, in practice, it is dubious that all confounders are ever truly accounted for, casting doubt on any resulting conclusion.

When some confounders are unobserved, an alternative identification strategy is to use *negative controls* (Miao et al., 2018a,b; Tchetgen et al., 2020; Cui et al., 2023; Deaner, 2021; Shi et al., 2020), which play a similar role as instrumental variables. Negative controls are observed covariates that have a more restricted relationship with the action and outcome: *negative control actions* do not directly impact the outcome of interest and *negative control outcomes* are not directly impacted by either the negative control actions or the primary action of interest. See Figure 1 for a typical causal diagram. When these negative control variables are sufficiently informative about the unmeasured confounders, there exist the so-called *bridge functions* that enable identification and estimation of causal quantities. These bridge functions are analogues to the propensity score functions and outcome regression functions one would use if all confounders were observed.

Learning these bridge functions, however, is a nontrivial task, as it no longer amounts to a regression problem as in the unconfounded case, which can be outsourced to standard machine learning methods. Many previous works focus on parametric settings to simplify the estimation of bridge functions. For example, Cui et al. (2023) recently proposed a doubly robust approach that may attain the semiparametric efficiency lower bound, but they focused on parametric estimates, which may be restrictive in practice. Moreover, their analysis for the estimation relied on the bridge functions' *unique* existence (although causal-effect identification itself does not require uniqueness), which may be dubious in practice and even refutable in many examples. It also requires certain completeness assumptions to achieve identification. Shi et al. (2020) studied the setting with discrete variables of finitely many possible values (see Example 3 below), so bridge functions are naturally parametric. Although they briefly mention two possible approaches to tackle nonunique bridge functions, these approaches are not fully developed, and the estimation procedure ultimately focuses on unique bridge functions. Completeness and uniqueness conditions are similarly required in other prior work (Miao et al., 2018a,b; Tchetgen et al., 2020; Deaner, 2021). In this paper, we propose an alternative identification strategy without completeness conditions and introduce new minimax estimators for the bridge functions that are amenable to general function approximation beyond parametric models. We catalog a variety of settings in which the functions are learnable and estimation with negative controls is practically feasible. In all estimation settings, we avoid assuming uniqueness of bridge functions as much as possible.

Our contributions are:

- We introduce a novel identification strategy, which requires the existence of two different types of bridge functions, but not any completeness condition. This provides an alternative to the previous identification strategies based on the existence of at least one type of bridge functions and some completeness conditions. Our identification assumptions are in general different from the existing ones, but in the discrete data setting and in some linear model settings they are strictly weaker than the existing identification assumptions. Our identification strategy complements the approaches in the existing literature, broadening the scope of causal identification with negative controls. We also consider a more general setting than average-effect estimation, where we allow actions to be possibly stochastic and the action space to be possibly continuous.

- We propose new estimators for bridge functions (even if nonunique) by introducing an adversarial critic function and formulating the learning problem as a minimax game. Our minimax approach accommodates a wide range of flexible function classes, including reproducing kernel Hilbert spaces (RKHSs) and neural networks. Then, by plugging these bridge function estimators into different estimating equations, we derive our estimators.

- We provide finite-sample convergence results under a variety of different assumptions, generally without assuming unique bridge functions (see Table 1). One important assumption is the well-specification of the hypothesis classes, which we call *realizability*. Another assumption, which we call *closedness*, ensures that the critic function classes in our minimax estimators are sufficiently rich. Depending on how much of these we are willing to assume, we obtain different convergence rates for the final estimator. One surprising result is that when we assume realizability of both bridge and critic function classes, our estimator is consistent even though the bridge functions themselves are not consistently estimated (see row I in Table 1). This result generally cannot be obtained by standard analyses in the instrumental variable literature (e.g., Ai and Chen, 2003; Newey and Powell, 2003; Dikkala et al., 2020). Thus, our finite-sample analysis exploits the special structure of double negative controls rather than simply importing existing instrumental-variables analyses.

The rest of the paper is organized as follows. Section 2 introduces our setup. Section 3 presents our new identification result and compares it with previous work. Section 4 introduces our estimation method. In Section 5–7, we present finite-sample results for the proposed estimators. These results are summarized at the end of Section 4. In Section 8, we review related literature. In Section 9, we present numerical results from simulation studies and a real-data analysis. In Appendix B, we compare our proposed identification strategy with existing ones in detail. All proofs are deferred to the appendix.

2. Setup

We consider an action $A \in \mathcal{A}$ that can be discrete or continuous. We associate $\mathcal{A}$ with a base measure μ, *e.g.*, the counting measure if $\mathcal{A}$ is finite or Lebesgue measure if $\mathcal{A}$ is continuous. Let $Y(a)$ denote the real-valued counterfactual outcome that would be observed if the action

were set to $a \in \mathcal{A}$ and $Y = Y(A)$ be the observed outcome corresponding to the actually observed action. Moreover, let $X \in \mathcal{X} \subseteq \mathbb{R}^d$ be a collection of observed covariates. For a given contrast function $\pi : \mathcal{A} \times \mathcal{X} \mapsto \mathbb{R}$, we are interested in estimating the *generalized average causal effect* (GACE):

$$J = \mathbb{E}\left[\int Y(a)\pi(a \mid X)\mathrm{d}\mu(a)\right]. \tag{1}$$

Example 1 (Average treatment effect)**.** Consider $\mathcal{A} = \{0, 1\}$. The counterfactual mean parameter $\mathbb{E}\left[Y(a_0)\right]$ for $a_0 \in \{0, 1\}$ is an example of Equation (1) with $\pi\left(a \mid x\right) = \mathbb{I}\left[a = a_0\right]$. If we are interested in the effect of "treatment" $A = 1$ compared to "control" $A = 0$, we can let $\pi(a \mid x) = 2a - 1$ and obtain $J = \mathbb{E}[Y(1) - Y(0)]$, the average treatment effect (ATE).

Example 2 (Policy evaluation)**.** If $\pi(a \mid x)$ is a density on $\mathcal{A}$ for each x with respect to (w.r.t.) μ, then J is the average outcome we experience when we follow the policy that assigns an action drawn from $\pi(\cdot \mid X)$ for an individual with covariates X (*e.g.*, Dudik et al., 2014; Tian, 2008). The measure μ can be the Lebesgue measure when the action space is continuous or the counting measure when the action space is discrete.

We *do not* assume that the observed covariates X include all confounders that affect both the action and the potential outcomes, and instead there exist some *unmeasured* confounders $U \in \mathcal{U} \subseteq \mathbb{R}^{p_u}$ (discrete, continuous, or mixed):

$$Y(a) \not\perp A \mid X, \quad \text{but} \quad Y(a) \perp A \mid U, X.$$

If U were observed, we could identify the GACE J simply by controlling for both X, U. However, in this paper we assume that confounders U *cannot* be observed, in which case, the GACE J is generally *unidentifiable* from the distribution of the observed variables (Y, X, A) alone. To overcome the challenge of unmeasured confounding, in this paper we employ the negative control framework proposed in Miao et al. (2018a); Deaner (2021); Cui et al. (2023). This framework involves two additional types of observed variables: negative control actions $Z \in \mathcal{Z} \subseteq \mathbb{R}^{p_z}$ and negative control outcomes $W \in \mathcal{W} \subseteq \mathbb{R}^{p_w}$, which can be discrete, continuous, or mixed. These variables are called "negative" controls due to the assumed absence of certain causal effects: negative control actions cannot directly affect the outcome Y, and neither the negative control actions Z nor the primary action A can affect the negative control outcomes W. At the same time, these variables remain useful controls because they are related to the unmeasured confounders. We can, in a sense, view them as proxies for the unmeasured confounders U.

Figure 1 shows a typical causal diagram for this setting. To formalize our setting and allow for more generality, however, we will employ potential outcome notation. Let $Y(a, z)$ and $W(a, z)$ denote the corresponding counterfactual outcomes one would observe had the primary action and negative control action taken value $(a, z) \in \mathcal{A} \times \mathcal{Z}$. We then formalize the negative control assumptions as follows.

Assumption 1 (Negative Controls)**.** 1. Consistency: $Y = Y(A, Z), W = W(A, Z)$.

2. Negative control actions: $Y(a, z) = Y(a)$, $\forall a \in \mathcal{A}$.

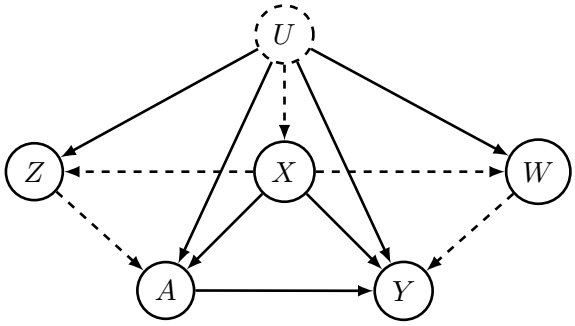


Figure 1: A typical causal diagram for negative controls. The dashed edges may be absent, and the dashed circle around U indicates that U is unobserved.

3. Negative control outcomes: $W(a,z) = W$, $\forall a \in \mathcal{A}, z \in \mathcal{Z}$.
4. Latent unconfoundedness: $(A, Z) \perp (Y(a), W) \mid U, X, \quad \forall a \in \mathcal{A}$.
5. Overlap: $|\pi(a|x)/f(a|x,u)| < \infty, \quad \forall a \in \mathcal{A}, x \in \mathcal{X}, u \in \mathcal{U}$.

Here, condition 1 states that the observed outcome Y is the realization of the potential outcome $Y(a)$ for units with $A = a$. This is a standard assumption in causal inference (Imbens and Rubin, 2015). Conditions 2 and 3 express the definition of negative controls in potential-outcome notation: the negative control action Z cannot affect the primary outcome Y, and the negative control outcome W cannot be affected by either the primary action A or the negative control actions Z. Condition 4 formalizes the assumption that the unmeasured variables U capture all common causes of (A, Z) and (Y, W) not included in X. Condition 4 may also be satisfied by causal diagrams other than Figure 1 (see table A.1 of Tchetgen et al., 2020). Condition 5 requires sufficient overlap between the contrast function π and the distribution of observed actions given both observed and unobserved confounders. This is a canonical assumption in causal inference and policy evaluation.

Our data consist of n independent and identically distributed (iid) observations of (Z, X, W, A, Y). Crucially, U is *unobserved*. Our aim is to estimate the GACE parameter J from these data.

Notation. We let $\mathbb{E}$ denote expectations w.r.t. $O = (Z, X, W, A, Y)$, and $\mathbb{E}_n$ denote empirical average over the n observations thereof. For a function g of (z, x, w, a, y) (or a subset thereof) we write g to mean the random variable $g(O)$. We let $L_2(O)$ denote the space of square-integrable functions of O. Similarly, $L_2(W, A, X)$ and $L_2(Z, A, X)$ denote the space of square-integrable functions of just W, A, X and Z, A, X, respectively. For a function g, we let $\|g\|_2$ denote the norm in these spaces. For a vector θ, we let $\|\theta\|$ denote the Euclidean norm. We let $\|\cdot\|_\infty$ denote the sup norm of either a variable or function, and for a class of functions we let it denote the supremum of sup norms in the class. For subsets $\mathbb{A}, \mathbb{B}$ of a field, we define $\mathbb{A} + \mathbb{B} = \{a + b : a \in \mathbb{A}, b \in \mathbb{B}\}, \mathbb{A}\mathbb{B} = \{ab : a \in \mathbb{A}, b \in \mathbb{B}\}$. We call a set S symmetric if $-s \in S$ for any $s \in S$, and we call S star-shaped (around the origin) if $\alpha s \in S$ for any $s \in S, \alpha \in [0, 1]$. Finally, we often use $O(\cdot)$ notation to denote rates w.r.t. the sample size n, unless otherwise specifically indicated. For a matrix A, we denote its Moore Penrose inverse by A^+.

3. Identifying GACE via Bridge Functions

We begin with the ideal unconfounded setting and then present identification results using negative controls.

3.1 The Ideal Unconfounded Setting

If the unobserved confounders U *were* observed, then the GACE could be identified, that is, it can be written as a function of the distribution of (Y, A, X, U). To illustrate this, define the regression function $k_0(a, u, x) = \mathbb{E}[Y \mid A = a, U = u, X = x]$, and define the generalized propensity score $f(a \mid u, x)$ as the conditional density of the distribution $A \mid U, X$ relative to the base measure μ (Hirano and Imbens, 2004). Based on these two functions, the following lemma shows the identification of J if U *were* observed.

Lemma 1. *If* $Y(a) \perp A \mid U, X$ *and* $|\pi(A \mid X)/f(A \mid X, U)| < \infty$, *then*

$$J = \mathbb{E}\phi_{\mathrm{IPW}}(Y, A, U, X; f) = \mathbb{E}\phi_{\mathrm{REG}}(Y, A, U, X; k_0) = \mathbb{E}\phi_{\mathrm{DR}}(Y, A, U, X; k_0, f),$$

where $\phi_{\mathrm{IPW}}(y, a, u, x; f) = \frac{\pi(a|x)}{f(a|x,u)} y$, $\phi_{\mathrm{REG}}(y, a, u, x; k_0) = \int k_0(a', u, x)\pi(a'|x)\mathrm{d}\mu(a')$,

$$\phi_{\mathrm{DR}}(y, a, u, x; k_0, f) = \frac{\pi(a|x)}{f(a|x,u)}\left(y - k_0(a, u, x)\right) + \int k_0(a', u, x)\pi(a'|x)\mathrm{d}\mu(a').$$

Lemma 1 suggests estimators for J if U were observed: we can first estimate the nuisance functions $k_0(a, u, x)$ and/or $f(a \mid u, x)$, and then estimate J by using any of the three estimating equations above with the estimated nuisance(s). The resulting three estimators are called the inverse propensity weighting (IPW) estimator, the regression-based (REG) estimator, and the doubly robust (DR) estimator, respectively (*e.g.*, Robins et al., 1994; Dudik et al., 2014).

3.2 The Negative-Control Setting

However, in this paper, we deal with the setting where U is *unobserved*, so estimators above are infeasible. In particular, neither $k_0(a, u, x)$ nor $f(a \mid u, x)$ can be identified. Instead, we can use their negative control analogues called *bridge functions* (Miao et al., 2018b; Cui et al., 2023).

Definition 1 (Bridge functions)**.** An outcome bridge function is $h_0 \in L_2(W, A, X)$ with

$$\mathbb{E}[h_0(W, A, X) \mid A, U, X] = k_0(A, U, X). \tag{2}$$

An action bridge function is q_0 with $\pi q_0 \in L_2(Z, A, X)$ and

$$\mathbb{E}[\pi(A \mid X) q_0(Z, A, X) \mid A, U, X] = \frac{\pi(A \mid X)}{f(A \mid U, X)}. \tag{3}$$

From Definition 1, we can observe that an outcome bridge function h_0 and an action bridge function q_0 can play a similar role as the regression function k_0 and the generalized propensity score f, respectively (see Lemma 2 below). These bridge functions are not necessarily unique. We thus define the sets of bridge functions as follows:

$$\begin{aligned}
\mathbb{H}_0 &= \{h \in L_2(W, A, X) : \mathbb{E}[Y - h(W, A, X) \mid A, U, X] = 0\}, \\
\mathbb{Q}_0 &= \{q : \pi q \in L_2(Z, A, X), \\
&\qquad \mathbb{E}[\pi(A \mid X)\left(q(Z, A, X) - 1/f(A \mid U, X)\right) \mid A, U, X] = 0\}.
\end{aligned} \tag{4}$$

The existence of bridge functions depends on the relationship between (Y, Z, W) and the unmeasured confounders U. Generally, such bridge functions exist when the negative control proxies Z, W are *sufficiently informative* for the unmeasured confounding, so that they can account for the impact of the unmeasured confounders U on the action A and the outcome Y, respectively.

Example 3 (Discrete setting)**.** We first consider a discrete data setting as in Miao et al. (2018a); Shi et al. (2020), where the variables W, Z, U are all discrete variables with values w_i, z_j, u_s for $i = 1, \ldots, |\mathcal{W}|, j = 1, \ldots, |\mathcal{Z}|, s = 1, \ldots, |\mathcal{U}|$. For any $(a, x) \in \mathcal{A} \times \mathcal{X}$, let $P(\mathbf{W} \mid \mathbf{U}, a, x)$ denote a $|\mathcal{W}| \times |\mathcal{U}|$ matrix whose (i, s)th element is $\mathbb{P}[W = w_i \mid U = u_s, A = a, X = x]$, $P(\mathbf{Z} \mid \mathbf{U}, a, x)$ a $|\mathcal{Z}| \times |\mathcal{U}|$ matrix whose (j, s)th element is $\mathbb{P}[Z = z_j \mid U = u_s, A = a, X = x]$, $\mathbb{E}[Y \mid \mathbf{U}, a, x]$ a $1 \times |\mathcal{U}|$ vector whose sth element is $\mathbb{E}[Y \mid U = u_s, A = a, X = x]$, $F(a \mid \mathbf{U}, x)$ a $|\mathcal{U}| \times |\mathcal{U}|$ diagonal matrix whose sth diagonal element is $f(a \mid u_s, x)$, and $\mathbf{e}$ an all-one column vector of length $|\mathcal{U}|$.

With these notations, Equations (2) and (3) translate into the following linear equation system:

$$\begin{aligned} h_0^\top(\mathbf{W}, a, x) P(\mathbf{W} \mid \mathbf{U}, a, x) &= \mathbb{E}[Y \mid \mathbf{U}, a, x], \\ q_0^\top(\mathbf{Z}, a, x) P(\mathbf{Z} \mid \mathbf{U}, a, x) F(a \mid \mathbf{U}, x) &= \mathbf{e}^\top. \end{aligned} \tag{5}$$

There exist solutions $h_0(\mathbf{W}, a, x)$ to the first equation if the vector $\mathbb{E}[Y \mid \mathbf{U}, a, x]$ belongs to the column space of the matrix $P(\mathbf{W} \mid \mathbf{U}, a, x)$. This requires that the proxy W can capture enough information in the unobserved confounder U, so that the column space of matrix $P(\mathbf{W} \mid \mathbf{U}, a, x)$ is rich enough to recover $\mathbb{E}[Y \mid \mathbf{U}, a, x]$. Similarly, there exist solutions $q_0(\mathbf{Z}, a, x)$ to the second equation if the proxy Z is informative enough such that the column space of $P(\mathbf{Z} \mid \mathbf{U}, a, x)$ can recover the vector $(1/f(a \mid u_1, x), \ldots, 1/f(a \mid u_{|\mathcal{U}|}, x))^\top$.

Clearly, these existence conditions would automatically hold if the matrices $P(\mathbf{W} \mid \mathbf{U}, a, x)$, $P(\mathbf{Z} \mid \mathbf{U}, a, x)$ have full column ranks (which implies that $|\mathcal{W}| \geq |\mathcal{U}|$ and $|\mathcal{Z}| \geq |\mathcal{U}|$). In this case, the proxies Z, W are both strongly informative for the unmeasured confounders U, such that they can handle arbitrary $\mathbb{E}[Y \mid \mathbf{U}, x]$ and $(1/f(a \mid u_1, x), \ldots, 1/f(a \mid u_{|\mathcal{U}|}, x))^\top$. However, even in this case, the solutions are generally nonunique when $|\mathcal{W}| > |\mathcal{U}|$ and $|\mathcal{Z}| > |\mathcal{U}|$. If we were to further assume that $|\mathcal{W}| = |\mathcal{Z}| = |\mathcal{U}|$, $P(\mathbf{W} \mid \mathbf{U}, a, x)$ and $P(\mathbf{Z} \mid \mathbf{U}, a, x)$ are invertible square matrices, then the bridge functions would uniquely exist.

Example 4 (Linear Model)**.** Suppose (Y, W, Z, A) are generated as follows:

$$\begin{aligned} Y &= \alpha_Y^\top U + \beta_Y^\top X + \gamma_Y^\top A + \omega_Y^\top W + \epsilon_Y, \\ Z &= \alpha_Z U + \beta_Z X + \gamma_Z A + \epsilon_Z, \\ W &= \alpha_W U + \beta_W X + \epsilon_W, \\ A &\sim \text{Unif}\left(\underline{\alpha}_A^\top U + \underline{\beta}_A^\top X, \overline{\alpha}_A^\top U + \overline{\beta}_A^\top X\right), \end{aligned}$$

where $\epsilon_Y, \epsilon_Z, \epsilon_W$ are independent mean-zero random noises that are also independent of (A, U, X). It can be shown that when there exist solutions θ_W, θ_Z to the equations $\theta_W^\top \alpha_W = \alpha_Y^\top$ and $\theta_Z^\top \alpha_Z = \overline{\alpha}_A^\top - \underline{\alpha}_A^\top$, then valid bridge functions h_0, q_0 exist (see Lemma 50):

$$h_0(W, A, X) = (\theta_W + \omega_Y)^\top W + \left(\beta_Y - \theta_W^\top \beta_W\right) X + \gamma_Y A, \quad \forall \theta_W \text{ s.t. } \theta_W^\top \alpha_W = \alpha_Y^\top,$$

$$q_0(Z, A, X) = \theta_Z^\top Z + \left(\overline{\beta}_A - \underline{\beta}_A - \theta_Z^\top \beta_Z\right) X - \theta_Z^\top \gamma_Z A, \quad \forall \theta_Z \text{ s.t. } \theta_Z^\top \alpha_Z = \overline{\alpha}_A^\top - \underline{\alpha}_A^\top.$$

The existence of solutions to equations $\theta_W^\top \alpha_W = \alpha_Y^\top$ and $\theta_Z^\top \alpha_Z = \overline{\alpha}_A^\top - \underline{\alpha}_A^\top$ means that the proxy variables Z, W capture relevant information in the unmeasured confounders U, such that the column spaces of loading matrices α_Z, α_W can cover the effects of the unmeasured confounders U on Y and A (characterized by the vectors α_Y and $\overline{\alpha}_A^\top - \underline{\alpha}_A^\top$ respectively).

In particular, if α_W, α_Z have full column rank (which implies $p_w \geq p_u$ and $p_z \geq p_u$), then the existence conditions always hold because the column spaces of α_W, α_Z can cover any α_Y and $\overline{\alpha}_A^\top - \underline{\alpha}_A^\top$, no matter how the unmeasured confounders U affect Y and A. Moreover, the bridge functions are nonunique if $p_w > p_u$ and $p_z > p_u$. In these cases, we have more proxy variables than necessary to account for the unmeasured confounders, so there exist many different ways to transform the proxy variables to tackle the confounding problem.

Example 5 (Nonparametric Model)**.** More generally, when the relevant variables can be continuous and we assume nonparametric models for the data generating process, we need to ensure the existence of solutions to conditional moment equations in Equations (2) and (3). The existence of bridge functions again requires the proxy variables Z, W to capture the relevant information in the unmeasured confounders U. They require that after conditioning on variables A, X, there exist some transformations (i.e., bridge functions) of Z, W whose information about the unmeasured confounders U can recover the confounding effects of U on the action A and outcome Y respectively.

For general nonparametric models, the conditional moment equations in Equations (2) and (3) correspond to equations defined by potentially infinite-dimensional and ill-posed linear operators. It is therefore difficult to explicitly characterize the existence of solutions to these equations. One set of sufficient (but not necessary) conditions that can ensure the existence of solutions to these equations is characterized by Picard's theorem in functional analysis. According to Picard's theorem, when the functions $k_0(A, U, X)$ and $1/f(A \mid U, X)$ satisfy the so-called Picard's condition (roughly meaning that they are sufficiently smooth relative to the equations), then the solutions exist when the following completeness conditions hold (see Lemmas 30 to 32 in Appendix C.2):

$$\mathbb{E}\left[g(U, A, X) \mid Z, A, X\right] = 0 \text{ only when } g(U, A, X) = 0, \tag{6}$$

$$\mathbb{E}\left[g(U, A, X) \mid W, A, X\right] = 0 \text{ only when } g(U, A, X) = 0. \tag{7}$$

These completeness conditions roughly mean that the negative controls Z, W can capture all variations of the unobserved confounders U, after conditioning on A, X.

In fact, when Z, W, U are all discrete variables as in Example 3, the completeness conditions are equivalent to the full column ranks of matrices $P(\mathbf{W} \mid \mathbf{U}, a, x)$, $P(\mathbf{Z} \mid \mathbf{U}, a, x)$. When the variables satisfy the linear models in Example 4, these completeness conditions often imply that the full column ranks of the matrices α_W, α_Z. As we discussed in Examples 3 and 4, these conditions are sufficient but not necessary conditions for the existence of bridge functions. It is possible that even when these matrices are rank deficient so that the completeness conditions fail, the equations in Examples 3 and 4 still have solutions because their right-hand-side quantities belong to the column spaces of these matrices (see concrete examples in Appendix B.2). In this paper, we will directly rely on the existence of bridge functions rather than the completeness conditions.

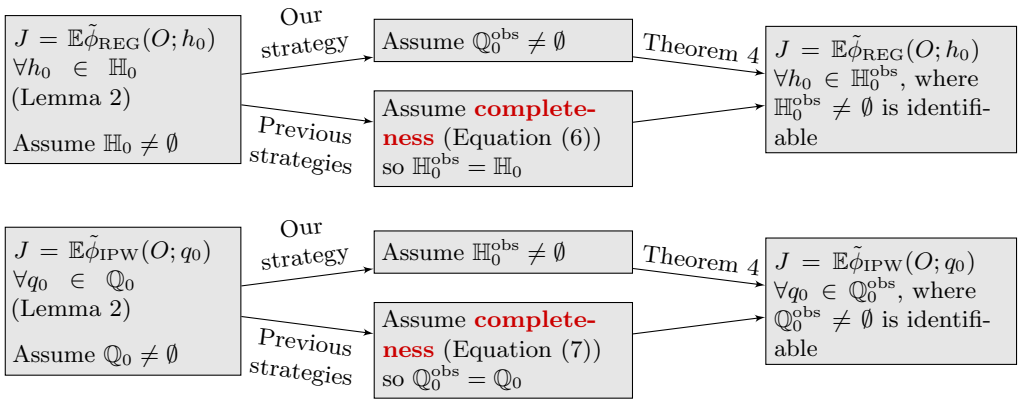


Figure 2: Different identification strategies. See relevant discussions below Theorem 4 and in Appendices A to C.

Examples 3 and 4 illustrate that bridge functions, besides existing, are nonunique if the negative control proxies carry more information than the unmeasured confounders, namely, when the negative controls have more values or dimensions than the unmeasured confounders. As unmeasured confounders are unobserved in practice, we may tend to use as many negative control variables as possible to safeguard the existence of bridge functions. But this may also cause bridge functions to be nonunique. Therefore, assuming uniqueness may be too strong. Our paper avoids the uniqueness assumption as much as possible and imposes uniqueness only when it is needed to derive stronger theoretical guarantees.

The lemma below shows that even when bridge functions are nonunique, *any* one of them can identify J. We first define an operator for the lemma below:

$$\mathcal{T} : L_2(W, A, X) \to L_2(W, X), \quad (\mathcal{T}h)(w, x) = \int h(w, a, x)\pi(a|x)\mathrm{d}\mu(a).$$

Lemma 2. *Suppose that Assumption 1 holds. For any $h_0 \in \mathbb{H}_0$ and $q_0 \in \mathbb{Q}_0$,*

$$J = \mathbb{E}\tilde{\phi}_{\text{IPW}}(O; q_0) = \mathbb{E}\tilde{\phi}_{\text{REG}}(O; h_0) = \mathbb{E}\tilde{\phi}_{\text{DR}}(O; h_0, q_0),$$
$$\textit{where} \quad \tilde{\phi}_{\text{IPW}}(O; q_0) = \pi(A \mid X)q_0(Z, A, X)Y, \quad \tilde{\phi}_{\text{REG}}(O; h_0) = (\mathcal{T}h_0)(X, W),$$
$$\tilde{\phi}_{\text{DR}}(O; h_0, q_0) = \pi(A \mid X)q_0(Z, A, X)\left(Y - h_0(W, A, X)\right) + (\mathcal{T}h_0)(W, X).$$

Note that the estimating equations in Lemma 2 simply replace the regression function $k_0(A, X, U)$ and the inverse propensity score weight $1/f(A \mid X, U)$ in Lemma 1 by the bridge functions $h_0(W, A, X)$ and $q_0(Z, A, X)$ respectively. Since the latter only depends on observed variables, as long as we can learn *any* pair of bridge functions, we can use estimating equations in Lemma 2 to estimate J.

3.3 Learning Bridge Functions from Observed Data

Definition 1 defines bridge functions in terms of conditional moment equations[1] given the unobserved confounders U, so we cannot directly use it to learn the bridge functions from the observed data. Nevertheless, the following lemma shows that the bridge functions also satisfy analogous conditional moment equations based only on observed data. This lemma extends the characterizations of bridge functions in Miao et al. (2018b); Cui et al. (2023).

Lemma 3. *Under Assumption 1, any $h_0 \in \mathbb{H}_0$ and $q_0 \in \mathbb{Q}_0$ satisfy*

$$\mathbb{E}[Y - h_0(W, A, X) \mid Z, A, X] = 0, \tag{8}$$

$$\mathbb{E}[\pi(A \mid X)\,(q_0(Z, A, X) - 1/f(A \mid W, X)) \mid W, A, X] = 0. \tag{9}$$

The conditional moment equations in Equations (8) and (9) give rise to the following alternative sets of functions that we can possibly learn from observed data, whose elements we call *observed-data* bridge functions:

$$\mathbb{H}_0^{\text{obs}} = \{h \in L_2(W, A, X) : \mathbb{E}[Y - h(W, A, X) \mid Z, A, X] = 0\},$$
$$\mathbb{Q}_0^{\text{obs}} = \{q : \pi q \in L_2(Z, A, X), \mathbb{E}[\pi(A \mid X)\,(q(Z, A, X) - 1/f(A \mid W, X)) \mid W, A, X] = 0\}.$$

According to Lemma 3, $\mathbb{H}_0 \subseteq \mathbb{H}_0^{\text{obs}}$ and $\mathbb{Q}_0 \subseteq \mathbb{Q}_0^{\text{obs}}$. So, when $\mathbb{H}_0 \neq \emptyset, \mathbb{Q}_0 \neq \emptyset$, we also have $\mathbb{H}_0^{\text{obs}} \neq \emptyset, \mathbb{Q}_0^{\text{obs}} \neq \emptyset$. Actually, it can be shown that $\mathbb{H}_0 \subseteq \mathbb{H}_0^{\text{obs}}$ and $\mathbb{Q}_0 \subseteq \mathbb{Q}_0^{\text{obs}}$ when the completeness conditions in Equations (6) and (7) hold respectively (see Lemma 24 in Appendix B). But without completeness conditions, the sets of observed-data bridge functions $\mathbb{H}_0^{\text{obs}}$ or $\mathbb{Q}_0^{\text{obs}}$ can be strictly bigger than the sets of true bridge functions $\mathbb{H}_0, \mathbb{Q}_0$. In the following lemma, we further show that observed-data bridge functions in $\mathbb{H}_0^{\text{obs}}$ or $\mathbb{Q}_0^{\text{obs}}$ can be directly used to identify J, even if they are not true bridge functions in $\mathbb{H}_0, \mathbb{Q}_0$.

Theorem 4. *Suppose Assumption 1 holds.*

1. *If $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$, then $J = \mathbb{E}\tilde{\phi}_{\text{REG}}(O; h_0)$ for any $h_0 \in \mathbb{H}_0^{\text{obs}}$.*
2. *If $\mathbb{Q}_0 \neq \emptyset$ and $\mathbb{H}_0^{\text{obs}} \neq \emptyset$, then $J = \mathbb{E}\tilde{\phi}_{\text{IPW}}(O; q_0)$ for any $q_0 \in \mathbb{Q}_0^{\text{obs}}$.*
3. *If the conditions in either statement 1 or statement 2 hold, then $J = \mathbb{E}\tilde{\phi}_{\text{DR}}(O; h_0, q_0)$ for any $h_0 \in \mathbb{H}_0^{\text{obs}}$ and $q_0 \in \mathbb{Q}_0^{\text{obs}}$.*

Theorem 4 suggests a straightforward way to estimate J: first estimate $h_0 \in \mathbb{H}_0^{\text{obs}}$ and $q_0 \in \mathbb{Q}_0^{\text{obs}}$ by solving Equations (8) and (9), and then estimate the GACE J by using any of the estimating equations in Lemma 2.

Note that by Lemma 3, existence of both true bridge functions, $\mathbb{H}_0 \neq \emptyset, \mathbb{Q}_0 \neq \emptyset$, is a sufficient condition to guarantee the assumptions of each of the three statements of Theorem 4. In Appendix A, we show that the existence of just one bridge function, *i.e.*, $\mathbb{H}_0 \neq \emptyset$ or $\mathbb{Q}_0 \neq \emptyset$, is insufficient for identification. Again by Lemma 3, this also means the

1. Equations (3) and (9) do not exactly fall into the usual conditional moment equation framework (*e.g.*, Ai and Chen, 2003), since they involve unknown density functions $f(A \mid U, X), f(A \mid W, X)$, respectively. We still call them conditional moment equations for simplicity, but estimating q_0 does require more care. See Section 4.2.1.

existence of just one observed data bridge function, *i.e.*, $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ or $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$, is similarly insufficient. In Appendix A, we prove that even the simultaneous existence of two observed data bridge functions, *i.e.*, $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$, is also insufficient for identification.

Interestingly, the identification formulae in Theorem 4 hold for any bridge functions $h_0 \in \mathbb{H}_0^{\text{obs}}$ and $q_0 \in \mathbb{Q}_0^{\text{obs}}$, even if they violate the conditional moment equations in Equations (2) and (3), *i.e.*, it holds for $h_0 \in \mathbb{H}_0^{\text{obs}} \backslash \mathbb{H}_0$ and $q_0 \in \mathbb{Q}^{\text{obs}} \backslash \mathbb{Q}_0$. In other words, even when the sets of bridge functions are unidentifiable (*i.e.*, $\mathbb{H}_0 \subsetneq \mathbb{H}_0^{\text{obs}}$ or $\mathbb{Q}_0 \subsetneq \mathbb{Q}_0^{\text{obs}}$), the GACE parameter J can still be identifiable, provided that conditions in Theorem 4 hold.

Our identification strategies in Theorem 4 are different from previous identification results based on observed-data bridge functions (Cui et al., 2023; Deaner, 2021; Miao et al., 2018a). This previous literature assumes the existence of one observed-data bridge function (*i.e.*, either $\mathcal{H}_0^{\text{obs}} \neq \emptyset$ or $\mathcal{Q}_0^{\text{obs}} \neq \emptyset$), and impose the *completeness conditions* in Equation (6) or Equation (7) to ensure $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$ or $\mathbb{Q}_0^{\text{obs}} = \mathbb{Q}_0$ (see Lemma 24). Take the identification via $\tilde{\phi}_{\text{REG}}$ as an example. Previous literature assume the completeness condition in Equation (6) and assume $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ (or equivalently $\mathbb{H}_0 \neq \emptyset$ under the completeness condition). Then any *observed-data* bridge function $h_0 \in \mathbb{H}_0^{\text{obs}}$ must also be a true bridge function in $\mathbb{H}_0$, and thus can identify GACE according to Lemma 2. In contrast, our identification result in Theorem 4 statement 1 assumes $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$, but not any completeness condition. In other words, our identification strategy replaces a completeness condition by the condition $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$. Our result is based on a completely *new* proof that allows us to directly identify GACE via any $h_0 \in \mathbb{H}_0^{\text{obs}}$ without requiring h_0 to also belong to $\mathbb{H}_0$. In this way, we can achieve identification without assuming any completeness conditions, and allow for $\mathbb{H}_0 \subsetneq \mathbb{H}_0^{\text{obs}}$. See Appendix B.1 for a detailed review of the identification strategies in Cui et al. (2023) and Figure 2 for an illustration for the difference in identification strategies.

One may hope to compare our new identification assumptions with those in the previous literature. Intuitively, our conditions of the existence of bridge functions require the proxy variables Z and W to capture the *relevant* variations of the unobserved confounders U, such that the effects of U on some bridge function transformations of these proxy variables can recover the effects of U on the action A and outcome Y respectively. In contrast, the completeness conditions in the existing literature require the proxy variables Z and W to capture *all* variations of U, so that no matter how U affect the action A and outcome Y, the proxy variables can always account for all unmeasured confounding. Therefore, our identification strategy does not rely on completeness (even implicitly), but instead replaces it with a different, more targeted requirement.

The differences between the two different types of conditions are particularly clear in the discrete data setting (Example 3) and linear model setting (Example 4). As we discussed in these two examples and Example 5, our assumptions about the existence of bridge functions require the existence of solutions to some linear equation systems, while the completeness conditions usually require the coefficient matrices of these linear equations to have full rank. The full rank conditions are sufficient but not necessary for the existence of solutions. There exist many instances where the completeness conditions fail but our proposed conditions still hold. Therefore, in the discrete data setting and linear model setting, our proposed conditions can be *strictly* weaker than the completeness conditions imposed in the existing literature. We refer to Appendix B.2 for more concrete discussions.

The comparisons between the two different types of conditions are more difficult in general nonparametric models (Example 5). In general models, the existence of bridge functions amounts to the existence of solutions to some linear equations defined by some potentially infinite-dimensional and ill-posed linear operators. In contrast, the completeness conditions require the adjoint of these linear operators to be injective mappings. For infinite-dimensional operators and equations, these two different types of conditions are closely related but are not directly comparable[2] without further regularity conditions. One approach to relate them is through Picard's theorem in functional analysis, as we discussed in Example 5. Under the regularity conditions therein, one can again prove that the completeness conditions are sufficient for our conditions. See Appendix C.2 for more details. Another setting with clear comparisons is to consider our combined identification conditions $\mathbb{H}_0 \neq \emptyset, \mathbb{Q}_0 \neq \emptyset$ that enable all three identification formulas in Theorem 4. To achieve the same identification goal, the existing literature would require $\mathbb{H}_0^{\text{obs}} \neq \emptyset, \mathbb{Q}_0^{\text{obs}} \neq \emptyset$ and also the two completeness conditions in Equations (6) and (7), which in turn imply our combined identification condition $\mathbb{H}_0 \neq \emptyset, \mathbb{Q}_0 \neq \emptyset$.

In summary, our proposed identification conditions and theory complement the identification results in the previous literature, thereby broadening the scope of causal identification with negative controls. But we note that these different identification strategies lead to the same identification formula (see Theorem 4), so their differences do not impact the statistical estimation and inference. In the rest of this paper, we study the estimation and inference of the GACE parameter in Equation (1) based on the unified identification formula.

4. Minimax Estimation of GACE

We first define estimators of GACE and then introduce minimax estimators for the bridge functions.

4.1 Estimation of GACE

Once we obtain bridge function estimators $\hat{h}, \hat{q}$ as explained in the next section, we can plug them into the estimating equations in Theorem 4 to construct the following estimators for the GACE parameter J:

$$\hat{J}_{\text{IPW}} = \mathbb{E}_n\Big[\tilde{\phi}_{\text{IPW}}(O; \hat{q})\Big] = \mathbb{E}_n[\pi(A \mid X)\hat{q}(Z, A, X)Y],$$
$$\hat{J}_{\text{REG}} = \mathbb{E}_n\Big[\tilde{\phi}_{\text{REG}}(O; \hat{h})\Big] = \mathbb{E}_n\Big[(\mathcal{T}\hat{h})(X, W)\Big],$$
$$\hat{J}_{\text{DR}} = \mathbb{E}_n\Big[\tilde{\phi}_{\text{DR}}(O; \hat{h}, \hat{q})\Big] = \mathbb{E}_n\Big[\pi(A \mid X)\hat{q}(Z, A, X)\Big(Y - \hat{h}(W, A, X)\Big) + (\mathcal{T}\hat{h})(X, W)\Big].$$

To use these estimators, we need that the corresponding equations identify GACE to begin with. Thus, throughout the rest of the paper, we assume Assumption 1 and $\mathbb{H}_0 \neq \emptyset, \mathbb{Q}_0 \neq \emptyset$ (see discussions below Theorem 4).

In this section, we discuss how to construct $\hat{h}, \hat{q}$, presenting two types of minimax estimators. In subsequent sections, we discuss resulting guarantees for GACE estimation under various estimators and assumptions, an overview of which we present in Section 4.4.

2. It is possible that one type of conditions hold while the other type of conditions fail, due to challenging properties of infinite-dimensional function spaces that differ from finite-dimensional spaces.

4.2 Minimax Estimators of Bridge Functions

Estimating bridge functions based on Equations (8) and (9) requires solving conditional moment equations, which is generally a difficult estimation problem. Estimation methods in the previous literature on negative controls mainly focus on parametric methods, sieve methods or Reproducing Kernel Hilbert Space methods (see Section 8 for a review). In this paper, we propose to use minimax approaches to estimate the bridge functions, which accommodate not only the hypothesis classes used in the previous literature, but also more flexible ones such as neural networks.

In this section, we introduce two minimax reformulations of the conditional moment equations in Equations (8) and (9). Each reformulation motivates an estimation strategy for bridge functions. Similar reformulations have also been used by previous literature for instrument variable (IV) estimation (see Section 8). To introduce these reformulations, we consider a generic conditional estimating equation problem,

$$\mathbb{E}\left[\rho\left(g_0(O_1), O_1\right) \mid O_2\right] = 0, \tag{10}$$

where g_0 is a function that we wish to solve for and O_1, O_2 are two sets of random variables.

4.2.1 Strategy I: Minimax Estimators without Stabilizers

We first note that Equation (10) has the following reformulation:

$$\begin{aligned}\mathbb{E}\left[\rho\left(g_0(O_1), O_1\right) \mid O_2\right] = 0 &\iff \mathbb{E}\left[g'(O_2)\rho\left(g_0(O_1), O_1\right)\right] = 0, \quad \forall g' \in L_2(O_2), \\ &\iff \sup_{g' \in L_2(O_2)} \left(\mathbb{E}\left[g'(O_2)\rho\left(g_0(O_1), O_1\right)\right]\right)^2 = 0. \end{aligned} \tag{11}$$

This motivates the following estimators for bridge functions

$$\hat{h} \in \arg\min_{h \in \mathbb{H}} \max_{q \in \mathbb{Q}'} \quad \left(\mathbb{E}_n[q(Z, A, X)\left(h(W, A, X) - Y\right)]\right)^2, \tag{12}$$

$$\hat{q} \in \arg\min_{q \in \mathbb{Q}} \max_{h \in \mathbb{H}'} \quad \left(\mathbb{E}_n[\pi(A|X)q(Z, A, X)h(W, A, X) - (\mathcal{T}h)(W, X)]\right)^2, \tag{13}$$

Our estimators $\hat{h}, \hat{q}$ can be viewed as solutions to minimax games, where an adversarial player picks elements from function classes $\mathbb{Q}', \mathbb{H}'$ to form the most difficult marginal moments while our estimators minimize the violations of such moments. Equations (12) and (13) involve two types of function classes: we call $\mathbb{Q}, \mathbb{H}$ the *bridge classes* and $\mathbb{Q}', \mathbb{H}'$ the *critic classes*, and we call elements of the latter *critic functions*. Note that throughout, $\mathbb{Q}, \mathbb{H}, \mathbb{Q}', \mathbb{H}'$ can change with n. We review some examples in Section 4.3 below.

Note that although Equation (9) involves the generalized propensity score $f(A \mid W, X)$, it does not appear in Equation (13) at all. In this sense, our estimation method for $\hat{q}$ is different from a naïve application of Equation (11) to Equation (9), wherein we would first get a preliminary generalized propensity score estimator $\hat{f}\left(A \mid W, X\right)$ and then solve

$$\arg\min_{q \in \mathbb{Q}} \max_{h \in \mathbb{H}'} (\mathbb{E}_n[h(W, A, X)\pi(A|X)\{q - 1/\hat{f}(A|W, X)\}])^2.$$

Instead, our estimator in Equation (13) exploits the fact that

$$\mathbb{E}\left[h(W, A, X)\pi\left(A \mid X\right)/f(A \mid W, X)\right] = \mathbb{E}\left[\left(\mathcal{T}h\right)(X, W)\right].$$

Thus, it obviates the need to estimate the generalized propensity score before estimating the bridge functions. This fact also characterizes the difference in the estimation of h_0 and q_0. The estimation of h_0 is analogous to the nonparametric IV regression problem (Newey and Powell, 2003; Darolles et al., 2010), and the estimator Equation (12) is analogous to minimax approaches therein. In contrast, the estimation of q_0 requires additional considerations.

Example 1, Cont'd (Average treatment effect). Consider binary action $A \in \{0,1\}$ in Example 1. In this case, the conditional moment equation for the action bridge function q_0 is equivalent to

$$\mathbb{E}\left[\mathbb{I}\left[A=a\right] q_0(Z,a,X) - 1 \mid W, X\right] = 0, \quad a \in \{0,1\}. \tag{14}$$

Apparently, this equation does not explicitly depend on the propensity score either. Note that when $Z = W = \emptyset$, the action bridge function given by this conditional moment equation is exactly the inverse propensity score weight $1/\mathbb{P}\left[A = a \mid X\right]$.

4.2.2 Strategy II: Minimax Estimators with Stabilizers

Again considering the generic estimating equation in Equation (10), note that for any constant $\lambda > 0$,

$$\begin{aligned}
&\mathbb{E}\left[\rho\left(g_0(O_1), O_1\right) \mid O_2\right] = 0 \iff \\
&0 = \frac{1}{4\lambda}\mathbb{E}(\mathbb{E}\left[\rho\left(g_0(O_1), O_1\right) \mid O_2\right])^2 = \sup_{g' \in L_2(O_2)} \mathbb{E}\left[g'(O_2)\rho\left(g_0(O_1), O_1\right)\right] - \lambda\|g'\|_2^2.
\end{aligned} \tag{15}$$

This motivates the following estimators for bridge functions

$$\hat{h} \in \arg\min_{h\in\mathbb{H}} \max_{q\in\mathbb{Q}'} \mathbb{E}_n[q\left(Z,A,X\right)\left(h(W,A,X) - Y\right)] - \lambda\mathbb{E}_n[q^2\left(Z,A,X\right)], \tag{16}$$

$$\begin{aligned}
\hat{q} \in \arg\min_{q\in\mathbb{Q}} \max_{h\in\mathbb{H}'} \mathbb{E}_n[\pi(A \mid X) q(Z,A,X) h(W,A,X) - (\mathcal{T}h)(W,X)] \\
- \lambda\mathbb{E}_n[h^2\left(W,A,X\right)].
\end{aligned} \tag{17}$$

We call the terms $\lambda\mathbb{E}_n[q^2], \lambda\mathbb{E}_n[h^2]$ *stabilizers*. Stabilizers are different from regularizers. Regularizers typically introduce estimation bias (*e.g.*, Knight and Fu, 2000; Carrasco et al., 2007), so we generally let them vanish when the sample size grows. Stabilizers, on the other hand, do not introduce bias and are merely a way to reformulate the conditional moment equations. In particular, the objective function of the stabilizer-based reformulation involves only marginal expectations that can be easily estimated by sample averages. In Section 7, we show that they generally should not vanish. One exception is when critic classes $\mathbb{Q}'$ and $\mathbb{H}'$ are symmetric: in this case, the objective functions in Equations (16) and (17) with $\lambda = 0$ are equivalent to their counterparts in Equations (12) and (13).

4.3 Examples of Bridge Function Estimators

In this part, we give examples of bridge function estimators based on three different critic function classes: linear class, RKHS, and neural networks. In particular, for the linear class and RKHS, the inner maximization problems in Equations (12), (13), (16) and (17) have closed-form solutions, so that the minimax problems can be solved by standard optimization techniques such as stochastic gradient descent.

In this subsection, we focus on minimax estimators without stabilizers given in Equations (12) and (13). We can similarly compute minimax estimators with stabilizers given in Equations (16) and (17), and defer the details to Appendix E.2.

4.3.1 Linear Classes

Given $\phi : \mathcal{Z} \times \mathcal{A} \times \mathcal{X} \to \mathbb{R}^{d_1}, \psi : \mathcal{W} \times \mathcal{A} \times \mathcal{X} \to \mathbb{R}^{d_2}$, set

$$\mathbb{Q}' = \{(z,a,x) \mapsto \alpha_1^\top \phi(z,a,x) : \alpha_1 \in \mathbb{R}^{d_1}, \|\alpha_1\| \le c_1\}, \tag{18}$$

$$\mathbb{H}' = \{(w,a,x) \mapsto \alpha_2^\top \psi(w,a,x) : \alpha_2 \in \mathbb{R}^{d_2}, \|\alpha_2\| \le c_2\}. \tag{19}$$

Typical classical examples of basis functions include splines, polynomials, and wavelets (Chen, 2007). Another example is random feature expansions for positive definite kernels (Bach, 2017). When d_1, d_2 grow with n, these function classes are also called *linear sieves*. See also discussions in Appendix G.1.

It is easy to show that with linear critic classes, the inner maximum objectives in Equations (12) and (13) have closed-form expressions, and the resulting bridge function estimators are:

$$\begin{aligned} \hat{h} &\in \arg\min_{h \in \mathbb{H}} \left(\mathbb{E}_n[(Y-h)\phi]\right)^\top \left(\mathbb{E}_n[(Y-h)\phi]\right), \\ \hat{q} &\in \arg\min_{q \in \mathbb{Q}} \left(\mathbb{E}_n[q\pi\psi - \mathcal{T}\psi]\right)^\top \left(\mathbb{E}_n[q\pi\psi - \mathcal{T}\psi]\right). \end{aligned} \tag{20}$$

Given basis functions $\tilde{\phi} : \mathcal{Z} \times \mathcal{A} \times \mathcal{X} \to \mathbb{R}^{\tilde{d}_1}$ and $\tilde{\psi} : \mathcal{W} \times \mathcal{A} \times \mathcal{X} \to \mathbb{R}^{\tilde{d}_2}$, if further $\mathbb{Q}$ and $\mathbb{H}$ are also linear classes (without norm constraints for simplicity), namely,

$$\begin{aligned} \pi\mathbb{Q} &= \{(z,a,x) \mapsto \alpha_1^\top \tilde{\phi}(z,a,x) : \alpha_1 \in \mathbb{R}^{\tilde{d}_1}\}. \\ \mathbb{H} &= \{(w,a,x) \mapsto \alpha_2^\top \tilde{\psi}(w,a,x) : \alpha_2 \in \mathbb{R}^{\tilde{d}_2}\}, \end{aligned} \tag{21}$$

Then, the corresponding IPW and REG estimators have s:

$$\begin{aligned} \hat{J}_{\mathrm{REG}} &= \mathbb{E}_n[\mathcal{T}\tilde{\psi}]^\top \{\mathbb{E}_n[\tilde{\psi}\phi^\top]\mathbb{E}_n[\phi\tilde{\psi}^\top]\}^+ \mathbb{E}_n[\tilde{\psi}\phi^\top]\mathbb{E}_n[Y\phi], \\ \hat{J}_{\mathrm{IPW}} &= \mathbb{E}_n[\mathcal{T}\psi]^\top \mathbb{E}_n[\psi\tilde{\phi}^\top]\{\mathbb{E}_n[\tilde{\phi}\psi^\top]\mathbb{E}_n[\psi\tilde{\phi}^\top]\}^+ \mathbb{E}_n[Y\tilde{\phi}]. \end{aligned}$$

When $\tilde{\phi} = \phi, \tilde{\psi} = \psi$, we have $\hat{J}_{\mathrm{REG}} = \hat{J}_{\mathrm{IPW}} = \hat{J}_{\mathrm{DR}}$. This extends a similar equivalence result for the unconfounded setting (Singh et al., 2020; Uehara et al., 2020; Kallus, 2020).

Lemma 5. *Suppose $\tilde{\phi} = \phi$ and $\tilde{\psi} = \psi$. Then, we have*

$$\hat{J}_{\mathrm{IPW}} = \hat{J}_{\mathrm{REG}} = \hat{J}_{\mathrm{DR}} = \mathbb{E}_n[\mathcal{T}\psi]^\top \mathbb{E}_n[\phi\psi^\top]^+ \mathbb{E}_n[Y\phi]. \tag{22}$$

Example 1, Cont'd (Average treatment effect). Consider $\mathcal{A} = \{0,1\}$ and the parameter $\mathbb{E}[Y(a_0)]$ for $a_0 \in \{0,1\}$ as in Example 1. Given $\phi_b : \mathcal{Z} \times \mathcal{X} \to \mathbb{R}^{d_1}$, $\psi_b : \mathcal{W} \times \mathcal{X} \to \mathbb{R}^{d_2}$, we set the basis functions $\phi = \tilde{\phi} = (\mathbb{I}[a=0]\,\phi_b^\top(z,x), \mathbb{I}[a=1]\,\phi_b^\top(z,x))^\top$ and $\psi = \tilde{\psi} = (\mathbb{I}[a=0]\,\psi_b^\top(w,x), \mathbb{I}[a=1]\,\psi_b^\top(w,x))^\top$. Then,

$$\begin{aligned} &\hat{J}_{\mathrm{IPW}} = \hat{J}_{\mathrm{REG}} = \hat{J}_{\mathrm{DR}} = \hat{\xi}(a_0), \\ &\hat{\xi}(a_0) \coloneqq \mathbb{E}_n[\mathbb{I}(A=a_0)\psi_b(W,X)]^\top \mathbb{E}_n[\mathbb{I}(A=a_0)\phi_b(Z,X)\psi_b^\top(W,X)]^+ \mathbb{E}_n[\mathbb{I}(A=a_0)Y\phi_b(Z,X)]. \end{aligned}$$

We can then estimate the average treatment effect by $\hat{\xi}(1) - \hat{\xi}(0)$.

Example 3, Cont'd (Discrete setting). In Example 3, we let W, Z, U be discrete. Now further assume that X, A, Y are also discrete, and let ϕ, ψ be atomic basis functions corresponding to all possible values of (Z, A, X) and (W, A, X) respectively. Then Equation (22) reduces to

$$\hat{J}_{\mathrm{IPW}} = \hat{J}_{\mathrm{REG}} = \hat{J}_{\mathrm{DR}} = \textstyle\sum_{y\in\mathcal{Y}, a\in\mathcal{A}, x\in\mathcal{X}} y\pi(a \mid x)\hat{P}(y, \mathbf{Z}, a, x)\hat{P}(\mathbf{W}, \mathbf{Z}, a, x)^{+}\hat{P}(\mathbf{W}, x),$$

where $\hat{P}(y, \mathbf{Z}, a, x), \hat{P}(\mathbf{W}, \mathbf{Z}, a, x), \hat{P}(\mathbf{W}, x)$ are vectors and matrices consisting of sample frequency estimates of the corresponding probabilities. For example, $\hat{P}(y, \mathbf{Z}, a, x) \in \mathbb{R}^{|Z|}$ is the column vector whose jth element is $\frac{1}{n}\sum_{i=1}^{n} \mathbb{I}(Y_i = y, Z_i = z_j, X_i = x, A_i = a)$.

4.3.2 RKHS

Consider two positive-semidefinite kernels $k_z : (\mathcal{Z}, \mathcal{A}, \mathcal{X}) \times (\mathcal{Z}, \mathcal{A}, \mathcal{X}) \to \mathbb{R}$ and $k_w : (\mathcal{W}, \mathcal{A}, \mathcal{X}) \times (\mathcal{W}, \mathcal{A}, \mathcal{X}) \to \mathbb{R}$, and denote the induced RKHSs by $\mathcal{L}_z$ and $\mathcal{L}_w$ with RKHS norms $\|\cdot\|_{\mathcal{L}_z}$ and $\|\cdot\|_{\mathcal{L}_w}$, respectively. We consider the following critic classes:

$$\mathbb{Q}' = \{q : q \in \mathcal{L}_z, \|q\|_{\mathcal{L}_z} \le c_1\}, \quad \mathbb{H}' = \{h : h \in \mathcal{L}_w, \|h\|_{\mathcal{L}_w} \le c_2\}. \tag{23}$$

Lemma 6. *For $\mathbb{Q}', \mathbb{H}'$ given in Equation (23), the estimators in Equations (12) and (13) are given by*

$$\hat{h} \in \arg\min_{h\in\mathbb{H}} \ (\psi_n(h))^{\top} K_{z,n}\psi_n(h), \tag{24}$$

$$\hat{q} \in \arg\min_{q\in\mathbb{Q}} \ (\phi_n(q))^{\top} K_{w1,n}\phi_n(q) - 2(\phi_n(q))^{\top} K_{w2,n}\mathbf{1}_n, \tag{25}$$

where $K_{z,n}$, $K_{w1,n}$, $K_{w2,n}$ are $n \times n$ Gram matrices whose (i,j)th entry is

$$k_z((Z_i, A_i, X_i), (Z_j, A_j, X_j)), \qquad k_w((W_i, A_i, X_i), (W_j, A_j, X_j)),$$
$$\mathbb{E}_{\pi(A_j|X_j)}[k_w((W_i, A_i, X_i), (W_j, A_j, X_j))],$$

respectively, and $\psi_n(h) \in \mathbb{R}^n$, $\phi_n(q) \in \mathbb{R}^n$ are column vectors whose ith elements are $(Y_i - h(W_i, A_i, X_i))$, $q(Z_i, A_i, X_i)\pi(A_i|X_i)$, respectively, and $\mathbf{1}_n \in \mathbb{R}^n$ is an all-ones vector.

In Equations (24) and (25), computing the final bridge estimators only involves minimization problems whose objectives are convex in h or q and only depend on h or q via their evaluation at data points. If $\mathbb{H}, \mathbb{Q}$ are linear hypothesis classes like those in Equation (21), then $\hat{h}$ and $\hat{q}$ also have closed-form solutions. If $\mathbb{H}$ is an RKHS hypothesis class with kernel $\tilde{k}$ with either norm constraints or norm regularizers, then the optimal solution to Equation (24) will have the form $\hat{h}(w, a, x) = \sum_{i=1}^{n} \alpha_i \tilde{k}((W_i, A_i, X_i), (w, a, x))$, leading to a convex quadratic program in α with a closed-form solution. The same holds for Equation (25) if $\mathbb{Q}$ is an RKHS hypothesis class. If the hypothesis classes $\mathbb{H}, \mathbb{Q}$ are some more complex classes such as neural networks, then we can use stochastic gradient descent methods to solve for $\hat{h}, \hat{q}$, which have been shown to be highly successful in many nonconvex applications (Jain and Kar, 2017).

Example 1, Cont'd (Average treatment effect). Again, consider binary action $\mathcal{A} = \{0, 1\}$ as in Example 1 and the estimation of $\mathbb{E}[Y(a_0)]$ for $a_0 \in \{0, 1\}$. Given kernels $\bar{k}_z : (\mathcal{Z} \times \mathcal{X}) \times (\mathcal{Z} \times \mathcal{X}) \to \mathbb{R}$ and $\bar{k}_w : (\mathcal{W}, \mathcal{X}) \times (\mathcal{W}, \mathcal{X}) \to \mathbb{R}$, we set

$$k_z((Z_i, A_i, X_i), (Z_j, A_j, X_j)) = \mathbb{I}[A_i = A_j]\, \bar{k}_z((Z_i, X_i), (Z_j, X_j)),$$

$$k_w((W_i, A_i, X_i), (W_j, A_j, X_j)) = \mathbb{I}[A_i = A_j]\, \bar{k}_w((W_i, X_i), (W_j, X_j)).$$

Then we can solve Equations (24) and (25) to get estimators $\hat{h}(\cdot, a_0, \cdot), \hat{q}(\cdot, a_0, \cdot)$, using only observations corresponding to the action a_0. For the detailed calculation, see Appendix E.1. Then we can use $\hat{J}_{\mathrm{REG}} = \mathbb{E}_n[\hat{h}(W, a_0, X)]$, $\hat{J}_{\mathrm{IPW}} = \mathbb{E}_n[\mathbb{I}[A = a_0]\, \hat{q}(Z, A, X)\, Y]$ or $\hat{J}_{\mathrm{DR}} = \mathbb{E}_n[\hat{h}(W, a_0, X) + \mathbb{I}[A = a_0]\, \hat{q}(Z, A, X)\, (Y - \hat{h}(W, A, X))]$.

4.3.3 Neural Networks

An L-layer neural network with input x can be generically written in the following form:

$$\sigma_L(W_L(\sigma_{L-1}(\ldots \sigma_2(W_1 x + b_1)\ldots) + b_{L-1}) + b_L)$$

where $W_l \in \mathbb{R}^{d_{l,1} \times d_{l,2}}, b_l \in \mathbb{R}^{d_{l,2}}, \sigma_l$ for $l = 1, 2, \ldots, L$ are called as weights, biases and activation functions, respectively. One standard choice for the activation functions is the ReLU function defined by $\sigma(x) = [\max(x_1, 0,), \ldots, \max(x_L, 0)]^\top$. When $\mathbb{Q}, \mathbb{H}, \mathbb{Q}', \mathbb{H}'$ are neural network classes, we need to solve non-convex minimax optimization problems to compute the minimax estimators. These can be solved by several types of simultaneous stochastic gradient descent methods. One method is the (simultaneous version of) Adam (Kingma and Ba, 2015), which is a variant of gradient descent with momentum and per-parameter adaptive learning rates. An improved approach is the Optimistic Adam (Daskalakis et al., 2018), which is an adaptation of Adam with additional negative momentum.

4.4 An Overview of the Estimation Theory for GACE

In Sections 5 to 7, we will show that GACE estimators based on different estimating equations and different minimax bridge functions estimators, either without stabilizers (Section 4.2.1) or with stabilizers (Section 4.2.2), have different theoretical properties. Each type of estimator has its merits depending on the assumptions.

In the rest of the paper, we will *always* assume that functions in classes $\mathbb{Q}, \mathbb{H}, \mathbb{Q}', \mathbb{H}'$ are square integrable. We further define linear operators $P_z : L_2(W, A, X) \to L_2(Z, A, X)$ and $P_w : L_2(Z, A, X) \to L_2(W, A, X)$ as follows:

$$P_z(h) = \mathbb{E}[h(W, A, X) \mid Z, A, X], \quad P_w(q) = \mathbb{E}[q(Z, A, X) \mid W, A, X]. \tag{26}$$

We will show that two additional types of assumptions will play an important role in theoretical guarantees for different estimators.

1. *Realizability*, which characterizes whether the classes $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$ contain some observed bridge functions, *i.e.*, $\mathbb{H} \cap \mathbb{H}_0^{\mathrm{obs}} \neq \emptyset$, $\mathbb{Q} \cap \mathbb{Q}_0^{\mathrm{obs}} \neq \emptyset$, or $\mathbb{H}' \cap \mathbb{H}_0^{\mathrm{obs}} \neq \emptyset$, $\mathbb{Q}' \cap \pi\mathbb{Q}_0^{\mathrm{obs}} \neq \emptyset$.

2. *Closedness*, which characterizes whether the critic classes $\mathbb{H}', \mathbb{Q}'$ are rich enough and P_z, P_w are smooth enough, requiring $P_z\left(\mathbb{H} - \mathbb{H}_0^{\mathrm{obs}}\right) \subseteq \mathbb{Q}'$ and/or $\pi P_w\left(\mathbb{Q} - \mathbb{Q}_0^{\mathrm{obs}}\right) \subseteq \mathbb{H}'$.

We remark that both assumptions are model specification assumptions. While it is difficult to validate these assumptions from data, specifying richer function classes generally make them more plausible. In this paper, we aim to develop estimation methods and theory that accommodate general function classes, thereby addressing the limitations of existing

approaches based on strong parametric restrictions. Moreover, our theory is general and can also characterize the potential errors due to partial violation of the realizability and closedness conditions.

The theoretical results for our GACE estimators are organized as follows:

- In Section 5, we derive finite-sample error bounds for GACE estimators based on minimax bridge function estimators *without stabilizers* (Section 4.2.1), under *double* realizability assumptions for *both* the bridge hypothesis class and the critic class[3]. Note that the bounds are derived for the GACE estimators directly, regardless of the convergence of the bridge function estimators.

- In Section 6, we assume a realizability assumption on the bridge hypothesis class and a closedness assumption on the critic class, under which we derive error bounds for the minimax bridge function estimators *without stabilizers* in terms of a projected mean-squared-error (MSE) metric. Then we translate these into the error bounds for the downstream plug-in GACE estimators.

- In Section 7, we derive *faster-rate* projected MSE error bounds for minimax bridge function estimators *with stabilizers* under the same assumption as Section 6. Then we also translate these error bounds into the error bounds for the downstream plug-in GACE estimators. Finally, we show that when we additionally assume that bridge functions are unique and that the conditional moment equations in Equations (8) and (9) are not too ill-posed, we have that the resulting doubly robust estimator $\hat{J}_{\mathrm{DR}}$ is asymptotically normal with asymptotic variance equal to the semiparametric efficiency bound.

In Table 1, we summarize our results for different estimators under the different assumptions. We remark that although the results in Table 1 are based on covering numbers, our theory in Sections 5 to 7 provides meta theorems that can accommodate a wide range of function classes (e.g., VC subgraph classes, nonparametric smooth classes, neural networks). These meta theorems are instantiated to different function classes to obtain concrete error bounds.

5. Finite-Sample Analysis of Estimators without Stabilizers under Realizability

In this section, we analyze different GACE estimators based on minimax estimation of bridge functions *without* stabilizers, *i.e.*, minimax estimators $\hat{h}, \hat{q}$ given in Equations (12) and (13). Throughout this section, we assume *only* realizability but *not* closedness.

5.1 Error Bounds of GACE Estimators

First, we bound the errors of the IPW and REG estimators without stabilizers. Then, we consider analogous results for DR.

3. In special cases such as linear models, we only require realizability assumptions for the bridge hypothesis class *or* the critic class. See discussions in Appendix F.

	Main Assumptions	Est	Convergence Rate	Sta	Uni	Notes
I	$\mathbb{Q}_0^{\text{obs}} \cap \mathbb{Q} \neq \emptyset, \mathbb{H}_0^{\text{obs}} \cap \mathbb{H}' \neq \emptyset$	IPW	$n^{-\min(\frac{1}{2},\beta)}$	N	N	$\hat{h}, \hat{q}$ need not converge to any point in $\mathbb{H}_0^{\text{obs}}, \mathbb{Q}_0^{\text{obs}}$
	$\mathbb{H}_0^{\text{obs}} \cap \mathbb{H} \neq \emptyset, \pi\mathbb{Q}_0^{\text{obs}} \cap \mathbb{Q}' \neq \emptyset$	REG				
	$\mathbb{Q}_0^{\text{obs}} \cap \mathbb{Q} \neq \emptyset, \mathbb{H}_0^{\text{obs}} \cap \{h : h - \mathbb{H} \subseteq \mathbb{H}'\} \neq \emptyset$	DR				
	$\mathbb{H}_0^{\text{obs}} \cap \mathbb{H} \neq \emptyset, \mathbb{Q}_0^{\text{obs}} \cap \{q : \pi(q - \mathbb{Q}) \subseteq \mathbb{Q}'\} \neq \emptyset$	DR				
II	$\mathbb{Q}_0^{\text{obs}} \cap \mathbb{Q} \neq \emptyset, \pi P_w(\mathbb{Q} - \mathbb{Q}_0^{\text{obs}}) \subseteq \mathbb{H}'$	IPW / DR	$n^{-\frac{\beta}{2\beta+1}}$ / $n^{-\min(\frac{1}{4},\frac{\beta}{2})}$	Y / N	N	$\|\pi P_w(\hat{q} - q_0)\|_2 \to 0$ for any $q_0 \in \mathbb{Q}_0^{\text{obs}}$
III	$\mathbb{H}_0^{\text{obs}} \cap \mathbb{H} \neq \emptyset$, $P_z(\mathbb{H} - \mathbb{H}_0^{\text{obs}}) \subseteq \mathbb{Q}'$	REG / DR	$n^{-\frac{\beta}{2\beta+1}}$ / $n^{-\min(\frac{1}{4},\frac{\beta}{2})}$	Y / N	N	$\|P_z(\hat{h} - h_0)\|_2 \to 0$ for any $h_0 \in \mathbb{H}_0^{\text{obs}}$
IV	Assumptions in rows II and III, $\|P_u(\hat{q}\pi - q_0\pi)\|_2 \leq \tau_1 \|P_w(\hat{q}\pi - q_0\pi)\|_2$, $\|P_u(\hat{h} - h_0)\|_2 \leq \tau_1 \|P_z(\hat{h} - h_0)\|_2$	DR	$\max(n^{-\frac{1}{2}}, \tau_1^2 n^{-\frac{2\beta}{2\beta+1}})$	Y	N	Faster than row I if $\tau_1^2 = o(n^{\frac{\beta(1-2\beta)}{2\beta+1}})$
	Assumptions in rows II and III, $\|\hat{q}\pi - q_0\pi\|_2 \leq \tau_2 \|P_w(\hat{q}\pi - q_0\pi)\|_2$, $\|\hat{h} - h_0\|_2 \leq \tau_2 \|P_z(\hat{h} - h_0)\|_2$	DR	$\max(n^{-\frac{1}{2}}, \tau_2 n^{-\frac{2\beta}{2\beta+1}})$	Y	Y	Achieves efficiency if $\tau_2 = o(n^{\frac{2\beta-1}{4\beta+2}})$

Table 1: Convergence rates of different GACE estimators under different assumptions. The identification assumptions in Theorem 4 are always assumed. The sets $\mathbb{H}_0^{\text{obs}}, \mathbb{Q}_0^{\text{obs}}$ denote the sets of observed bridge functions (see Lemma 3). The function classes $\mathbb{H}, \mathbb{Q}'$ are used to construct the estimator $\hat{h}$ and $\mathbb{Q}, \mathbb{H}'$ to construct the estimator $\hat{q}$. We here summarize our conclusions (ignoring polylogs) when all relevant function classes are nonparametric with log covering number at radius ϵ scaling as $\epsilon^{-\beta}$. "Est" refers to which estimating equation is used for the final estimator (see Section 4.4). "Sta" refers to whether the bridge function estimators use stabilizers (see Section 4). "Uni" refers to whether we assume bridge functions are unique. The projection operators P_z, P_w, P_u are defined in Equations (26) and (33).

Theorem 7 (Analysis of IPW and REG estimators). *For any* $h_0 \in \mathbb{H}_0^{\text{obs}}, q_0 \in \mathbb{Q}_0^{\text{obs}}$,

$$|\hat{J}_{\text{IPW}} - J| \leq \sup_{q\in\mathbb{Q}} |(\mathbb{E}_n - \mathbb{E})[q\pi Y]| + 2 \sup_{q\in\mathbb{Q}, h\in\mathbb{H}'} |(\mathbb{E}_n - \mathbb{E})[-q\pi h + \mathcal{T}h]| + \\ + \inf_{h\in\mathbb{H}',} \sup_{q\in\mathbb{Q}} |\mathbb{E}[(q - q_0)\pi(h_0 - h)]| + \inf_{q\in\mathbb{Q}} \sup_{h\in\mathbb{H}'} |\mathbb{E}[(q_0 - q)\pi h]|. \tag{27}$$

$$|\hat{J}_{\text{REG}} - J| \leq \sup_{h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[\mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q}', h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q(Y - h)]| \\ + \inf_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} |\mathbb{E}[\{h_0 - h\}q]| + \inf_{q\in\mathbb{Q}'} \sup_{h\in\mathbb{H}} |\mathbb{E}[(q_0\pi - q)(h - h_0)|. \tag{28}$$

The first two empirical process terms on the right-hand sides of each of Equations (27) and (28) account for the "variance" between the empirical and population estimating equations. In Section 5.2, we will show that these terms converge to 0 for some choices of $\mathbb{H}, \mathbb{H}', \mathbb{Q}, \mathbb{Q}'$. The last two terms on the right-hand sides of each of Equations (27) and (28) are "bias" terms accounting for how well our bridge and critic classes satisfy the realizability conditions. If the realizability conditions are violated up to a certain amount, then these two terms can be bounded accordingly. More generally, we can grow the function classes to make "bias" term vanish eventually and even to balance the "bias" and "variance" terms (see Example 9 below for sieve classes and neural network classes). If the realizability conditions hold exactly, then the two bias terms become zero, leading to the simplified bounds given below.

Corollary 8 (IPW and REG under realizability). *If* $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset, \mathbb{H}' \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset$, *then*

$$|\hat{J}_{\text{IPW}} - J| \leq \sup_{q\in\mathbb{Q}} |(\mathbb{E}_n - \mathbb{E})[q\pi Y]| + 2 \sup_{q\in\mathbb{Q}, h\in\mathbb{H}'} |(\mathbb{E}_n - \mathbb{E})[-q\pi h + \mathcal{T}h]|.$$

If $\mathbb{H} \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset, \mathbb{Q}' \cap \pi\mathbb{Q}_0^{\text{obs}} \neq \emptyset$, *then*

$$|\hat{J}_{\text{REG}} - J| \leq \sup_{h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[\mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q}', h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q(Y - h)]|.$$

Interestingly, and in stark contrast to the unconfounded setting, these bounds show that the GACE estimators $\hat{J}_{\text{REG}}, \hat{J}_{\text{IPW}}$ may converge to the true GACE even when the bridge function estimators $\hat{h}$, $\hat{q}$ do not converge to any valid observed bridge functions in $\mathbb{H}_0^{\text{obs}}$, $\mathbb{Q}_0^{\text{obs}}$, respectively. This is possible because the loss functions for the minimax estimators $\hat{h}, \hat{q}$ without stabilizers are related to the errors of estimators $\hat{J}_{\text{IPW}}$ and $\hat{J}_{\text{REG}}$. In our proofs for Theorem 7 and Corollary 8, we decompose the errors of $\hat{J}_{\text{IPW}}$ and $\hat{J}_{\text{REG}}$ into some vanishing empirical process terms and some terms directly related to the minimax objectives without stabilizers. Then we only use the fact that $\hat{h}, \hat{q}$ minimize those objectives to handle those terms, without requiring $\hat{h}, \hat{q}$ to converge to any valid bridge functions. This is how we manage to prove the consistency of the IPW and REG estimators without consistency of the bridge function estimators. Below we provide a simple example for such phenomenon.

Example 6. Suppose $\mathbb{H}_0^{\text{obs}} = \{h_0\}$, $\mathbb{H} = \{a_1 + a_2 h_0 : a_1 \in \mathbb{R}, a_2 \in \mathbb{R}\}$, $\mathbb{Q}_0^{\text{obs}} = \{q_0\}$, $\mathbb{Q}' = \{\pi q_0\}$[4]. Then we can solve the minimax optimization problem in Equation (12) to get the outcome bridge function estimator $\hat{h}$. The bridge function estimator can be then used to

4. The classes $\mathbb{H}$ and $\mathbb{Q}'$ are designed for analytic tractability. In practice, one may hardly specify such classes because they involve the ground truth h_0, q_0. But this simple example conveniently illustrates the key phenomenon.

form the estimator $\hat{J}_{\mathrm{REG}}$. To understand the behavior of $\hat{h}$, we can analyze the population analogue of the minimax optimization problem. It is easy to verify that the minimizers for the population minimax objective are given by

$$\arg\min_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} \left(\mathbb{E}\left[q\left(h-Y\right)\right]\right)^2 = \{a_1 + a_2 h_0 : a_1 = (1-a_2)\mathbb{E}[\pi q_0 h_0]/\mathbb{E}[\pi q_0], a_2 \in \mathbb{R}\}.$$

This means that the truth h_0 is only one of the possible population minimizers. Therefore, the estimator $\hat{h}$ that minimizes the empirical analog of the minimax objective generally would not converge to h_0. Nevertheless, it is easy to show that the double realizability conditions $\mathbb{H}_0^{\mathrm{obs}} \cap \mathbb{H} \neq \emptyset, \pi\mathbb{Q}_0^{\mathrm{obs}} \cap \mathbb{Q}' \neq \emptyset$ are satisfied. Then Corollary 8 guarantees that $|\hat{J}_{\mathrm{REG}} - J|$ converges to 0 as n goes to infinity. This provides an example where the GACE estimator is consistent but the bridge function estimator is not.

Theorem 9 (Analysis of DR estimator)**.** *For any $h_0 \in \mathbb{H}_0^{\mathrm{obs}}, q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$,*

$$\begin{aligned}|\hat{J}_{\mathrm{DR}} - J| \leq & \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q\pi Y - q\pi h + \mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |(\mathbb{E}_n - \mathbb{E})[-q\pi h + \mathcal{T}h]| + \\ & + \inf_{q\in\mathbb{Q}} \sup_{h\in\mathbb{H}'} |\mathbb{E}[(q_0 - q)\pi h]| + \inf_{h\in\mathbb{H}'} \sup_{h'\in\mathbb{H}} \sup_{q\in\mathbb{Q}} |\mathbb{E}[(q - q_0)\pi(h_0 - h' - h)]|, \\ |\hat{J}_{\mathrm{DR}} - J| \leq & \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q\pi Y - q\pi h + \mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q}',h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q(Y-h)]| \\ & + \inf_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} |\mathbb{E}[\{h_0 - h\}q]| + \inf_{q\in\mathbb{Q}'} \sup_{q'\in\mathbb{Q}} \sup_{h\in\mathbb{H}} |\mathbb{E}[(q_0\pi - q'\pi - q)(h - h_0)|.\end{aligned}$$

In each bound, the first two terms are "variance" terms and the last two terms are "bias" terms analogous to those in Theorem 7. Under realizability, these bias terms again become zero.

Corollary 10 (DR under realizability)**.** *If $\mathbb{Q}_0^{\mathrm{obs}} \cap \mathbb{Q} \neq \emptyset, \mathbb{H}_0^{\mathrm{obs}} \cap \{h : h - \mathbb{H} \subseteq \mathbb{H}'\} \neq \emptyset$,*

$$|\hat{J}_{\mathrm{DR}} - J| \leq \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q\pi Y - q\pi h + \mathcal{T}h]| + 4\sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |(\mathbb{E}_n - \mathbb{E})[-q\pi h + \mathcal{T}h]|.$$

If $\mathbb{H}_0^{\mathrm{obs}} \cap \mathbb{H} \neq \emptyset, \mathbb{Q}_0^{\mathrm{obs}} \cap \{q : \pi(q - \mathbb{Q}) \subseteq \mathbb{Q}'\} \neq \emptyset$, then

$$|\hat{J}_{\mathrm{DR}} - J| \leq \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q\pi Y - q\pi h + \mathcal{T}h]| + 4\sup_{q\in\mathbb{Q}',h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q(Y-h)]|.$$

Corollary 10 suggests that if *either* $\mathbb{Q}$ *or* $\mathbb{H}$ is well-specified, and the associated $\mathbb{H}'$ or $\mathbb{Q}'$ is rich enough relative to $\mathbb{H}$ or $\mathbb{Q}$, then $\hat{J}_{\mathrm{DR}}$ is consistent, provided these function classes have limited complexity to ensure convergence of the empirical process terms above. In particular, $\hat{J}_{\mathrm{DR}}$ is consistent when either empirical process conditions on $\mathbb{Q}, \mathbb{H}'$ or on $\mathbb{H}, \mathbb{Q}'$ hold. Moreover, when $\mathbb{H} = \{0\}, \mathbb{Q}' = \{0\}$ so that $\hat{J}_{\mathrm{DR}}$ reduces to $\hat{J}_{\mathrm{IPW}}$, the first equation exactly recovers the result for $\hat{J}_{\mathrm{IPW}}$ in Corollary 8. Analogously, when $\mathbb{Q} = \{0\}, \mathbb{H}' = \{0\}$, the second equation recovers the result for $\hat{J}_{\mathrm{REG}}$ in Corollary 8.

5.2 Convergence Rates of GACE Estimators for Common Function Classes

We next analyze the convergence rates of $\hat{J}_{\mathrm{REG}}$ and $\hat{J}_{\mathrm{IPW}}$ by further bounding the "variance" terms in Theorem 7 and corollary 8 in terms of the complexity of some common function classes. For simplicity and brevity, we set $\mathbb{H}' = \mathbb{H}$ and $\mathbb{Q}' = \pi\mathbb{Q}$. The results can easily be specialized to each of $\hat{J}_{\mathrm{REG}}$ or $\hat{J}_{\mathrm{IPW}}$ without these restrictions when only the corresponding pair of bridge and critic classes are realizable. Each result we present also implies the same convergence rate for $\hat{J}_{\mathrm{DR}}$ if we simply modify the critic class realizability condition to be as in Theorem 9 and corollary 10. We omit these for brevity.

Example 7 (VC-subgraph classes)**.** VC-subgraph classes are function classes whose subgraph sets have bounded VC dimension (van der Vaart, 1998, Chapter 19). For example, $\{\theta^\top \phi(\cdot) : \|\theta\|_2 \le 1, \theta \in \mathbb{R}^d\}$ has VC-subgraph dimension at most $d+1$.

Corollary 11. *Let $\mathbb{H}, \mathbb{Q}$ have finite VC-subgraph dimensions $V(\mathbb{H}), V(\mathbb{Q})$, respectively. Assume $\mathbb{H} \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset$, $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset$, and $\|\mathbb{H}\|_\infty, \|\mathbb{Q}\|_\infty, \|\pi(A \mid X)\|_\infty, \|Y\|_\infty < \infty$. Then, letting $O(\cdot)$ be the order w.r.t. $n, V(\mathbb{H})$, $V(\mathbb{Q})$, and δ, with probability $1-\delta$, we have*

$$\max\{|\hat{J}_{\text{REG}} - J|,\ |\hat{J}_{\text{IPW}} - J|\} = O(\sqrt{(V(\mathbb{H}) + V(\mathbb{Q}) + 1 + \log(1/\delta))/n}).$$

Example 8 (Nonparametric classes characterized by metric entropy)**.** Many common nonparametric classes cannot be characterized by VC-subgraph dimensions. Instead, their complexity is characterized by their metric entropies (Wainwright, 2019). For example, a Hölder ball $\mathbb{W}$ with smoothness level α and an input dimension d has metric entropy under infinity norm $\log \mathcal{N}(\varepsilon, \mathbb{W}, \|\cdot\|_\infty) = O(\varepsilon^{-d/\alpha})$. Nonparametric function classes in RKHS such as an RKHS with Matérn kernels and an RKHS with Gaussian kernels also have metric entropy characterizations (Kühn, 2011).

Corollary 12. *Suppose*

$$\max\{\log \mathcal{N}(\varepsilon, \mathbb{H}, \|\cdot\|_\infty), \log \mathcal{N}(\varepsilon, \mathbb{Q}, \|\cdot\|_\infty)\} \le c_0 \varepsilon^{-\beta}$$

for $\beta > 0$ and $c_0 > 0$. Further assume $\mathbb{H} \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset$, $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset$, and

$$\|\mathbb{H}\|_\infty, \|\mathbb{Q}\|_\infty, \|\pi(A \mid X)\|_\infty, \|Y\|_\infty < \infty.$$

Then, letting $O(\cdot)$ be the order w.r.t. n and δ, with probability $1-\delta$, we have

$$\max\left\{|\hat{J}_{\text{REG}} - J|,\ |\hat{J}_{\text{IPW}} - J|\right\} = \begin{cases} O(n^{-1/2} + \sqrt{\log(1/\delta)/n}) & \beta < 2 \\ O(n^{-1/2}\log(n) + \sqrt{\log(1/\delta)/n}) & \beta = 2 \\ O(n^{-1/\beta} + \sqrt{\log(1/\delta)/n}) & \beta > 2 \end{cases}$$

Corollary 12 states that the convergence rates of $\hat{J}_{\text{REG}}$ and $\hat{J}_{\text{IPW}}$ are determined by the worse of metric entropies of $\mathbb{H}$ and $\mathbb{Q}$. For example, when we use Hölder balls for both, the convergence rate is $O(n^{-\min(1/2, \alpha/d)})$. This implies that when these two function classes are Donsker classes (van der Vaart, 1998), *i.e.*, $\beta < 2$, both estimators have parametric convergence rates. But if either function class is non-Donsker, *i.e.*, $\beta \ge 2$, then the estimators typically have slower convergence rates.

Example 9 (Neural networks and linear sieves)**.** Many nonparametric function classes (*e.g.*, Hölder balls as in Example 8) are not amenable to direct optimization. A practical solution to this is to use sieves, which approximate the target nonparametric classes by classes that grow as the sample size n grows. For example, splines and polynomial sieves have been widely used in sieve estimators (Chen, 2007). Neural networks can be also regarded as sieves (Yarotsky, 2017; Gribonval et al., 2021; Suzuki, 2018). When $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$ are sieve classes they depend on n and this dependence is implicit in our notation. Since sieve classes approach target nonparametric classes only as sample size grows infinitely, it is not suitable

to impose realizability on these sieve classes for any finite n. Instead, we must consider the convergence rate of the "bias" terms in Theorems 7 and 9. In Appendix G, we show that GACE estimation error bounds scale as $O(n^{-\alpha/(2\alpha+d)})$ when we use linear sieves to approximate Hölder balls or when we use neural networks to approximate Sobolev balls, each with smoothness α in d input dimensions and when each limiting nonparametric class satisfies realizability.

6. Finite-Sample Analysis of Estimators without Stabilizers under Realizability and Closedness

In the previous section, we studied the convergence of our GACE estimators based on minimax bridge function estimators by assuming the realizability of $\mathbb{H}, \mathbb{Q}'$ (or of $\mathbb{H}', \mathbb{Q}$). In particular, our convergence analysis does not require the convergence of bridge function estimators. In this section, we show that under a different set of assumptions, realizability, and closedness, bridge function estimators do converge in a suitable notion, which then implies convergence of GACE estimators.

6.1 Convergence of Bridge Function Estimators

For any observed bridge functions $h_0 \in \mathbb{H}_0^{\text{obs}}$ and $q_0 \in \mathbb{Q}_0^{\text{obs}}$, we quantify the estimation errors of bridge function estimators $\hat{h}$ and $\hat{q}$ by $\|P_z(\hat{h} - h_0)\|_2^2$ and $\|\pi P_w(\hat{q} - q_0)\|_2$, which we call projected mean-squared errors (projected MSE). Projected MSE measures how much $\hat{h}$ and $\hat{q}$ violate Equations (8) and (9):

$$\|P_z(\hat{h} - h_0)\|_2^2 = \mathbb{E}[(\mathbb{E}[Y - \hat{h}(W, A, X) \mid Z, A, X])^2], \tag{29}$$

$$\|\pi P_w(\hat{q} - q_0)\|_2^2 = \mathbb{E}[(\mathbb{E}[\pi(A \mid X)\,(\hat{q}(Z, A, X) - 1/f(A \mid W, X)) \mid W, A, X])^2]. \tag{30}$$

Obviously, these estimation errors are invariant to the choice of h_0 and q_0 so they are particularly relevant when bridge functions are nonunique. Note that even when $\|P_z(\hat{h} - h_0)\|_2 \to 0$ and $\|\pi P_w(\hat{q} - q_0)\|_2 \to 0$, $\hat{h}$ and $\pi\hat{q}$ may not necessarily converge to any fixed limits in terms of the $\|\cdot\|_2$ norm since $\mathbb{H}_0^{\text{obs}}$ and $\mathbb{Q}_0^{\text{obs}}$ need not be singletons.

Theorem 13. *Consider the bridge function estimators without stabilizers, $\hat{h}, \hat{q}$ in Equations* (12) *and* (13)*.*

1. *Suppose $\mathbb{H} \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset$ and take some $h_0 \in \mathbb{H} \cap \mathbb{H}_0^{\text{obs}}$. If $P_z(\mathbb{H} - h_0) \subseteq \mathbb{Q}'$, then*

$$\|P_z(\hat{h} - h_0)\|_2^2 \leq 2 \sup_{q \in \mathbb{Q}', h \in \mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[\{y - h\}q]|.$$

2. *Suppose $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset$ and take some $q_0 \in \mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}}$. If $\pi P_w(\mathbb{Q} - q_0) \subseteq \mathbb{H}'$, then*

$$\|\pi P_w(\hat{q} - q_0)\|_2^2 \leq 2 \sup_{q \in \mathbb{Q}, h \in \mathbb{H}'} |(\mathbb{E}_n - \mathbb{E})[-q\pi h + \mathcal{T} h]|.$$

In Theorem 13 statement 1, we assume the realizability condition for $\mathbb{H}$, and the closedness condition $P_z(\mathbb{H} - h_0) \subset \mathbb{Q}'$. The latter closedness condition is invariant to choice of $h_0 \in \mathbb{H}_0^{\text{obs}} \cap \mathbb{H}$, by following Equation (8). This closedness condition intuitively indicates

that the critic class $\mathbb{Q}'$ is rich enough relative to the hypothesis class $\mathbb{H}$ and the operator P_z is smooth enough. A similar observation is made in Theorem 13 statement 2.

We can further bound the empirical process terms in Theorem 13 for specific bridge classes and critic classes by following the same calculations as in Section 5.2, which then leads to the projected MSE convergence rates of bridge function estimators without stabilizers. For example, when we use VC-subgraph classes for all function classes, the rate is $O(n^{-1/4})$, and when we use Hölder balls for all function classes, the rate is $O(\max(n^{-1/4}, n^{-\alpha/(2d)}))$.

6.2 Convergence of GACE Estimators

In this part, we bound the errors of GACE estimators based on Theorem 13. First, we first introduce Rademacher complexity, a widely used function class complexity measure (Wainwright, 2019).

Definition 2 (Rademacher Complexity)**.** Given a class $\mathcal{G}$ of functions of variables O, let

$$\mathcal{R}(\mathcal{G}) = \frac{1}{2^n} \sum_{\epsilon \in \{-1,+1\}^n} \sup_{g \in \mathcal{G}} \frac{1}{n} \sum_{i=1}^n \epsilon_i g(O_i).$$

Note that $\mathcal{R}(\mathcal{G})$ is random as it depends on our sample $O_1, \dots, O_n$.

The following gives the convergence rates of $\hat{J}_{\mathrm{REG}}$ and $\hat{J}_{\mathrm{IPW}}$ in terms of the Rademacher complexity of some function classes, as well as the projected MSEs of bridge function estimators $\hat{h}, \hat{q}$. We then derive the convergence rates of the doubly robust estimator $\hat{J}_{\mathrm{DR}}$.

Theorem 14. *Take arbitrary $h_0 \in \mathbb{H}_0^{\mathrm{obs}}$ and $q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$ and assume $\|\pi q_0\|_\infty < \infty$, $\|\pi \mathbb{Q}\|_2 < \infty$, $\|h_0\|_\infty < \infty$, $\|\mathbb{H}\|_2 < \infty$, $\|Y\|_\infty < \infty$. Then, for some universal constants c_1 and c_2, with probability at least $1 - \delta$,*

$$|\hat{J}_{\mathrm{REG}} - J| \le \mathcal{R}(\mathcal{T}\mathbb{H}) + \|\pi q_0\|_2 \|P_z(\hat{h} - h_0)\|_2 + c_1 \sqrt{(\log(c_2/\delta))/n}, \tag{31}$$

$$|\hat{J}_{\mathrm{IPW}} - J| \le \mathcal{R}(\pi \mathbb{Q}) + \|h_0\|_2 \|\pi P_w(\hat{q} - q_0)\|_2 + c_2 \sqrt{(\log(c_2/\delta))/n}. \tag{32}$$

In the right-hand sides of Equations (31) and (32), the Rademacher complexity terms bound the stochastic equicontinuity term due to nuisance plug-in (Newey and Mcfadden, 1994). The projected MSE terms correspond to the biases due to plug-in nuisance estimates. These terms can be bounded according to Theorem 13. Typically, these bias terms dominate in Equations (31) and (32). Therefore, the convergence rates of $\hat{J}_{\mathrm{REG}}, \hat{J}_{\mathrm{IPW}}$ are typically the same as the projected MSE convergence rates of bridge function estimators. Notably, these convergence rates are often slower than the convergence rates in Corollary 8. For example, for nonparametric classes with a metric entropy exponent β (see Example 8), the convergence rates in Theorem 14 become $O\left(\max(n^{-1/4}, n^{-1/2\beta})\right)$, while the convergence rates in Corollary 8 become $O\left(\max(n^{-1/2}, n^{-1/\beta})\right)$ according to Corollary 12. This shows that even when using the same minimax estimators without stabilizers, different assumptions (realizability in Corollary 8 and realizability plus closedness in Theorem 14) can lead to different convergence rates.

Theorem 15. *Take arbitrary $h_0 \in \mathbb{H}_0^{\mathrm{obs}}$ and $q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$ and assume conditions in Theorem 14. Then, for some universal constants c_1 and c_2, with probability at least $1 - \delta$,*

$$|\hat{J}_{\mathrm{DR}} - J| \le \mathcal{R}(\pi \mathbb{Q}\{y - \mathbb{H}\} + \mathcal{T}\mathbb{H}) + \sup_{q \in \mathbb{Q}} \|\pi(q_0 - q)\|_2 \|P_z(\hat{h} - h_0)\|_2 + c_1 \sqrt{\log(c_2/\delta)/n}$$

$$|\hat{J}_{\mathrm{DR}} - J| \le \mathcal{R}(\pi \mathbb{Q}\{y - \mathbb{H}\} + \mathcal{T}\mathbb{H}) + \sup_{h \in \mathbb{H}} \|h_0 - h\|_2 \|\pi P_w(\hat{q} - q_0)\|_2 + c_1 \sqrt{\log(c_2/\delta)/n}.$$

The error bounds above are typically dominated by the bias terms. Thus, Theorem 15 suggests that $\hat{J}_{\mathrm{DR}}$ is consistent when the projected MSE of either $\hat{h}$ *or* $\hat{q}$ vanishes to 0, which is reminiscent of double robustness in the unconfounded setting (*e.g.*, Bang and Robins, 2005).

7. Finite-Sample Analysis of Estimators with Stabilizers

In the previous sections, we studied the convergence of our GACE estimators based on minimax bridge function estimators *without* stabilizers (Section 4.2.1). In this section, we derive projected MSE bounds of the estimators *with* stabilizers (Section 4.2.2), *also* under realizability and closedness. Again, according to Theorems 14 and 15, these projected MSE bounds translate into convergence rates of resulting GACE estimators. We further show that the DR estimator is asymptotically normal and semiparametrically efficient when additionally assuming unique bridge functions and limiting the ill-posedness of associated inverse problems.

We first derive the error bounds for the minimax bridge function estimators $\hat{h}$, $\hat{q}$ *with* stabilizers (see Equations (16) and (17)). Our analysis is adapted[5] from the general theory in Dikkala et al. (2020). The analyses use the function class complexity measure called the *critical radius*, which is widely used in statistical learning theory (e.g., Bartlett et al., 2005; Wainwright, 2019).

Definition 3 (Critical Radius)**.** Given a function class $\mathcal{G}$, its empirical critical radius is the smallest $\eta > 0$ that satisfies $\mathcal{R}(\mathcal{G}^{|\eta}) \leq \eta^2/\|\mathcal{G}\|_\infty$ where $\mathcal{G}^{|\eta} = \{g \in \mathcal{G} : \mathbb{E}_n[g^2] \leq \eta^2\}$.

Theorem 16 (Convergence rate of $\hat{h}$ with stabilizer)**.** *Assume* $\|Y\|_\infty < \infty$, $\|\mathbb{H}\|_\infty < \infty$. *Let* $\hat{h}$ *be the minimax estimator with stabilizer in Equation* (16). *Fix any* $h_0 \in \mathbb{H}_0^{\mathrm{obs}}, h' \in \mathbb{H}$. *Assume:*

1. $\mathbb{Q}'$ *is symmetric and star-shaped.*
2. η_h *upper bounds the critical radii of* $\mathbb{Q}'$ *and the star hull of the following class* $\mathcal{G}_h$*:*

$$\mathcal{G}_h := \left\{(w,z,a,x) \mapsto \big(h(w,a,x) - h'(w,a,x)\big)\, q(z,a,x) : h \in \mathbb{H}, q \in \mathbb{Q}'\right\}.$$

Then, letting $O(\cdot)$ *be the order w.r.t* $\lambda, \delta, n, \eta_h$*, with probability* $1-\delta$*, we have*

$$\begin{aligned}\|P_z(\hat{h} - h_0)\|_2 = O\big((1+\lambda+\lambda^{-1})(\eta_h + \sqrt{1+\log(1/\delta)/n}) + \sup_{h\in\mathbb{H}} \inf_{q\in\mathbb{Q}'} \|q - P_z(h-h')\|_2 \\ + (\lambda(\eta_h + \sqrt{1+\log(1/\delta)/n}))^{-1}\|P_z(h'-h_0)\|_2^2 + \|P_z(h'-h_0)\|_2.\end{aligned}$$

The first term in the bound above is a statistical "variance" term. The second to fourth terms are "bias" terms due to approximate closedness or approximate realizability. When the realizability condition and closedness condition are violated up to a certain amount,

5. Our Theorems 16 and 18 are largely adapted from theory in Dikkala et al. (2020). There are mainly two differences: first, we consider unpenalized estimators for simplicity while Dikkala et al. (2020) considers penalized estimators; second, our problems for the action bridge function estimation do not fit into the standard conditional moment equation problem (as explained in footnote 1) so we need to adapt accordingly.

those terms can be bounded accordingly. For example, in Example 9 Cont'd, we consider linear sieve and neural network classes that do not exactly satisfy these conditions but the errors due to violation vanish in the limit. When the realizability and closedness conditions hold exactly, the error bound can be simplified to the bound given below.

Corollary 17 ($\hat{h}$ with stabilizer under realizability and closedness)**.** *Assume the conditions in Theorem 16, $h' = h_0 \in \mathbb{H}_0^{\text{obs}} \cap \mathbb{H} \neq \emptyset$, and $P_z(\mathbb{H} - h_0) \subseteq \mathbb{Q}'$. With probability $1-\delta$,*

$$\|P_z(\hat{h} - h_0)\|_2 = O((1 + \lambda + \lambda^{-1})(\eta_h + \sqrt{(1 + \log(1/\delta))/n})).$$

Corollary 17 states that under realizability and closedness, the convergence rate of $\hat{h}$ is determined by the critical radii η_h of the critic class $\mathbb{Q}'$ and the class $\mathcal{G}_h$ induced by both $\mathbb{H}$ and $\mathbb{Q}'$. In Section 7.1, we will further bound η_h for common function classes.

We can analogously bound the estimation error of the action bridge function estimator.

Theorem 18 (Convergence rate of $\hat{q}$ with stabilizer)**.** *Assume $\|\pi(A \mid X)/f(A \mid X, W)\|_2 < \infty$, $\|\pi\mathbb{Q}\|_\infty$. Let $\hat{q}$ be the minimax estimator with stabilizer in Equation* (17)*. Fix any $q_0 \in \mathbb{Q}_0^{\text{obs}}, q' \in \mathbb{Q}$. Assume:*

1. *$\mathbb{H}'$ is symmetric and star-shaped.*
2. *η_q upper bounds the critical radii of $\mathbb{H}'$ and the star hull of the following class $\mathcal{G}_q$:*

$$\mathcal{G}_q \coloneqq \{(w, z, a, x) \mapsto \pi(a \mid x)\big(q(z, a, x) - q'(z, a, x)\big)\, h(w, a, x) : q \in \mathbb{Q}, h \in \mathbb{H}'\}.$$

Then, letting $O(\cdot)$ be the order w.r.t. $\lambda, \delta, n, \eta_q$, with probability $1-\delta$, we have

$$\begin{aligned}\|\pi P_w(\hat{q} - q_0)\|_2 = O\big((1 + \lambda + \lambda^{-1})(\eta_q + \sqrt{1 + \log(1/\delta)/n}) + \sup_{q \in \mathbb{Q}} \inf_{h \in \mathbb{H}'} \|h - \pi P_w(q - q')\|_2 \\ + (\lambda(\eta_q + \sqrt{1 + \log(1/\delta)/n}))^{-1}\|\pi P_w(q' - q_0)\|_2^2 + \|\pi P_w(q' - q_0)\|_2.\end{aligned}$$

Corollary 19 ($\hat{q}$ with stabilizer under realizability and closedness)**.** *Assume the conditions in Theorem 18, $q' = q_0 \in \mathbb{Q}_0^{\text{obs}} \cap \mathbb{Q} \neq \emptyset$, and $\pi P_w(\mathbb{Q} - q_0) \subseteq \mathbb{H}'$. With probability $1-\delta$,*

$$\|\pi P_w(\hat{q} - q_0)\|_2 = O((1 + \lambda + \lambda^{-1})(\eta_q + \sqrt{\{1 + \log(1/\delta)\}/n})).$$

The bounds in Corollaries 17 and 19 suggest we can choose λ constant. In the rest of this section, we set $\lambda = 1$ for simplicity.

In Section 7.1, we discuss specific rates for some standard function classes. In Section 7.3, we use Theorems 14 and 15 to translate projected MSE bounds in Corollaries 17 and 19 into convergence rates of IPW, REG, and DR estimators of GACE. In Section 7.2, we give a tighter characterization of the doubly robust estimator under some additional assumptions.

7.1 Convergence Rates of Bridge Function Estimators for Common Function Classes

Corollaries 17 and 19 show that the convergence rates of minimax estimators $\hat{h}, \hat{q}$ with stabilizers (Section 4.2.2) are determined by critical radii of relevant function classes. We next provide concrete rates for these function classes. We focus on analyzing $\|P_z(\hat{h} - h_0)\|_2$ as an example. The projected MSE $\|\pi P_w(\hat{q} - q_0)\|_2$ can be analyzed analogously.

Example 7, Cont'd (VC-subgraph classes). We first consider VC-subgraph classes.

Corollary 20. *Assume that the conditions in Corollary 17 hold for some* $h_0 \in \mathbb{H} \cap \mathbb{H}_0^{\text{obs}}$. *Suppose* $\mathbb{H}$ *and* $\mathbb{Q}'$ *are VC-subgraph classes with VC-subgraph dimensions* $V(\mathbb{H}), V(\mathbb{Q}')$, *respectively. Then, letting* $O(\cdot)$ *be the order w.r.t.* $n, \delta, V(\mathbb{H}), V(\mathbb{Q}')$, *with probability* $1-\delta$,

$$\|P_z(\hat{h} - h_0)\|_2 = O(\sqrt{(V(\mathbb{Q}') + V(\mathbb{H}))\log n + 1 + \log(1/\delta)}/\sqrt{n}).$$

Example 8, Cont'd (Nonparametric classes characterized by metric entropy). We next analyze classes with limited metric entropy (*e.g.*, Hölder balls, Sobolev balls, and RKHSs).

Corollary 21. *Assume that the conditions in Corollary 17 hold for some* $h_0 \in \mathbb{H} \cap \mathbb{H}_0^{\text{obs}}$. *Suppose that* $\max(\log \mathcal{N}(\varepsilon, \mathbb{H}, \|\cdot\|_\infty),\ \log \mathcal{N}(\varepsilon, \mathbb{Q}', \|\cdot\|_\infty)) \leq c_0(1/\varepsilon)^\beta$. *Then, letting* $O(\cdot)$ *be the order with respect to* n, β, δ, *with probability* $1-\delta$,

$$\|P_z(\hat{h} - h_0)\|_2 = \begin{cases} O(n^{-1/(2+\beta)} + \sqrt{\log(1/\delta)/n}) & \beta < 2 \\ O(n^{-1/4}\log n + \sqrt{\log(1/\delta)/n}) & \beta = 2 \\ O(n^{-1/(2\beta)} + \sqrt{\log(1/\delta)/n}) & \beta > 2 \end{cases}$$

Example 9, Cont'd (Neural networks and linear sieves). If bridge and critic classes grow with sample size and approach nonparametric classes only in the limit, it is inappropriate to directly impose realizability and closedness on the bridge and critic classes. Instead, we must consider the "bias" terms due to approximate realizability and closedness. In Appendix G, we show projected MSE rates of $\tilde{O}(n^{-\alpha/(2\alpha+d)})$ when we use linear sieves to approximate Hölder balls or when we use neural networks to approximate Sobolev balls, each with smoothness α in d input dimensions and when each limiting nonparametric class satisfies both realizability and closedness.

We note that the projected MSE error bounds derived here for the bridge function estimators with stabilizers typically vanish at faster rates than the error bounds in Section 6 for estimators without stabilizers, when both assuming the same realizability and closedness conditions. For example, the estimators with stabilizers achieve projected MSE error rates $O(n^{-1/2})$ for VC-subgraph classes and $O(n^{-\alpha/(2\alpha+d)})$ for Hölder balls (corresponding to $\beta = d/\alpha < 2$ in Corollary 21). In contrast, the estimators without stabilizers can only achieve $O(n^{-1/4})$ and $O(\max(n^{-1/4}, n^{-\alpha/2d}))$ error rates, respectively.

Given the projected MSE error bounds for minimax estimators, we can also derive the convergence rates for the downstream GACE estimators by plugging them into the bounds for plug-in GACE estimators (Theorems 14 and 15). The convergence rates of GACE estimators are typically dominated by the corresponding projected MSE convergence rates. Therefore, the faster projected MSE error rates for the minimax estimators with stabilizers would also translate into faster error rates for the corresponding GACE estimators.

7.2 Tighter Analysis of the Doubly Robust Estimator

In Theorem 15, we showed that the convergence rate of the doubly robust estimator $\hat{J}_{\text{DR}}$ is determined by the projected MSE convergence rate of a single bridge function estimator, when conditions in either Corollary 17 *or* Corollary 19 are satisfied. In this section, we show that if the conditions in *both* theorems are satisfied, and the related inverse problems are not too ill-posed, then the doubly robust estimator $\hat{J}_{\text{DR}}$ can converge at a faster rate than the bridge function estimators.

We first introduce measures of the ill-posedness of inverse problems defined by the conditional moment equations in Equations (8) and (9), relative to the bridge classes $\mathbb{H}, \mathbb{Q}$.

Definition 4 (Measures of Ill-posedness)**.** Fix $h_0 \in \mathbb{H}_0^{\text{obs}} \cap \mathbb{H}, q_0 \in \mathbb{Q}_0^{\text{obs}} \cap \mathbb{Q}$. Define

$$P_u(h) = \mathbb{E}\left[h(W, Z, A, X) \mid U, A, X\right], \tag{33}$$

$$\tau_1^{\mathbb{H}} = \sup_{h \in \mathbb{H}} \frac{\|P_u(h - h_0)\|_2}{\|P_z(h - h_0)\|_2}, \quad \tau_1^{\mathbb{Q}} = \sup_{q \in \mathbb{Q}} \frac{\|P_u(\pi q - \pi q_0)\|_2}{\|P_w(\pi q - \pi q_0)\|_2}. \tag{34}$$

$$\tau_2^{\mathbb{H}} = \sup_{h \in \mathbb{H}} \frac{\|h - h_0\|_2}{\|P_z(h - h_0)\|_2}, \quad \tau_2^{\mathbb{Q}} = \sup_{q \in \mathbb{Q}} \frac{\|\pi q - \pi q_0\|_2}{\|P_w(\pi q - \pi q_0)\|_2}. \tag{35}$$

In this definition, we follow the convention $\frac{0}{0} = 0$.

To interpret the ill-posedness measures above, let us focus on $\tau_1^{\mathbb{H}}$ and $\tau_2^{\mathbb{H}}$ as examples and note that $\mathbb{E}\left[h - h_0 \mid Z, A, X\right] = \mathbb{E}\left[\mathbb{E}\left[h - h_0 \mid A, U, X\right] \mid Z, A, X\right]$ according to Assumption 1 condition (4). This means that we have two different levels of inverse problems: the first involves inverting $\mathbb{E}\left[h - h_0 \mid A, U, X\right]$ from $\mathbb{E}\left[h - h_0 \mid Z, A, X\right]$, and the second involves inverting $h - h_0$ from $\mathbb{E}\left[h - h_0 \mid Z, A, X\right]$ for $h \in \mathbb{H}$. The degrees of ill-posedness of these two levels relative to the hypothesis class $\mathbb{H}$ are quantified by $\tau_1^{\mathbb{H}}$ and $\tau_2^{\mathbb{H}}$, respectively. And, we have $1 \leq \tau_1^{\mathbb{H}} \leq \tau_2^{\mathbb{H}}$. Note that if the completeness condition in Equation (6) holds[6], $\tau_1^{\mathbb{H}}$ is invariant to the choice of $h_0 \in \mathbb{H}_0^{\text{obs}} \cap \mathbb{H}$, since $\mathbb{E}\left[h_0 \mid Z, A, X\right] = \mathbb{E}\left[Y \mid Z, A, X\right]$ and $\mathbb{E}\left[h_0 \mid U, A, X\right] = \mathbb{E}\left[Y \mid U, A, X\right]$ according to Equations (2) and (8). So $\tau_1^{\mathbb{H}}$ may be finite even when the bridge function is not uniquely identified. In contrast, $\tau_2^{\mathbb{H}}$ is $+\infty$ when $\mathbb{H}_0^{\text{obs}} \cap \mathbb{H}$ is not a singleton. Moreover, note that the definition of $\tau_1^{\mathbb{H}}$ and $\tau_2^{\mathbb{H}}$ depend on the choice of hypothesis class $\mathbb{H}$ (which can change with n). When specialized to linear sieves (Appendix G.1), these measures are related to the sieve ill-posedness measures introduced in Blundell et al. (2003); Chen and Pouzo (2012).

The ill-posedness measures in Definition 4 directly influence the convergence behavior of minimax bridge function estimators in different error measures. For example, if $\tau_1^{\mathbb{H}} < +\infty$ and $\tau_2^{\mathbb{H}} < +\infty$ for all n, then for the minimax estimator $\hat{h}$, the error bound on $\|P_z(\hat{h} - h_0)\|$ derived in Corollary 17 easily translates into bounds on $\|\mathbb{E}[\hat{h} - h_0 \mid A, U, X]\|_2$ and $\|\hat{h} - h_0\|_2$, with additional inflation factors of $\tau_1^{\mathbb{H}}$ and $\tau_2^{\mathbb{H}}$ respectively. So larger $\tau_1^{\mathbb{H}}$ and $\tau_2^{\mathbb{H}}$, *i.e.*, more ill-posed inverse problems, correspond to larger error in terms of $\|\mathbb{E}[\hat{h} - h_0 \mid A, U, X]\|_2$ and $\|\hat{h} - h_0\|_2$. The ill-posedness measures depend both on the data generating process and on the choice of hypothesis classes $\mathbb{H}, \mathbb{Q}$. As is discussed in Blundell et al. (2003) for linear sieves, larger hypothesis classes typically correspond to bigger ill-posedness measures.

Next, we slightly revise our doubly robust estimator by cross-fitting the bridge function estimators, a technique that has been widely used to remove restrictive Donsker conditions on hypothesis classes (*e.g.*, Zheng and van der Laan, 2011; Chernozhukov et al., 2018). For simplicity, we focus on two-fold cross-fitting; it is straightforward to extend to more folds.

Definition 5 (Cross-fitted DR Estimator)**.** Randomly split the whole sample into two halves denoted as $\mathcal{D}_0, \mathcal{D}_1$. Then for $j = 0, 1$, fit $\hat{h}^{(j)}, \hat{q}^{(j)}$ based on $\mathcal{D}_j$ according to Equations (16)

6. Under this completeness condition, we have $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$ according to Appendix B Lemma 24. Thus any $h_0 \in \mathbb{H}_0^{\text{obs}} \cap \mathbb{H}$ must satisfy both Equations (2) and (8).

and (17), respectively. Finally, denoting $I_i = \mathbb{I}[i \in \mathcal{D}_0]$, redefine the DR estimator by

$$\hat{J}_{\mathrm{DR}} = \frac{1}{n}\sum_{i=1}^{n} \pi(A_i|X_i)\hat{q}^{(I_i)}(Z_i, A_i, X_i)(Y_i - \hat{h}^{(I_i)}(W_i, A_i, X_i)) + (\mathcal{T}\hat{h}^{(I_i)})(W_i, X_i).$$

Now we derive the property of the cross-fitted DR estimator in terms of the ill-posedness measures $\tau_1^{\mathbb{H}}, \tau_1^{\mathbb{Q}}$ in Equation (34), *without* assuming unique bridge functions.

Theorem 22 (DR estimator without bridge function uniqueness)**.** *Suppose the conditions in Corollaries 17 and 19 hold. Then, letting $O(\cdot)$ be the order w.r.t. n, $\tau_1^{\mathbb{H}}$, $\tau_1^{\mathbb{Q}}$, η_h, η_q, and δ, with probability at least $1-\delta$, we have*

$$|\hat{J}_{\mathrm{DR}} - J| = O\left(\tau_1^{\mathbb{H}}\tau_1^{\mathbb{Q}}\eta_h\eta_q + \sqrt{\log(1/\delta)/n}\right). \tag{36}$$

Note that the convergence rate in Theorem 22 can be faster than the rate in Corollary 10. For example, suppose all relevant function classes are Hölder balls of order α in d dimensions. In this case, the convergence rate of $\hat{J}_{\mathrm{DR}}$ based on bridge function estimators *with* stabilizers given in Theorem 22 reduces to $\max(\tau_1^{\mathbb{H}}\tau_1^{\mathbb{Q}} n^{-2\alpha/(2\alpha+d)}, n^{-1/2})$, while the convergence rate in Corollary 10 for the estimator *without* stabilizers is $\max(n^{-\alpha/d}, n^{-1/2})$. The former converges faster than the latter if the ill-posedness measures satisfy $\tau_1^{\mathbb{H}}\tau_1^{\mathbb{Q}} = o\left(n^{\alpha(d-2\alpha)/(2\alpha d+d^2)}\right)$. This faster convergence is possible primarily because Theorem 22 assumes both realizability and closedness, under which the bridge function estimators with stabilizers converge to certain valid observed bridge functions $h_0 \in \mathbb{H}_0^{\mathrm{obs}}, q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$ in terms of errors $\|\mathbb{E}[\hat{h} - h_0 \mid A, U, X]\|_2$, $\|\mathbb{E}\left[\pi\left(\hat{q} - q_0\right) \mid A, U, X\right]\|_2$. In contrast, Corollary 10 assumes only realizability, under which the bridge function estimators may not converge to any true bridge function at all. Actually, under realizability and closedness, we can also tailor Theorem 22 to estimators without stabilizers by leveraging the projected MSE bounds in Theorem 13 and the ill-posedness measures $\tau_1^{\mathbb{H}}, \tau_1^{\mathbb{Q}}$ defined in Equation (34). But since the projected MSE rates of bridge function estimators without stabilizers are typically slower than their counterparts with stabilizers, the convergence rate of the resulting DR estimator will again be slower than the rate in Theorem 22.

Next, we show $\hat{J}_{\mathrm{DR}}$ is asymptotically normal when additionally assuming unique bridge functions. Note that with nonunique bridge functions, the minimax estimators $\hat{h}, \hat{q}$ may not converge to any fixed limit, even when $\|\mathbb{E}[\hat{h} - h_0 \mid A, U, X]\|_2$ and $\|\mathbb{E}\left[\pi\left(\hat{q} - q_0\right) \mid A, U, X\right]\|_2$ vanish to 0 for all $h_0 \in \mathbb{H}_0, q_0 \in \mathbb{Q}_0$. In this nonunique case, even if $\tau_1^{\mathbb{H}}\tau_1^{\mathbb{Q}}\eta_h\eta_q = o(n^{-1/2})$ so that $\hat{J}_{\mathrm{DR}}$ is $\sqrt{n}$-consistent, certain challenging stochastic equicontinuity term (*i.e.*, Equation (63) in Appendix I.7) is non-negligible and may distort the asymptotic normality. This intractable term makes it very difficult to derive the asymptotic distribution of $\hat{J}_{\mathrm{DR}}$. In the following lemma, we further assume that the bridge classes single out *unique* bridge functions, and show $\hat{J}_{\mathrm{DR}}$ has an asymptotically normal distribution with a closed-form variance.

Theorem 23 (DR estimator with unique bridge functions)**.** *Suppose that conditions in Corollaries 17 and 19 hold, $\mathbb{H} \cap \mathbb{H}_0^{\mathrm{obs}} = \{h_0\}$ and $\pi\mathbb{Q} \cap \pi\mathbb{Q}_0^{\mathrm{obs}} = \{\pi q_0\}$. If $\tau_2^{\mathbb{H}}\eta_h = o(1)$, $\tau_2^{\mathbb{Q}}\eta_q = o(1)$, and $\min(\tau_2^{\mathbb{H}}, \tau_2^{\mathbb{Q}})\eta_q\eta_h = o(n^{-1/2})$, then*

$$\sqrt{n}(\hat{J}_{\mathrm{DR}} - J) \xrightarrow{d} \mathcal{N}(0, V_{\mathit{eff}}),$$

$$V_{\mathit{eff}} = \mathbb{E}\left[\left(\tilde{\phi}_{\mathrm{DR}}(O; h_0, q_0) - J\right)^2\right].$$

For simplicity, we focus here on the case of exact realizability and closedness of the bridge and critic classes. If, following Example 9, we use classes that grow with n and slowly approach nonparametric classes that satisfy realizability and closedness, then we can state the unique bridge function condition to be in terms of the limiting classes. Note that the condition $\min(\tau_2^{\mathbb{H}}, \tau_2^{\mathbb{Q}})\eta_h\eta_q = o(n^{-1/2})$ in Theorem 23 is generally incomparable to $\tau_1^{\mathbb{H}}\tau_1^{\mathbb{Q}}\eta_h\eta_q = o(n^{-1/2})$ in Theorem 22, while both conditions are implied by $\tau_2^{\mathbb{H}}\tau_2^{\mathbb{Q}}\eta_h\eta_q = o(n^{-1/2})$.

Finally, we remark that the asymptotic variance V_{eff} in Theorem 23 coincides with the semiparametric efficiency bound derived in Appendix D under the assumption that $\mathbb{H}_0^{\mathrm{obs}} = \{h_0\}$ and some additional regularity conditions. These efficiency results extend the asymptotic efficiency results in Cui et al. (2023). See Appendix D for details.

7.3 Comparing Estimators with and without Stabilizers

So far we have analyzed theoretical guarantees for the minimax bridge function estimators with and without stabilizers and the corresponding GACE estimators. The theory is based on different assumptions and analyses: Section 5 assumes double realizability for both the bridge function class and the critic function class and directly derives the error bounds for the GACE estimators with no stabilizers, while Sections 6 and 7 assume realizability for the bridge function class and closedness for the critic function class and analyze the error bounds for the minimax bridge function estimators and plug-in GACE estimators.

Under the realizability and closedness assumptions, our analyses in Sections 6 and 7 show that the minimax bridge function estimators with stabilizers have clear advantages because their projected MSE error bounds have faster convergence rates than their counterparts without stabilizers. The faster rates for bridge function estimators can also translate into faster error rates for the downstream plug-in GACE estimators. Moreover, the faster rates of the bridge function estimators with stabilizers also mean that they are more likely to satisfy the product rate conditions in Theorem 23, so that the resulting double robust estimator is more likely to achieve asymptotic normality.

However, this does not mean that using stabilizers always leads to better convergence rates. In fact, the GACE estimators $\hat{J}_{\mathrm{REG}}, \hat{J}_{\mathrm{IPW}}$ based on bridge function estimators without stabilizers can sometimes achieve better convergence rates under the double realizability conditions. For example, when all relevant function classes are Hölder balls, the error convergence rates for $\hat{J}_{\mathrm{REG}}, \hat{J}_{\mathrm{IPW}}$ without stabilizers are $O\left(\max(n^{-1/2}, n^{-\alpha/d})\right)$ according to Corollaries 8 and 12. In contrast, the error bounds for estimators $\hat{J}_{\mathrm{REG}}, \hat{J}_{\mathrm{IPW}}$ with stabilizers require the realizability and closedness conditions, and when instantiated to Hölder balls, the error bounds are the slower $\tilde{O}(n^{-\alpha/(2\alpha+d)})$ rates. This shows the potential advantages of minimax bridge function estimators *without* stabilizers in GACE estimation. However, we remark that this comparison is also subtle because the two different convergence rates are based on quite different assumptions and analyses.

Therefore, there is no universal dominance of one approach over the other in terms of theoretical error rates. In our empirical study in Section 9, we find that these two different approaches tend to produce similar estimation results in nearly all simulation and real data

experiments, suggesting that both are viable in practice. In practice, the practitioners can try both but if they want to start with a particular type of estimators, they can choose the estimators with stabilizers for the fast convergence rate guarantees and statistical inference guarantees under the corresponding realizability and closedness conditions.

8. Related Literature

We review related work on negative controls and on instrumental variables and minimax estimation.

8.1 Negative Controls

Miao et al. (2018a) first develop sufficient conditions for identifying counterfactual distributions with negative controls. Shi et al. (2020) focus on multiply robust estimation of the average treatment effect in a discrete setting. Miao et al. (2018b) propose to estimate the average treatment effect by first estimating the outcome bridge function using standard Generalized Method of Moments (GMM) methods (Hansen, 1982). Cui et al. (2023) derive the semiparametric efficiency bound for average treatment effect based on both the outcome bridge function and action bridge function. Most previous work focuses on the average treatment effect with discrete treatment, assumes completeness conditions, requires unique bridge functions for estimation[7], and restricts attention to parametric bridge-function models.. In contrast, our paper studies a generalized average causal effect with general actions (continuous or discrete), develops new identification results without imposing completeness conditions, allows for nonunique bridge functions in estimation, and proposes flexible minimax learning approaches to accommodate both nonparametric and parametric estimation of bridge functions. More details on comparing identification strategies is given in Appendix B.

Deaner (2021) studies the identification of the counterfactual mean on the treated, based on the outcome bridge function or the action bridge function. They also require completeness conditions and their estimation is restricted to two-stage sieve estimators for the outcome bridge function (Chen and Pouzo, 2012, 2015). Later, Singh (2020); Mastouri et al. (2021) propose to use a two-stage kernel estimator for the outcome bridge function, and Xu et al. (2021) propose a neural network extension with adaptive features, all of which require a unique outcome bridge function and assume completeness conditions. In contrast, our paper considers minimax estimation approaches accommodating a wider array of nonparametric machine learning methods, and assumes weaker assumptions.

A few recent papers also propose minimax estimators for the bridge functions. Ghassami et al. (2021) studies doubly robust estimation of average treatment effect, and Qi et al. (2021) studies evaluating and learning optimal treatment regimes. These two papers focus on discrete treatments and also require completeness conditions. Moreover, they study

7. Shi et al. (2020) studies the discrete data setting in Example 3 and considers negative controls that can take more values than the unobserved confounders. In this case, the bridge functions are nonunique. They briefly describe two possible approaches to tackle nonunique bridge functions: one by considering the pseudo-inverse of certain non-invertible conditional probability matrix, and the other by coarsening the values of negative controls according to the number of values of the unobserved confounders. But they do not develop the two approaches in full detail. Instead, they still focus on unique bridge functions in estimation.

minimax estimators *with* stabilizers, whose theoretical guarantees are in line with some of our results in Section 7. Under both completeness and uniqueness conditions, Mastouri et al. (2021) proposes a Maximum Moment Restriction approach, which is a special example of our minimax estimation *without* stabilizers. Both Ghassami et al. (2021) and Mastouri et al. (2021) focus on RKHS hypothesis classes. In contrast, our paper assumes weaker assumptions, analyzes both types of minimax estimators, and considers a wider variety of hypothesis classes. Like Ghassami et al. (2021); Qi et al. (2021), our analysis for minimax estimators *with* stabilizers also builds on Dikkala et al. (2020), but our analysis for the minimax estimators *without* stabilizers is completely different from other literature. Moreover, our definitions of ill-posedness measures $\tau_1^{\mathbb{H}}, \tau_1^{\mathbb{Q}}$ that allow for nonunique bridge functions are also new.

8.2 Instrumental Variables and Minimax Estimation

Estimation of the outcome bridge function h_0 is closely related to the nonparametric instrumental variable (IV) regression problem (e.g., Newey, 2013). Our paper nonetheless differs significantly from the previous literature. First, the conditional moment equation for the action bridge function q_0 in Equation (9) is distinct from the IV conditional moment equation since it includes an unknown density. See the discussions in Section 4.2.1. Second, our target estimand is the GACE rather than bridge functions, so most of our analysis is different from previous literature on the estimation of the IV regression function. For example, our new analysis shows that our GACE estimator with bridge functions without stabilizers remains consistent even when the bridge function estimators are inconsistent (Example 6). Despite these differences, as we mentioned, the estimation of h_0 itself is similar to the IV regression problem, so we next summarize the relevant literature on IV regression.

The nonparametric IV regression problem is typically cast into conditional moment equations (*e.g.*, Ai and Chen, 2003). One classical approach to this problem is a nonparametric analogue of the two-stage least squares (2SLS) method based on sieve estimators (Newey and Powell, 2003) or kernel density estimators (Darolles et al., 2010; Hall and Horowitz, 2005). Another classical approach is to use sieves to convert the conditional moments into unconditional moments of increasing dimension (*e.g.*, Chen, 2007; Chen and Qiu, 2016), and then combine all unconditional moments via standard GMM method (Hansen, 1982). Later, Hartford et al. (2017); Singh et al. (2019) extend the two-stage approach by employing neural network density estimator or conditional mean embedding in RKHS respectively in the first stage. It remains unclear how to incorporate more general hypothesis classes while still providing rigorous theoretical guarantees for nonparametric estimators.

Recently, there has been intense interest in minimax approaches that reformulate the nonparametric IV regression problem via Equation (11), Equation (15), or other related variants (Lewis and Syrgkanis, 2018; Zhang et al., 2020; Dikkala et al., 2020; Chernozhukov et al., 2020; Liao et al., 2020; Bennett et al., 2019; Bennett and Kallus, 2020; Muandet et al., 2020). The resulting minimax formulation is more analogous to the empirical risk minimization framework predominant in machine learning, which permits more general function classes. Notably, Dikkala et al. (2020) provide a thorough analysis of the convergence rates of their minimax estimators. Our theoretical analysis for the minimax estimators with stabilizers builds on their analysis (see Corollaries 17 and 19). Minimax estimators have also been employed in some other contexts, such as causal-effect estimation (Chernozhukov et al.,

	IPW	IPW(sta)	REG	REG(sta)	DR	DR(sta)	Linear	Plain REG	DR (no W,Z)
MSE	0.0399	0.0353	0.0798	0.0892	0.0771	0.0844	0.1348	0.7718	0.7771
Squared Bias	0.0005	0.0322	0.0693	0.0794	0.0735	0.0800	0.1289	0.7660	0.7732
Variance	0.0394	0.0031	0.0105	0.0098	0.0037	0.0044	0.0059	0.0059	0.0039

Table 2: $g(t) = \sin(t)$, $n = 400$.

	IPW	IPW(sta)	REG	REG(sta)	DR	DR(sta)	Linear	Plain REG	DR (no W,Z)
MSE	0.0324	0.0288	0.0724	0.0757	0.0691	0.0739	0.1654	1.0011	1.0015
Squared Bias	0.0063	0.0276	0.0658	0.0697	0.0679	0.0721	0.1639	0.9984	0.9999
Variance	0.0262	0.0012	0.0066	0.0060	0.0012	0.0018	0.0015	0.0027	0.0015

Table 3: $g(t) = \sin(t)$, $n = 1200$.

	IPW	IPW(sta)	REG	REG(sta)	DR	DR(sta)	Linear	Plain REG	DR (no W,Z)
MSE	0.0202	0.0469	0.0037	0.0033	0.0029	0.0025	0.0655	0.5037	0.5142
Squared Bias	0.0124	0.0430	0.0004	0.0002	0.0006	0.0003	0.0514	0.4966	0.5092
Variance	0.0077	0.0038	0.0033	0.0031	0.0023	0.0022	0.0141	0.0072	0.0050

Table 4: $g(t) = t^3$, $n = 400$.

	IPW	IPW(sta)	REG	REG(sta)	DR	DR(sta)	Linear	Plain REG	DR (no W,Z)
MSE	0.0122	0.0431	0.0026	0.0018	0.0025	0.0012	0.0915	0.8313	0.8352
Squared Bias	0.0059	0.0417	0.0008	0.0001	0.0016	0.0002	0.0894	0.8265	0.8327
Variance	0.0064	0.0014	0.0018	0.0017	0.0009	0.0009	0.0021	0.0049	0.0025

Table 5: $g(t) = t^3$, $n = 1200$.

2020; Hirshberg and Wager, 2019) and reinforcement learning policy evaluation (Feng et al., 2019; Yang et al., 2020; Uehara et al., 2021), when all confounders are measured. Unlike this literature, our paper focuses on the more challenging setting with unmeasured confounders and addresses the confounding problem by using negative controls.

9. Numerical Experiments

In this section, we illustrate our estimators with and without stabilizers in both numerical simulations and a real-data analysis. In particular, we feature the performance of minimax bridge function estimators using neural networks. The code for all experiments can be found in `https://github.com/ZHzhang01/Minimax_JMLR`.

9.1 Simulation Study

In our simulations, we adjust the data generating process (DGP) used in Cui et al. (2023) to allow for multi-dimensional variables and nonlinear bridge functions. Concretely, we first start with the DGP in Cui et al. (2023), which uses one-dimensional variables, to generate multi-dimensional variables $U, X', Z', W' \in \mathbb{R}^d$ with $d = 60$ (and $A \in \{0, 1\}$). This DGP ensures the existence of bridge functions $h_0(W', A, X')$ and $q_0(Z', A, X')$ that are linear in (W', A, X') and (Z', A, X') respectively. To introduce nonlinearity, we transform (W', Z', X') into (W, Z, X) via $X = g(GX')$, $Z = g(GZ')$, $W = g(GW')$ where $G \in \mathbb{R}^{d\times d}$ is an invertible matrix and $g(\cdot)$ is a nonlinear function applied elementwise to GX', GZ', GW' respectively.

	IPW	IPW(sta)	REG	REG(sta)	DR	DR(sta)	Linear	Plain REG	DR (no W,Z)
MSE	0.0118	0.0153	0.0148	0.0139	0.0083	0.0081	0.0447	0.0183	0.0151
Squared Bias	0.0024	0.0012	0.0006	0.0006	0.0005	0.0005	0.0000	0.0000	0.0000
Variance	0.0094	0.0141	0.0142	0.0133	0.0077	0.0077	0.0447	0.0183	0.0151

Table 6: $g(t) = \sin(t)$, $n = 400$, ablation study-I.

	IPW	IPW(sta)	REG	REG(sta)	DR	DR(sta)	Linear	Plain REG	DR (no W,Z)
MSE	0.0330	0.0211	0.0170	0.0155	0.0168	0.0132	0.0594	0.0199	0.0173
Squared Bias	0.0235	0.0061	0.0054	0.0040	0.0067	0.0042	0.0035	0.0000	0.0000
Variance	0.0095	0.0150	0.0116	0.0115	0.0101	0.0090	0.0558	0.0199	0.0173

Table 7: $g(t) = t^3$, $n = 400$, ablation study-II.

In the final data, we only observe (W, Z, X) but not (W', Z', X'). The corresponding bridge functions $h_0(W, A, X)$ and $q_0(Z, A, X)$ exist (when $g()$ is invertible) and are nonlinear. See Appendix H for details on parameter specifications.

Our goal is to estimate the counterfactual mean parameter $J = \mathbb{E}[Y(1)]$, which is an example of the GACE parameter with $\pi(a \mid x) = \mathbb{I}(a = 1)$ (see Example 1). We compare the performance of four different types of estimators. *(i)* The first type of estimators (`IPW, IPW(sta), REG, REG(sta), DR, DR(sta)`) are our proposed IPW, REG, and DR estimators based on *nonlinear* minimax bridge function estimators without and with stabilizers. To construct the outcome bridge function estimator $\hat{h}$ (both with and without stabilizers), we use three-layer neural networks as the bridge class $\mathbb{H}$, and an RKHS with an RBF kernel as the critic class $\mathbb{Q}'$. To construct the action bridge function estimator $\hat{q}$ (both with and without stabilizers), we again use three-layer neural networks as the bridge class $\mathbb{Q}$ and an RKHS as the critic class $\mathbb{H}'$. Throughout we set the stabilizer parameter as $\lambda = 1$. See Appendix H for more details of model specifications and hyperparameter choices. *(ii)* The second type of estimator (`Linear`) is the DR estimator based on *linear* minimax bridge function estimators without stabilizers, where all bridge and critic classes are simple linear classes described in Section 4.3.1 with basis functions $\phi(z, a, x) = \tilde{\phi}(z, a, x) = (z^\top, a, x^\top)^\top$ and $\psi(w, a, x) = \tilde{\psi}(w, a, x) = \left(w^\top, a, x^\top\right)^\top$. This estimator has a closed-form given in Lemma 5. *(iii)* The third baseline is a regression-adjustment estimator (`Plain REG`) that fits a neural network to approximate the conditional outcome regression function $m(a, x) = \mathbb{E}[Y \mid A = a, X = x]$ using only the treatment A and covariates X, without accounting for the unmeasured confounding. In particular, we train the same multilayer network architecture used in our bridge-function models, but we restrict its inputs to (A, X) and optimize it by minimizing mean squared error to predict Y; negative-control variables (W, Z) are deliberately excluded from both the inputs and the loss. *(iv)* The fourth type of estimator (`DR (no W,Z)`) is the regular doubly robust estimator in the unconfounded setting. Like the `Plain REG` estimator, this estimator also ignores the unmeasured confounding and does not use the negative controls.

In Tables 2 to 5, we show the performance of different estimators over 1000 replications of experiments, with sample sizes $n = 400$ and 1200, and variable transformation maps $g(t) = \sin(t)$ and $g(t) = t^3$, respectively. We report the normalized MSE, normalized squared bias, and normalized variance, each scaled by the squared true estimand, over the $1,000$ replications.

Across all settings, the benchmark estimators `Plain REG` and `DR (no W,Z)` use flexible neural network nuisance estimators but they ignore the unmeasured confounding, resulting in high bias and MSE. The benchmark `Linear` does account for unmeasured confounding but uses restrictive linear bridge function estimators, which also lead to high bias and MSE because of significant misspecification. In contrast, our proposed minimax `IPW`, `REG`, and `DR` estimators learn nonlinear bridge functions from flexible neural networks together with RKHS critics without and with stabilizers (`sta`), and explicitly use the negative controls to adjust for unmeasured confounding. These six estimators achieve substantially lower errors. In particular, `DR` and `DR(sta)` achieve robust performance: they typically attain the lowest or near-lowest MSE, combining bias reduction with improved stability, especially in the case $g(t) = t^3$. Moreover, the REG and DR estimators perform broadly similarly with and without stabilizers, whereas stabilization has a more noticeable effect on IPW under $g(t) = t^3$.

To isolate the role of unmeasured confounding, we further run ablation studies designed to remove unmeasured confounding. Concretely, in this setting, the treatment A is drawn from a logistic model that depends only on the pre-transformation covariates X', and all other variables (W, Z, U) and the outcome Y are generated afterward given (A, X'). This construction ensures that $A \perp (Y(1), Y(0)) \mid X'$. For the observed X, ignorability $A \perp (Y(1), Y(0)) \mid X$ holds exactly under the cubic transformation; under the sine transformation, global non-invertibility induces only negligible residual dependence at the reported precision. We defer the details to Appendix H and report the estimation results in Tables 6 and 7. Under these settings, Plain REG and DR (no W, Z) exhibit essentially zero squared bias, consistent with exact or near-exact ignorability, but their variances keep their MSEs above that of DR(sta), which attains the lowest MSE in both Tables 6 and 7 (0.0081 and 0.0132 , respectively). The proposed DR and DR(sta) estimators therefore remain highly competitive and, in these two settings, outperform the unconfounded baselines in terms of MSE.

9.2 Real Data Analysis

We also apply our methods to the right heart catheterization (RHC) dataset from the Study to Understand Prognoses and Preferences for Outcomes and Risks for Treatments (SUPPORT). The treatment variable A indicates whether a patient received an RHC within 24 hours of admission or not. The binary outcome Y indicates whether the patient survives for 30 days after admission. This dataset includes 5735 participants with 72 covariates. See Hirano et al. (2003) Table 2 for more details.

A number of papers have analyzed the RHC dataset assuming no unmeasured confounding (Hirano et al., 2003; Tan, 2006; Vermeulen and Vansteelandt, 2015). In contrast, Cui et al. (2023) and Tchetgen et al. (2020) further address unmeasured confounding that may arise as the patients self-selected or were selected into the treatment. They view some measurements in a blood test prior to the treatment as noisy snapshots of the unobserved underlying physiological state. In particular, they propose a practical strategy to determine some plausible proxies, ending up with the negative control actions $Z = (\texttt{pafi1}, \texttt{paco21})$ and negative control outcome proxies $W = (\texttt{ph1}, \texttt{hema1})$ (see Appendix H for detailed descriptions). We follow their practice and use the same proxies in our analysis. As in those

papers, the validity and informativeness of these proxies remain substantive identifying assumptions and cannot be established from the observed data alone.

Table 8 shows point estimates and 95% confidence intervals for the average treatment effect. The first two columns correspond to our proposed `DR` and `DR (sta)`, respectively. The bridge classes are also specified as three-layer neural networks, and the critic classes are specified analogously to those in Section 9.1 (albeit with different covariate and negative control dimensions). The third column `Linear(closed-form)` corresponds to the linear case with closed-form as in (22). The fourth column `Linear(moment-based)` uses the linear moment estimator for the outcome bridge function used in Cui et al. (2023). Moreover, in Vermeulen and Vansteelandt (2015), they assume no unmeasured confounding and report an estimate of -0.0612 with standard error 0.0141. All estimates in Table 8 suggest that applying RHC causes higher 30-day mortality (or lower survival probability) than not applying RHC. We can observe that when using simple linear bridge function estimators (the third and fourth columns), the estimates are similar to the results previously reported in Vermeulen and Vansteelandt (2015). However, when using more flexible models for the bridge functions (the first two columns), the results suggest that applying RHC might not be as harmful as previous methods suggest.

Because the ground-truth effect is unknown in this real-data study, it is impossible to conclusively determine whether the estimates from our proposed methods based on flexible minimax bridge function estimation are closer to the true effects than simpler linear parametric estimators. Instead, we conduct some additional sensitivity analyses to check the robustness of our proposed method and the linear parametric estimators with respect to the model specifications. In Table 9, we assess the robustness of the minimax estimators with neural network bridge functions by modifying the neural-network architecture. Specifically, we separately halve the width and the depth of the networks. The estimates remain quite stable across these specifications. In Table 10, we further change the specifications of the linear (moment-based) bridge function estimators, by adding interaction terms, quadratic terms, or both, using the negative controls and a subset of the covariates. These seemingly mild changes in parametric form lead to large shifts in the treatment effect estimates and confidence intervals. This illustrates that the linear specification is sensitive to how the functional form is hand-crafted, whereas the neural-network-based minimax DR estimator is more stable.

We further conduct two complementary robustness analyses specific to the RHC application. First, the allocation $Z = (\texttt{pafi1}, \texttt{paco21})$ and $W = (\texttt{ph1}, \texttt{hema1})$ is only one of six possible assignments of the four candidate proxy variables. Across all six assignments, the proposed neural-network minimax DR estimators remain negative and range only from -0.0630 to -0.0429. In contrast, the simple linear closed-form estimator is positive in three of the six allocations (ranging from 0.0299 to 1.0435), and the linear moment-based estimator is positive in two; see Appendix H.2.1.

Second, we keep the proxy allocation fixed and systematically omit blocks of observed baseline covariates along two pre-specified rankings, yielding 28 increasingly severe omission scenarios. DR and DR(sta) retain negative estimates with 95% confidence intervals below zero in every constructed scenario. The linear estimators are much less stable: along the treatment–outcome predictive omission path, removing only the top four covariate blocks changes the linear moment-based estimate from -0.0828 in the full-X analysis to

	DR	DR(sta)	Linear(closed-form)	Linear(moment-based)
Treatment effects (SEs)	-0.0528 (0.01194)	-0.0535 (0.01064)	-0.0860 (0.02408)	-0.0828 (0.01977)
95% CIs	[-0.0762, -0.0294]	[-0.0743, -0.0326]	[-0.1332, -0.0388]	[-0.1216, -0.0441]

Table 8: Treatment effect estimates (standard errors) and 95% confidence intervals of the average treatment effect.

0.0828, and the linear closed-form estimate from -0.0860 to 0.3822; both corresponding confidence intervals include zero. More elaborate linear basis specifications can also produce very wide intervals under more severe omissions. Details are provided in Appendix H.2.2. Together, these exercises provide empirical evidence that the flexible minimax estimators are more stable under the specified proxy-allocation and covariate-omission perturbations. We emphasize that these analyses are empirical stress tests rather than formal tests of the proxy identification assumptions. Withholding observed covariates can change the effective unmeasured-confounding structure and need not preserve the proxy assumptions. Therefore, the observed stability does not establish proxy validity or robustness to arbitrary unmeasured confounding. Rather, it provides practically useful evidence that, under the specified proxy-allocation and covariate-omission perturbations, the flexible minimax estimators are less sensitive than the linear alternatives.

10. Conclusions and Future Work

In this paper, we tackled a central challenge in doing causal inference using negative controls: estimating the bridge functions. We developed an alternative identification strategy that does not rely on the completeness conditions widely assumed in the previous literature. We proposed new minimax estimators for the bridge functions that were amenable to general function approximation. We studied the behavior of these bridge function estimators and the resulting GACE estimators under a range of different assumptions.

Our work can be extended to tackle complex estimation in other settings. For example, both Tennenholtz et al. (2020); Lee and Bareinboim (2021) study complex settings where a causal estimand is identified using a proxy or negative-control approach – the first in a partially observable reinforcement-learning setting, and the second in a more general directed-acyclic-graph setting – but both focus on the setting of *discrete* data distributions for simplicity. By leveraging our work, which allows flexible hypothesis classes, we may be able to tackle these more complicated settings on continuous spaces.

	Half Width		Half Depth	
	DR	DR(sta)	DR	DR(sta)
Treatment effects (SEs)	-0.0595 (0.00982)	-0.0524 (0.01052)	-0.0436 (0.01009)	-0.0605 (0.01017)
95% CIs	[-0.0787, -0.0403]	[-0.0730, -0.0318]	[-0.0634, -0.0238]	[-0.0804, -0.0406]

Table 9: Treatment effect estimates (standard errors) and 95% confidence intervals for our proposed DR minimax estimators based on neural network bridge function estimation, across two different specifications of the neural networks. The "Half Width" setting reduces the number of hidden neurons in the neural network bridge function estimators by a half, while the "Half Depth" setting reduces the depth of the neural networks by a half.

	Linear (moment-based)	Linear (Interaction)	Linear (Quadratic)	Linear (Both)
Treatment effects (SEs)	-0.0828 (0.01977)	-0.0372 (0.07466)	-0.0859 (0.02427)	+0.1053 (0.97969)
95% CIs	[-0.1216, -0.0441]	[-0.1835, +0.1091]	[-0.1334, -0.0383]	[-1.8149, +2.0255]

Table 10: Treatment effect estimates (standard errors) and 95% confidence intervals of the ATE estimator with linear bridge functions, across three different specifications of the linear (moment-based) bridge functions. The "Interaction" specification additionally incorporates product interaction terms of variables $(WX_{\text{sub}}, ZX_{\text{sub}})$ in the linear bridge function estimation (where X_{sub} is a subset of covariates X that consist of the two mostly predictive covariates for the outcome Y), "Quadratic" specification incorporates quadratic terms (W^2, Z^2), while "Both" incorporates both interaction and quadratic terms.

Acknowledgments

The authors thank the Editor and anonymous reviewers for constructive feedback. Xiaojie Mao acknowledges support from the National Natural Science Foundation of China (Grant No. 72322001) and Shenzhen Science and Technology Program (Grant No. AI2026036). Zhiheng Zhang is supported by "the Fundamental Research Funds for the Central Universities" (Grant No. 2025110602) of Shanghai University of Finance and Economics and the Shanghai Engineering Research Center of Finance Intelligence (Grant No. 19DZ2254600).

Appendix A. Some Weaker Conditions That Fail to Ensure Identification

In this section, we prove three statements: (1) the existence of action bridge functions alone ($\mathbb{Q}_0 \neq \emptyset$) is not enough to identify GACE; (2) the existence of outcome bridge functions alone ($\mathbb{H}_0 \neq \emptyset$) is not enough to identify GACE; (3) the existence of two observed data bridge functions ($\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$) is not enough for identification.

A.1 The Existence of Action Bridge Functions Alone Is Not Enough

Suppose A, Z, U, Y are binary variables and X, W are empty (or constant) variables. We want to construct two different instances that satisfy the following conditions:

- The joint distributions of (Z, A, Y) are the same.
- Both instances satisfy $Z \perp Y \mid A, U$.
- The 2×2 matrices corresponding to $P(\mathbf{Z} \mid \mathbf{U}, A = 0)$ and $P(\mathbf{Z} \mid \mathbf{U}, A = 1)$ are both invertible, which ensures the existence of action bridge functions.
- The values of GACE corresponding to the contrast function $\pi(a \mid x) = I(a = 0)$, namely the counterfactual mean parameter $\mathbb{E}[Y(0)]$, are different in the two instances. This target parameter can be written as

$$\mathbb{E}[Y(0)] = \mathbb{E}[\mathbb{E}[Y \mid A = 0, U]].$$

Table 11 describes two different instances satisfying the above conditions where the parameter $\mathbb{E}[Y(0)]$ is equal to 0.5532 in one and to 0.543896 in the other.

A.2 The Existence of Outcome Bridge Functions Alone Is Not Enough

Suppose A, W, U, Y are binary variables and X, Z are empty (or constant) variables. We want to construct two different instances that satisfy the following conditions:

- The joint distributions of (W, A, Y) are the same.
- Both instances satisfy $W \perp A \mid U$.
- The 2×2 matrices corresponding to $P(\mathbf{W} \mid \mathbf{U}, A = 0)$ and $P(\mathbf{W} \mid \mathbf{U}, A = 1)$ are both invertible, which ensures the existence of outcome bridge functions.
- The GACE parameter $\mathbb{E}[Y(0)]$ takes different values in two instances.

Table 12 describes two different instances satisfying the above conditions where the parameter $\mathbb{E}[Y(0)]$ is equal to $1/2$ in one and to $72/175 \approx 0.411429$ in the other.

For the two instances in Table 12, direct summation over U gives the same distribution of (W, A, Y). Moreover, for $a = 0, 1$,

$$P_1(\mathbf{W} \mid \mathbf{U}, A = a) = \begin{pmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{pmatrix}, \qquad P_2(\mathbf{W} \mid \mathbf{U}, A = a) = \begin{pmatrix} 0.86 & 0.26 \\ 0.14 & 0.74 \end{pmatrix}.$$

Both matrices have determinant 0.6, and neither depends on a, so $W \perp A \mid U$ and outcome bridge functions exist in both instances.

U	A	Z	Y	P (instance 1)	P (instance 2)
0	0	0	0	0.075	0.025
0	0	0	1	0.05625	0.04375
0	0	1	0	0.025	0.075
0	0	1	1	0.01875	0.13125
0	1	0	0	0.0375	0.09
0	1	0	1	0.01875	0.12
0	1	1	0	0.1125	0.03
0	1	1	1	0.05625	0.04
1	0	0	0	0.025	0.075
1	0	0	1	0.04375	0.05625
1	0	1	0	0.075	0.025
1	0	1	1	0.13125	0.01875
1	1	0	0	0.1125	0.06
1	1	0	1	0.13125	0.03
1	1	1	0	0.0375	0.12
1	1	1	1	0.04375	0.06

Table 11: Two instances where the action bridge functions exist and the distributions of the observables (Z, A, Y) are identical but the parameter $\mathbb{E}[Y(0)]$ has different values.

U	Y	A	W	P (instance 1)	P (instance 2)
0	0	0	0	0.256	0.241
0	0	0	1	0.064	0.029
0	0	1	0	0.056	0.023
0	0	1	1	0.014	0.003
0	1	0	0	0.064	0.060
0	1	0	1	0.016	0.020
0	1	1	0	0.024	0.020
0	1	1	1	0.006	0.004
1	0	0	0	0.004	0.019
1	0	0	1	0.016	0.051
1	0	1	0	0.024	0.057
1	0	1	1	0.096	0.107
1	1	0	0	0.016	0.020
1	1	0	1	0.064	0.060
1	1	1	0	0.056	0.060
1	1	1	1	0.224	0.226

Table 12: Two instances where the outcome bridge functions exist and the distributions of the observables (W, A, Y) are identical but the parameter $\mathbb{E}[Y(0)]$ has different values.

A.3 The Existence of Observed Data Bridge Functions Is Not Enough

In Theorem 4, we prove that the target GACE parameter is identifiable if either $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ or $\mathbb{Q}_0 \neq \emptyset$ and $\mathbb{H}_0^{\text{obs}} \neq \emptyset$. In particular, the identification conditions are satisfied if $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0 \neq \emptyset$. However, one may wonder whether the weaker condition $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ would suffice for the identification. In this part, we show this is not enough for identification.

Specifically, we ignore the covariates for simplicity and construct two instances of joint distributions of (Z, W, U, A, Y), where all variables are zero-one binary random variables. These two instances are observationally equivalent, in the sense that their distributions of (Z, W, A, Y) are identical. Moreover, they both satisfy the conditions $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ of interest, but do not satisfy our proposed conditions (i.e., they have $\mathbb{H}_0 = \emptyset$ and $\mathbb{Q}_0 = \emptyset$). We will show that their corresponding average treatment effect parameters are different, so the causal identification fails.

In both instances, we impose the restrictions $(Z, W) \perp (A, Y, U)$, and for simplicity we let $P(Z = 1) = P(Z = 0) = P(W = 1) = P(W = 0) = 0.5$. Then we only need to specify the distributions of (A, Y, U). Note that as long as our specifications satisfy $A \not\perp U$ and $Y \not\perp U \mid A$, we have $\mathbb{H}_0 = \emptyset$ and $\mathbb{Q}_0 = \emptyset$ but $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ always hold.

- $\mathbb{H}_0 = \emptyset$ and $\mathbb{Q}_0 = \emptyset$. Indeed, if $\mathbb{H}_0 \neq \emptyset$, then there exists $h(A, W)$ such that $\mathbb{E}[h(A, W) \mid A, U] = \mathbb{E}[Y \mid A, U]$. Because $W \perp (A, U)$, this requires

$$\sum_{w \in \{0,1\}} h(a, w) P(W = w) = P(Y = 1 \mid A = a, U = u), \quad u = 0, 1.$$

  Notably, there does not exist such a h function, because the right hand side of the equation above varies with $u \in \{0, 1\}$ as long as $Y \not\perp U \mid A$ but the left hand side does not. Similarly, for $\mathbb{Q}_0 \neq \emptyset$, we need $q(A, Z)$ such that $\mathbb{E}[q(A, Z) \mid A, U] = 1/P(A \mid U)$. Because $Z \perp (A, U)$, this requires

$$\sum_{z \in \{0,1\}} q(a, z) P(Z = z) = 1/P(A = a \mid U = u), \quad u = 0, 1.$$

  Again, there does not exist such q function because $A \not\perp U$. Therefore, we have $\mathbb{H}_0 = \emptyset$ and $\mathbb{Q}_0 = \emptyset$ as long as $A \not\perp U, Y \not\perp U \mid A$.

- $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$. Specifically, we need to find $h(A, W)$ such that $\mathbb{E}[Y \mid A, Z] = \mathbb{E}[h(A, W) \mid A, Z]$. By $(W, Z) \perp (A, Y, U)$, this is equivalent to

$$\mathbb{E}[Y \mid A] = \mathbb{E}[\mathbb{E}[h(A, W) \mid A, U] \mid A, Z] = \mathbb{E}[\mathbb{E}[h(A, W) \mid A, U] \mid A] = \mathbb{E}[h(A, W) \mid A]$$

  Namely, we need to find a function $h(a, w)$ such that $\mathbb{E}[Y \mid A = a] = \mathbb{E}[h(a, W)]$ for $a \in \{0, 1\}$, which obviously always exists so that $\mathbb{H}_0^{\text{obs}} \neq \emptyset$. Similarly, for $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$, we only need to find $q(a, z)$ such that

$$\mathbb{E}[q(a, Z)] = \frac{1}{P(A = a)}, \ a \in \{0, 1\}.$$

  Again, such a q always exists as long as $P(A = 1) > 0, P(A = 0) > 0$.

Therefore, we only need to construct two instances of joint distributions of (A, Y, U) such that: (1) $A \not\perp U$ and $Y \not\perp U \mid A$; (2) their corresponding distributions of (A, Y) are identical; (3) their corresponding average treatment effects are different. Below we present two possible instances.

Instance One. The following give the distributions for the first instance:

$$\begin{aligned} P(U=1) &= 0.3725465, \\ P(A=1 \mid U=0) = 0.7306441, \quad P(A=1 \mid U=1) &= 0.4271815, \\ P(Y=1 \mid A=0, U=0) = 0.8064139, \quad P(Y=1 \mid A=0, U=1) &= 0.8523738, \\ P(Y=1 \mid A=1, U=0) = 0.1364452, \quad P(Y \mid A=1, U=1) &= 0.5224844. \end{aligned}$$

These imply that the observed variables (A, Y) have the following distributions:

$$P(A=1) = 0.6175902, \quad P(Y=1, A=1) = 0.1457034, \quad P(Y=1, A=0) = 0.3181885.$$

The average treatment effect under this instance is

$$\mathbb{E}\left[\mathbb{E}\left[Y \mid A=1, U\right] - \mathbb{E}\left[Y \mid A=0, U\right]\right] = -0.5432734.$$

Instance Two. The following give the distributions for the second instance:

$$\begin{aligned} P(U=1) &= 0.7354514, \\ P(A=1 \mid U=0) = 0.541148, \quad P(A=1 \mid U=1) &= 0.6450871, \\ P(Y=1 \mid A=0, U=0) = 0.8654667, \quad P(Y=1 \mid A=0, U=1) &= 0.8165264, \\ P(Y=1 \mid A=1, U=0) = 0.4652918, \quad P(Y \mid A=1, U=1) &= 0.16671. \end{aligned}$$

These imply that the observed variables (A, Y) have the same distributions:

$$P(A=1) = 0.6175902, \quad P(Y=1, A=1) = 0.1457034, \quad P(Y=1, A=0) = 0.3181885.$$

The average treatment effect under this instance is

$$\mathbb{E}\left[\mathbb{E}\left[Y \mid A=1, U\right] - \mathbb{E}\left[Y \mid A=0, U\right]\right] = -0.5837741.$$

Appendix B. Comparing Identification Strategies

One notable difference between our paper and previous literature is that our paper does not assume any completeness condition but previous papers do. In this section, we take the identification strategy in Cui et al. (2023) as an example to illustrate the role of completeness conditions in the previous literature and compare the previous identification strategies with our proposed identification strategy.

B.1 Conditions in Cui et al. (2023)

We first recall the sets of bridge functions given by the conditional moment equations in Equations (2) and (3):

$$\mathbb{H}_0 = \{h \in L_2(W, A, X) : \mathbb{E}[Y - h(W, A, X) \mid A, U, X] = 0\},$$

$$\mathbb{Q}_0 = \{q : \pi q \in L_2(Z, A, X), \mathbb{E}[\pi(A \mid X)\,(q(Z, A, X) - 1/f(A \mid U, X)) \mid A, U, X] = 0\}\,,$$

and the sets of bridge functions given by the observed data conditional moment equations in Equations (8) and (9):

$$\mathbb{H}_0^{\text{obs}} = \{h \in L_2(W, A, X) : \mathbb{E}[Y - h(W, A, X) \mid Z, A, X] = 0\}\,,$$
$$\mathbb{Q}_0^{\text{obs}} = \{q : \pi q \in L_2(Z, A, X), \mathbb{E}[\pi(A \mid X)\,(q(Z, A, X) - 1/f(A \mid W, X)) \mid W, A, X] = 0\}\,.$$

In Lemma 2, we already show that any bridge functions in $\mathbb{H}_0$ and $\mathbb{Q}_0$ can identify the causal parameter. However, we cannot directly learn functions in $\mathbb{H}_0$ and $\mathbb{Q}_0$ from the observed data, because they depend on the unmeasured confounders U. Instead, we can at most learn functions in $\mathbb{H}_0^{\text{obs}}$ and $\mathbb{Q}_0^{\text{obs}}$ from the observed data. Although in Lemma 3, we already show that $\mathbb{H}_0 \subseteq \mathbb{H}_0^{\text{obs}}$ and $\mathbb{Q}_0 \subseteq \mathbb{Q}_0^{\text{obs}}$, the converse may not be true (e.g., Appendix A.3 provides some examples). In other words, functions in $\mathbb{H}_0^{\text{obs}}$ (or $\mathbb{Q}_0^{\text{obs}}$) that we can possibly learn from observed data may not necessarily belong to $\mathbb{H}_0$ (or $\mathbb{Q}_0$). Thus without further conditions, we cannot use Lemma 2 for identification, since this lemma only applies to functions in $\mathbb{H}_0$ and $\mathbb{Q}_0$.

Cui et al. (2023) handles this problem by assuming the following completeness conditions.

Assumption 2. Consider the following conditions:

1. For any $g(A, U, X) \in L_2(A, U, X)$, $\mathbb{E}\left[g(A, U, X) \mid W, A, X\right] = 0$ only when $g(A, U, X) = 0$.

2. For any $g(A, U, X) \in L_2(A, U, X)$, $\mathbb{E}\left[g(A, U, X) \mid Z, A, X\right] = 0$ only when $g(A, U, X) = 0$.

With these completeness conditions, we can show that $\mathbb{Q}_0^{\text{obs}} = \mathbb{Q}_0$ and $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$, namely, any bridge functions that we can learn from the observed data (*i.e.*, functions in $\mathbb{Q}_0^{\text{obs}}$ and $\mathbb{H}_0^{\text{obs}}$) are indeed bridge functions defined by unmeasured confounders (*i.e.*, functions in $\mathbb{Q}_0$ and $\mathbb{H}_0$). These conclusions have appeared as intermediate steps in the identification theorems of Miao et al. (2018a); Cui et al. (2023).

Lemma 24. *Suppose Assumption 1 holds.*

1. *If Assumption 2 condition 1 further holds, then* $\mathbb{Q}_0^{obs} = \mathbb{Q}_0$.

2. *If Assumption 2 condition 2 further holds, then* $\mathbb{H}_0^{obs} = \mathbb{H}_0$.

Therefore, by assuming completeness conditions in Assumption 2, previous literature can use Lemma 2 to identify the causal parameter by any $q_0 \in \mathbb{Q}_0^{\text{obs}} = \mathbb{Q}_0$ and/or $h_0 \in \mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$.

Remark 1. Although Shi et al. (2020); Miao et al. (2018b) do not exactly assume the completeness conditions in Assumption 2, their identification strategy is the same as those when assuming Assumption 2: namely, they impose conditions that ensure $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$, so that they can still use Lemma 2 to achieve identification. Miao et al. (2018b) assumes that $\mathbb{H}_0 \neq \emptyset$ and that $\mathbb{H}_0^{\text{obs}}$ is a singleton. Since $\mathbb{H}_0 \subseteq \mathbb{H}_0^{\text{obs}}$ (Lemma 2), we must have $\mathbb{H}_0 = \mathbb{H}_0^{\text{obs}}$ is equal to the singleton. Shi et al. (2020) studies discrete negative controls and unmeasured confounders (see Example 3), and focuses on the setting where these variables have the

same number of values, *i.e.*, $|\mathcal{W}| = |\mathcal{Z}| = |\mathcal{U}|$. In this case, they assume that the matrix $P(\mathbf{W} \mid \mathbf{Z}, a, x)$ is invertible for any $a \in \mathcal{A}, x \in \mathcal{X}$ (see Assumption 4 condition 1 below). It is easy to show that this condition implies that $P(\mathbf{U} \mid \mathbf{Z}, a, x)$ is also invertible, namely, Assumption 2 condition 2 holds. Thus, Shi et al. (2020) implicitly requires $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$ as well in their main theory[8].

As a concrete example, we tailor the identification results in Cui et al. (2023) to the GACE parameter J in the following proposition.

Proposition 25 (Cui et al. (2023))**.** *Suppose that Assumption 1 holds.*

1. *If $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ (or, if $\mathbb{H}_0 \neq \emptyset$) and Assumption 2 condition 2 holds, then $J = \mathbb{E}\tilde{\phi}_{\text{REG}}(O; h_0)$ for any $h_0 \in \mathbb{H}_0^{\text{obs}}$.*

2. *If $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ (or, if $\mathbb{Q}_0 \neq \emptyset$) and Assumption 2 condition 1 holds, then $J = \mathbb{E}\tilde{\phi}_{\text{IPW}}(O; q_0)$ for any $q_0 \in \mathbb{Q}_0^{\text{obs}}$.*

3. *If conditions in either statement 1 or statement 2 hold, then $J = \mathbb{E}\tilde{\phi}_{\text{DR}}(O; h_0, q_0)$ for any $h_0 \in \mathbb{H}_0^{\text{obs}}$ and any $q_0 \in \mathbb{Q}_0^{\text{obs}}$.*

The proof for Proposition 25 goes as follows: given the conditions in statement 1, we must have $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$ according to Lemma 24, so any $h_0 \in \mathbb{H}_0^{\text{obs}}$ must satisfy $h_0 \in \mathbb{H}_0$ and thus can be used to identify J according to Lemma 2. Statement 2 can be proved analogously by noting $\mathbb{Q}_0^{\text{obs}} = \mathbb{Q}_0$ given the conditions therein. The completeness conditions in Proposition 25 are crucial since they ensure $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$ and $\mathbb{Q}_0^{\text{obs}} = \mathbb{Q}_0$. Note that as Cui et al. (2023) remarked, Proposition 25 can also start with $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0 \neq \emptyset$. These two conditions are equivalent to $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$, respectively, under the completeness conditions in Assumption 2.

B.2 Comparing Our Conditions and Conditions in Cui et al. (2023)

In the following proposition, we consider the discrete setting in Example 3. We show that our identification strategies in Theorem 4 require *strictly* weaker conditions than the counterparts in Proposition 25. In particular, we show that conditions in Proposition 25 statement 1 and statement 2 imply conditions in Theorem 4 statement 1 and statement 2, respectively. However, the converse is not true: there exist instances where conditions in Theorem 4 statement 1 and statement 2 hold but conditions in Proposition 25 statement 1 and statement 2 are violated.

Proposition 26. *Consider the discrete setting in Example 3 or the linear model setting in Example 4. Suppose that Assumption 1 holds.*

8. Shi et al. (2020) also discuss the case with $|\mathcal{W}| > |\mathcal{U}|$ and $|\mathcal{Z}| > |\mathcal{U}|$ so that bridge functions are nonunique and the matrix $P(\mathbf{W} \mid \mathbf{Z}, a, x)$ is not invertible. They remark that we can address the n'o'nu'ni'queness of bridge functions by considering the pseudo-inverse of the non-invertible matrix $P(\mathbf{W} \mid \mathbf{Z}, a, x)$ or by coarsening the negative controls W and Z so that their numbers of values are identical to that of U. But these approaches are not fully developed, and they still primarily focus on the setting with $|\mathcal{W}| = |\mathcal{U}|, |\mathcal{Z}| = |\mathcal{U}|$ in their paper. In this main setting, their identification again requires $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0$ and $\mathbb{Q}_0^{\text{obs}} = \mathbb{Q}_0$.

1. *If* $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ *and the completeness condition in Assumption 2 condition 2 holds, then* $\mathbb{H}_0 \neq \emptyset$ *and* $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$*. However, there exist instances such that* $\mathbb{H}_0 \neq \emptyset$ *and* $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ *but Assumption 2 condition 2 does not hold and* $\mathbb{H}_0 \subsetneq \mathbb{H}_0^{\text{obs}}$*.*

2. *If* $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ *and the completeness condition in Assumption 2 condition 1 holds, then* $\mathbb{Q}_0 \neq \emptyset$ *and* $\mathbb{H}_0^{\text{obs}} \neq \emptyset$*. However, there exist instances such that* $\mathbb{Q}_0 \neq \emptyset$ *and* $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ *but Assumption 2 condition 1 does not hold and* $\mathbb{Q}_0 \subsetneq \mathbb{Q}_0^{\text{obs}}$*.*

Proof. We prove Proposition 26 statement 1 here. Statement 2 can be proved analogously. *Discrete Data Setting.* We first consider the discrete setting in Example 3 where W, Z, U are all discrete variables. For simplicity, we omit the covariates.

First, according to Lemma 24, under Assumption 2 condition 2, $\mathbb{H}_0 = \mathbb{H}_0^{\text{obs}}$ so $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ implies $\mathbb{H}_0 \neq \emptyset$. Moreover, in the discrete setting given in Example 3, the completeness condition in Assumption 2 condition 2 amounts to $P(\mathbf{U} \mid \mathbf{Z}, a)$ having full row rank for any a. By Bayes rule, this implies that $P(\mathbf{Z} \mid \mathbf{U}, a)$ has full column rank for any a. It then follows from Equation (5) that $\mathbb{Q}_0 \neq \emptyset$. Since $\mathbb{Q}_0 \subseteq \mathbb{Q}_0^{\text{obs}}$ according to Lemma 3, we must have $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ as well. This proves the first part of statement 1.

Now we consider the second part of statement 1. We construct instances such that $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ but the completeness condition in Assumption 2 condition 2 does not hold. The idea is as follows: according to Equation (5), $\mathbb{H}_0 \neq \emptyset, \mathbb{Q}_0 \neq \emptyset$ if there exist solutions $h_0(\mathbf{W}, a)$ and $q_0(\mathbf{Z}, a)$ to the linear equation systems below for any $a \in \{0, 1\}$:

$$h_0^\top(\mathbf{W}, a) P(\mathbf{W} \mid \mathbf{U}, a) = \mathbb{E}[Y \mid \mathbf{U}, a],\ q_0^\top(\mathbf{Z}, a) P(\mathbf{Z} \mid \mathbf{U}, a) F(a \mid \mathbf{U}) = \mathbf{e}^\top.$$

Thus once we specify $P(\mathbf{W} \mid \mathbf{U}, a), \mathbb{E}[Y \mid \mathbf{U}, a], P(\mathbf{Z} \mid \mathbf{U}, a), F(a \mid \mathbf{U})$ in such a way that solutions to the linear equations above exist, we have $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ (because $\mathbb{Q}_0 \subseteq \mathbb{Q}_0^{\text{obs}}$ according to Lemma 3). At the same time, we can specify them in such a way that $P(\mathbf{Z} \mid \mathbf{U}, a)$ has deficient column rank, which in turn implies that $P(\mathbf{U} \mid \mathbf{Z}, a)$ has deficient row rank so the completeness condition given in Assumption 2 condition 2 does not hold. This is not difficult as long as $F(a \mid \mathbf{U})$ is chosen appropriately such that $\mathbf{e}^\top F^{-1}(a \mid \mathbf{U})$ belongs to the row space of $P(\mathbf{Z} \mid \mathbf{U}, a)$ and thus $q_0(\mathbf{Z}, a)$ exists. In other words, we construct an instance where the proxy Z does not capture all variations of U so the row space of $P(\mathbf{Z} \mid \mathbf{U}, a)$ is deficient. Nevertheless, Z still captures the relevant information in U, such that the row space of $P(\mathbf{Z} \mid \mathbf{U}, a)$ can still cover the vector $\mathbf{e}^\top F^{-1}(a \mid \mathbf{U})$ that characterizes the unobserved confounding in the action A.

Let us consider a simplified setting where $Y \in \{0, 1\}$, $U, W, Z \in \{0, 1, 2\}$, $A \in \{0, 1\}$. We provide a concrete instance for $a = 0$ that violates Assumption 2 condition 2 but satisfies $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$. It is easy to construct an instance for $a = 1$ analogously so we omit it here. We specify the following: for $a = 0$,

$$P(\mathbf{U} \mid \mathbf{Z}, a) = \begin{bmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{bmatrix}, \quad P(\mathbf{W} \mid \mathbf{U}, a) = \begin{bmatrix} \frac{1}{3} & \frac{1}{6} & \frac{1}{4} \\ \frac{1}{3} & \frac{1}{2} & \frac{5}{12} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix},$$

$$P(\mathbf{Z} \mid a) = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{3} \\ \frac{1}{6} \end{bmatrix}, \quad P(Y = 1 \mid \mathbf{U}, a) = \begin{bmatrix} \frac{1}{5} \\ \frac{2}{5} \\ \frac{3}{10} \end{bmatrix}, \quad F(a \mid \mathbf{U}) = \begin{bmatrix} \frac{3}{8} & 0 & 0 \\ 0 & \frac{3}{8} & 0 \\ 0 & 0 & \frac{3}{8} \end{bmatrix}.$$

In this instance, the rank of $P(\mathbf{U} \mid \mathbf{Z}, a)$ is 1 so Assumption 2 condition 2 does not hold. However, we can easily check that

$$q_0(\mathbf{Z}, a) = \begin{bmatrix} 2 \\ 3 \\ 4 \end{bmatrix} \in \mathbb{Q}_0 \subseteq \mathbb{Q}_0^{\text{obs}}, \quad h_0(\mathbf{W}, a) = \begin{bmatrix} -\frac{3}{10} \\ \frac{9}{10} \\ 0 \end{bmatrix} \in \mathbb{H}_0.$$

Thus $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ but Assumption 2 condition 2 does not hold. In this instance, actually $\mathbb{H}_0 \subsetneq \mathbb{H}_0^{\text{obs}}$. For example:

$$h_0(\mathbf{W}, a) = \begin{bmatrix} \frac{6}{5} \\ 0 \\ 0 \end{bmatrix} \in \mathbb{H}_0^{\text{obs}} \setminus \mathbb{H}_0.$$

In contrast, if the condition in Proposition 25 statement 1 were true, we must have $\mathbb{H}_0 = \mathbb{H}_0^{\text{obs}}$. This again confirms that conditions in our Theorem 4 statement 1 are strictly weaker than the counterparts in Proposition 25 statement 1.

Linear Model Setting. We now consider the linear model setting in Example 4. In that example, we assume the following data generating process:

$$\begin{aligned} Y &= \alpha_Y^\top U + \beta_Y^\top X + \gamma_Y^\top A + \omega_Y^\top W + \epsilon_Y, \\ Z &= \alpha_Z U + \beta_Z X + \gamma_Z A + \epsilon_Z, \\ W &= \alpha_W U + \beta_W X + \epsilon_W, \\ A &\sim \text{Unif}\left(\underline{\alpha}_A^\top U + \underline{\beta}_A^\top X, \overline{\alpha}_A^\top U + \overline{\beta}_A^\top X\right), \end{aligned}$$

We first prove that if the completeness condition in Assumption 2 condition 2 holds, then the matrix α_Z must have full column rank. We argue this by contraposition.

Let $d_U = \dim(U)$. Suppose $\text{rank}(\alpha_Z) = r < d_U$. Then there exist matrices $C \in \mathbb{R}^{r \times d_U}$ and $D \in \mathbb{R}^{(d_U - r) \times d_U}$ such that the map $u \mapsto (Cu, Du)$ is bijective, and the row space of C coincides with that of α_Z. Define $T = CU, S = DU$. Then $U \leftrightarrow (T, S)$ bijectively, and there exists a matrix M such that $\alpha_Z U = MT$. Hence

$$Z = MT + \beta_Z X + \gamma_Z A + \epsilon_Z, \text{ and } Z \perp S \mid (A, T, X).$$

Define a nonzero $V = v^\top S \in L^2$ for some $v \in \mathbb{R}^{d_U - r}$ such that $\text{Var}(V \mid A, T, X) > 0$ on a set of positive probability. Further define $g(A, U, X) = V - \mathbb{E}[V \mid A, T, X]$. Then $g(A, U, X) \neq 0$ with positive probability, and $\mathbb{E}[g(A, U, X) \mid A, T, X] = 0$. Using $Z \perp S \mid (A, T, X)$, we obtain

$$\mathbb{E}[g(A, U, X) \mid A, Z, X, T] = \mathbb{E}[g(A, S, T, X) \mid A, Z, X, T] = \mathbb{E}[g(A, U, X) \mid A, T, X] = 0,$$

and hence, even though $g(A, U, X)$ is nonzero with positive probability, we still have

$$\mathbb{E}[g(A, U, X) \mid A, Z, X] = \mathbb{E}\left[\mathbb{E}[g(A, U, X) \mid A, Z, X, T] \mid A, Z, X\right] = 0.$$

This contradicts completeness. Therefore $\text{rank}(\alpha_Z) = d_U$. This means that completeness condition implies the full column rank of matrix α_Z.

Then it follows from the completeness condition that there exists solution θ_Z to the linear equation system $\theta_Z^\top \alpha_Z = \overline{\alpha}_A^\top - \underline{\alpha}_A^\top$, which in turn implies the existence of bridge function $q_0 \in \mathbb{Q}_0$ according to the discussions in Example 4: $q_0(Z, A, X) = \theta_Z^\top Z + \left(\overline{\beta}_A - \underline{\beta}_A - \theta_Z^\top \beta_Z\right) X - \theta_Z^\top \gamma_Z A \in \mathbb{Q}_0$. This proves the first part of statement 1.

To prove the second part of statement 1, we only need to construct instances such that $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ but the completeness condition in Assumption 2 condition 2 does not hold. To this end, we construct $\alpha_W, \alpha_Z, \alpha_Y, \overline{\alpha}_A, \underline{\alpha}_A$ such that α_W, α_Z both have deficient column rank but α_Y is in the column space of α_W and $\overline{\alpha}_A - \underline{\alpha}_A$ is in the column space of α_Z. Then there exist solutions θ_W, θ_Z to the equations $\theta_W^\top \alpha_W = \alpha_Y^\top$ and $\theta_Z^\top \alpha_Z = \overline{\alpha}_A^\top - \underline{\alpha}_A^\top$ respectively. According to the discussions in Example 4, we have

$$h_0(W, A, X) = (\theta_W + \omega_Y)^\top W + \left(\beta_Y - \theta_W^\top \beta_W\right) X + \gamma_Y A \in \mathbb{H}_0,$$
$$q_0(Z, A, X) = \theta_Z^\top Z + \left(\overline{\beta}_A - \underline{\beta}_A - \theta_Z^\top \beta_Z\right) X - \theta_Z^\top \gamma_Z A \in \mathbb{Q}_0.$$

So for any such instance, we have $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0 \neq \emptyset$. This again implies that $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$. Moreover, because α_Z has deficient column rank, the completeness condition in Assumption 2 condition 2 does not hold. In summary, as long as α_Z does not have full column rank but $\alpha_Y, \overline{\alpha}_A, \underline{\alpha}_A$ are specified such that α_Y is in the column space of α_W and $\overline{\alpha}_A - \underline{\alpha}_A$ is in the column space of α_Z, then $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ but the completeness condition in Assumption 2 condition 2 fails. Again, this is because the proxy variables Z do not capture variations of U along all directions so that α_Z has deficient column rank and the completeness condition fails. However, it captures the variation of U along the relevant direction, such that the column space of α_Z can still cover $\overline{\alpha}_A - \underline{\alpha}_A$ that characterizes the unobserved confounding in action A.

As a concrete example, consider the instance with the following parameter specifications:

$$\alpha_W = \begin{bmatrix} 1 & 0 \end{bmatrix}, \alpha_Y = \begin{bmatrix} 2 \\ 0 \end{bmatrix}$$

and

$$\alpha_Z = \begin{bmatrix} 0 & 1 \end{bmatrix}, \quad \overline{\alpha}_A = \begin{bmatrix} 1 \\ 3 \end{bmatrix}, \quad \underline{\alpha}_A = \begin{bmatrix} 1 \\ 2 \end{bmatrix}.$$

Note in this instance, the dimension of proxy variables $p_w = p_z = 1$ is smaller than the dimension of the unobserved confounders $p_u = 2$. Under this case our proposed identification conditions still hold but the previous identification conditions fail. ■

Remark 2. In Proposition 26, we focus on the discrete data setting in Example 3 and the linear model setting in Example 4 to compare the IPW identification assumptions and REG identification assumptions in Theorem 4 and Proposition 25, separately. In the discrete data setting and linear model setting, the completeness conditions in Assumption 2 are sufficient but not necessary for the existence of bridge functions. So our identification strategies in Theorem 4 that only require existence of bridge functions (and observed bridge functions) are strictly weaker than those based on completeness conditions. For more general settings, the existence of bridge functions involve equations defined by potentially infinite-dimensional and

ill-posed linear operators. Because of technical challenges with infinite-dimensional function spaces, the existence of bridge functions and the corresponding completeness condition are closely related but not directly comparable without further regularity conditions. In Appendix C.2, we show that the completeness condition can imply the existence of bridge functions if additional regularity conditions are assumed.

Remark 3. Often one may hope to assume that identification assumptions in both statements 1 and 2 of Theorem 4 (or assumptions in both statements 1 and 2 of Proposition 25) hold simultaneously, rather than separately. Then one can apply all of the three identification formulae and the corresponding estimators, or consider semiparametrically efficient estimation (see Appendix D). In this case, our combined identification assumptions are $\mathbb{H}_0 \neq \emptyset, \mathbb{Q}_0 \neq \emptyset$, while the combined identification assumptions in Proposition 25 are both $\mathbb{H}_0^{\text{obs}} \neq \emptyset, \mathbb{Q}_0^{\text{obs}} \neq \emptyset$ and the two completeness conditions in Assumption 2. The latter are sufficient conditions for the former, since Assumption 2 implies that $\mathbb{H}_0^{\text{obs}} = \mathbb{H}_0, \mathbb{Q}_0^{\text{obs}} = \mathbb{Q}_0$. This means that when considering combined identification assumptions, our identification assumptions are never stronger than existing ones.

Appendix C. Completeness Conditions

In Appendix B, we show that completeness conditions in Assumption 2 play an important role in the identification strategy in previous literature. In this section, we review some related completeness conditions, discuss their relations, and further describe how our assumptions in Section 2 differ from those in previous literature. For completeness conditions in other settings, such as nonparametric instrumental variable models, see review and discussions in Hu and Shiu (2011); D'Haultfoeuille (2011); Darolles et al. (2010); Newey and Powell (2003). Throughout this section, we always assume Assumption 1 so we suppress this assumption in all statements.

C.1 Different Versions of Completeness Conditions

In the following assumptions, we list two other completeness conditions that also involve the unobserved confounders U. Although these conditions are not directly assumed in previous literature, we will show in Lemma 28 that they are implied by some other conditions that appear in the previous literature.

Assumption 3. Consider the following conditions:

1. For any $g(W, A, X) \in L_2(W, A, X)$, $\mathbb{E}[g(W, A, X) \mid A, U, X] = 0$ only when $g(W, A, X) = 0$.
2. For any $g(Z, A, X) \in L_2(Z, A, X)$, $\mathbb{E}[g(Z, A, X) \mid A, U, X] = 0$ only when $g(Z, A, X) = 0$.

In the following lemma, we further show that completeness conditions in Assumption 3 can ensure the uniqueness of bridge functions.

Lemma 27. *If completeness condition (1) in Assumption 3 holds, then $\mathbb{H}_0$ is either empty or a singleton. If completeness condition (2) in Assumption 3 holds, then $\mathbb{Q}_0$ is either empty or a singleton.*

Next, we introduce two completeness conditions that involve only observed variables.

Assumption 4. Consider the following conditions:

1. For any $g(W,A,X) \in L_2(W,A,X)$, $\mathbb{E}\left[g(W,A,X) \mid Z,A,X\right] = 0$ only when $g(W,A,X)=0$.
2. For any $g(Z,A,X) \in L_2(Z,A,X)$, $\mathbb{E}\left[g(Z,A,X) \mid W,A,X\right] = 0$ only when $g(Z,A,X) = 0$.

In the following lemma, we further show the relationship among Assumptions 2 to 4.

Lemma 28. *Assume Assumption 1 holds.*

1. *If Assumptions 2 and 3 hold, then Assumption 4 holds.*
2. *If Assumption 4 holds, then Assumption 3 holds.*

Lemma 28 shows that Assumptions 2 and 3 are sufficient for Assumption 4, and Assumption 4 is sufficient for Assumption 3. Since Assumption 3 ensure unique bridge functions according to Lemma 27, assuming Assumption 4 implicitly requires the bridge functions to be unique. For example, Cui et al. (2023) shows that Assumption 4 can be used to justify the existence of observed bridge functions, *i.e.*, $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$. Our discussion shows that this assumption does not only have implications for the existence of bridge functions, but also has indirect implications for the uniqueness of bridge functions. Our paper avoids assuming completeness conditions so we do not risk implicitly imposing uniqueness of the bridge functions.

C.2 Completeness Conditions and Existence of Bridge Functions

In this section, we show the existence of bridge functions in Equations (2) and (3) under the completeness conditions in Assumption 2 and some additional regularity conditions, using the singular value decomposition approach in Kress (2014); Miao et al. (2018a); Carrasco et al. (2007).

C.2.1 Characterizing Linear Operators

Let $K_{W|a,x} : L_2(W \mid A=a, X=x) \to L_2(U \mid A=a, X=x)$ and $K_{Z|a,x} : L_2(Z \mid A=a, X=x) \to L_2(U \mid A=a, X=x)$ be the linear operators defined as follows:

$$[K_{W|a,x}h](a,u,x) = \mathbb{E}\left[h(W,a,x) \mid A=a, U=u, X=x\right] = \int K(w,a,u,x)h(w,a,x)f(w \mid a,x)\mathrm{d}\mu(w),$$

$$[K_{Z|a,x}q](a,u,x) = \mathbb{E}\left[q(Z,a,x) \mid A=a, U=u, X=x\right] = \int K'(z,a,u,x)q(z,a,x)f(z \mid a,x)\mathrm{d}\mu(z),$$

where $K(w,a,u,x)$ and $K'(z,a,u,x)$ are the corresponding kernel functions defined as follows:

$$K(w,a,u,x) = \frac{f(w,u \mid a,x)}{f(u \mid a,x)f(w \mid a,x)}, \quad K'(z,a,u,x) = \frac{f(z,u \mid a,x)}{f(u \mid a,x)f(z \mid a,x)}.$$

Their adjoint operators $K^*_{W|a,x} : L_2(U \mid A=a, X=x) \to L_2(W \mid A=a, X=x)$ and $K^*_{Z|a,x} : L_2(U \mid A=a, X=x) \to L_2(Z \mid A=a, X=x)$ are given as follows:

$$[K^*_{W|a,x}g](w,a,x) = \int K(w,a,u,x)g(u,a,x)f(u \mid a,x)\mathrm{d}\mu(u) = \mathbb{E}\left[g(U,a,x) \mid W=w, A=a, X=x\right],$$
$$[K^*_{Z|a,x}g](z,a,x) = \int K'(z,a,u,x)g(u,a,x)f(u \mid a,x)\mathrm{d}\mu(u) = \mathbb{E}\left[g(U,a,x) \mid Z=z, A=a, X=x\right].$$

The existence of bridge functions is equivalent to existence of solutions to the following equations of the first kind:

$$[K_{W|a,x}h](a,u,x) = k_0(a,u,x), \quad [K_{Z|a,x}q](a,u,x) = 1/f(a \mid u,x), \quad \text{a.e. } u,a,x \text{ w.r.t } \mathbb{P}.$$

Moreover, the completeness conditions in Assumption 2 is equivalent to the injectivity of the adjoint operators $K^*_{Z|a,x}$ and $K^*_{W|a,x}$.

Proposition 29. *When the completeness condition in Assumption 2 condition 1 holds, the adjoint operator $K^*_{W|a,x}$ is injective and its null space is $\mathcal{N}(K^*_{W|a,x}) = \{g : K^*_{W|a,x}g = 0\} = \{0\}$ for almost all a, x. When the completeness condition in Assumption 2 condition 2 holds, the adjoint operator $K^*_{Z|a,x}$ is injective and its null space is $\mathcal{N}(K^*_{Z|a,x}) = \{g : K^*_{Z|a,x}g = 0\} = \{0\}$ for almost all a, x.*

We further assume the following regularity conditions.

Assumption 5. Assume that for almost every a, x:

1. $\iint f(w \mid u,a,x)f(u \mid w,a,x)\mathrm{d}\mu(w)\mathrm{d}\mu(u) < \infty$.
2. $\iint f(z \mid u,a,x)f(u \mid z,a,x)\mathrm{d}\mu(z)\mathrm{d}\mu(u) < \infty$.

According to Example 2.3 in Carrasco et al. (2007, P5659), Assumption 5 ensures that both $K_{W|a,x}$ and $K_{Z|a,x}$ are compact operators. Then by Theorem 2.41 in Carrasco et al. (2007, P5660), both $K_{W|a,x}$ and $K_{Z|a,x}$ admit singular value decomposition: there exist

$$\left(\lambda^j_{W|a,x}, \varphi^j_{W|a,x}, \psi^j_{W|a,x}\right)_{j=1}^{+\infty}, \quad \left(\lambda^j_{Z|a,x}, \varphi^j_{Z|a,x}, \psi^j_{Z|a,x}\right)_{j=1}^{+\infty}$$

with orthonormal sequences

$$\left\{\varphi^j_{W|a,x} \in L_2(W \mid a,x)\right\}, \qquad \left\{\varphi^j_{Z|a,x} \in L_2(Z \mid a,x)\right\},$$
$$\left\{\psi^j_{W|a,x} \in L_2(U \mid a,x)\right\}, \qquad \left\{\psi^j_{Z|a,x} \in L_2(U \mid a,x)\right\}$$

such that

$$K_{W|a,x}\varphi^j_{W|a,x} = \lambda^j_{W|a,x}\psi^j_{W|a,x}, \quad K_{Z|a,x}\varphi^j_{Z|a,x} = \lambda^j_{Z|a,x}\psi^j_{Z|a,x}.$$

C.2.2 Existence of Bridge Functions

Following Miao et al. (2018a), we use Picard's Theorem (Kress, 2014, Theorem 15.18) to characterize the existence of solutions to equations of the first kind by the singular value decomposition of the associated operators.

Lemma 30 (Picard's Theorem)**.** *Let $K : H_1 \to H_2$ be a compact operator with singular system $(\lambda_j, \varphi_j, \psi_j)_{j=1}^{+\infty}$, and ϕ be a given function in H_2. Then the equation of the first kind $Kh = \phi$ has solutions if and only if*

1. *$\phi \in \mathcal{N}(K^*)^{\perp}$, where $\mathcal{N}(K^*) = \{h : K^*h = 0\}$ is the null space of the adjoint operator K^*.*
2. *Picard's condition $\sum_{n=1}^{+\infty} \lambda_n^{-2} |\langle \phi, \psi_n \rangle|^2 < +\infty$ is satisfied.*

In the following two lemmas, we apply Picard's theorem to show that the completeness conditions can imply the existence of bridge functions under regularity conditions. According to Proposition 29, the completeness conditions ensure that the adjoint operators $K^*_{W|a,x}$ and $K^*_{Z|a,x}$ have null spaces $\mathcal{N}(K^*_{W|a,x}) = \{0\}, \mathcal{N}(K^*_{Z|a,x}) = \{0\}$. Then the first condition in Lemma 30 holds automatically. Hence we mainly need to additionally impose Picard's condition in Lemma 30. Under Picard's condition and other regularity conditions, the completeness conditions imply the existence of bridge functions.

Lemma 31. *Assume the operator $K^*_{W|a,x}$ satisfies Assumption 5 condition 1 and Picard's condition, i.e., $\sum_{j=1}^{+\infty} \left(\lambda^j_{W|a,x}\right)^{-2} \left\{\int k_0(a,u,x)\psi^j_{W|a,x}(u,a,x)\mathrm{d}\mu(u)\right\}^2 < \infty$ for almost all a, x. If the completeness condition in Assumption 2 condition 1 holds, then there exists function $h_0 \in L_2(W|A = a, X = x)$ such that for almost all a, x*

$$\mathbb{E}\left[Y - h_0(W, a, x) \mid A = a, U, X = x\right] = 0.$$

Lemma 32. *Assume the operator $K^*_{Z|a,x}$ satisfies Assumption 5 condition 2 and Picard's condition, i.e., $\sum_{j=1}^{+\infty} \left(\lambda^j_{Z|a,x}\right)^{-2} \left\{\int \frac{\pi(a|x)}{f(a|u,x)}\psi^j_{Z|a,x}(u,a,x)\mathrm{d}\mu(u)\right\}^2 < \infty$. If the completeness condition in Assumption 2 condition 2 holds, then there exists function q_0 such that $\pi(a \mid x)q_0(Z,a,x) \in L_2(Z \mid A = a, X = x)$ for almost all a, x and*

$$\mathbb{E}\left[\pi\left(a \mid x\right) q_0(Z, a, x) \mid a, U, x\right] = \frac{\pi\left(a \mid x\right)}{f(a \mid U, x)}.$$

Appendix D. Semiparametric Efficiency Bound

In this section, we derive the semiparametric efficiency bound for $J = \mathbb{E}\left[\int Y(a)\pi(a|X)\mathrm{d}(a)\right]$ under the nonparametric model $\mathcal{M}_{np}$ that does not restrict the data distributions other than require[9] $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\text{obs}} \neq \emptyset$. Our efficiency analysis generalizes Theorem 3.1 in Cui et al. (2023) for average treatment effect with discrete treatments.

9. These correspond to the identification assumptions in Theorem 4 statement 1. Alternatively, we can also follow Cui et al. (2023) and require the model to satisfy conditions in Proposition 25 statement 1, namely $\mathbb{H}_0^{\text{obs}} \neq \emptyset$ and the completeness conditions in Assumption 2 condition 2.

Theorem 33. *Let $\mathbb{P}$ be a data generating distribution such that $\mathbb{H}_0^{\text{obs}} = \{h_0\}, \mathbb{Q}_0^{\text{obs}} = \{q_0\}$, and the corresponding conditional expectation operator P_z defined in Equation (26) is bijective. Under Assumption 1, the efficient influence function of J under the model $\mathcal{M}_{np}$ locally at the data generating distribution $\mathbb{P}$ is given as follows:*

$$\text{EIF}(J) = \pi(A \mid X) q_0(Z, A, X)[Y - h_0(W, A, X)] + \mathcal{T} h_0(W, X) - J.$$

The corresponding semiparametric efficiency bound of J is $V_{\text{eff}} = \mathbb{E}\left[\text{EIF}^2(J)\right]$.

In Theorem 23, we show that our GACE estimator proposed in Section 7.2 can attain the efficiency bound in Theorem 33 when the bridge functions are unique.

Appendix E. Additional Examples of Minimax Estimators

This appendix gives additional examples of bridge function estimators based on minimax formulations.

E.1 RKHS Minimax Estimators without Stabilizers in the Discrete Action Setting

Consider Example 1 with a binary action $A \in \{0, 1\}$. We aim to estimate the counterfactual mean parameter $\mathbb{E}[Y(a)]$ for $a \in \{0, 1\}$, which is a special example of the GACE parameter corresponding to $\pi = \mathbb{I}(A = a)$.

We first consider estimating the outcome bridge function $h_0(W, a, X)$ based on a bridge class $\mathbb{H}$ over $\mathcal{W} \times \mathcal{A} \times \mathcal{X}$. For the critic class $\mathbb{Q}' : \mathcal{Z} \times \mathcal{A} \times \mathcal{X} \mapsto \mathbb{R}$, we can first specify a kernel function $\bar{k}_z : (\mathcal{Z} \times \mathcal{X}) \times (\mathcal{Z} \times \mathcal{X}) \to \mathbb{R}$ and then use the RKHS induced by the following product kernel:

$$k_z((Z_i, A_i, X_i), (Z_j, A_j, X_j)) = \mathbb{I}[A_i = A_j] \bar{k}_z((Z_i, X_i), (Z_j, X_j)).$$

Then we can apply Equation (24) in Lemma 6 to construct the estimator $\hat{h}$ as follows:

$$\arg\min_{h \in \mathbb{H}} (\psi_n(h))^\top \bar{K}_{z,n} \psi_n(h),$$

where

$$\{\bar{K}_{z,n}\}_{i,j} = \mathbb{I}(A_i = a)\mathbb{I}(A_j = a)\bar{k}_z((Z_i, X_i), (Z_j, X_j)), \{\psi_n(h)\}_i = \mathbb{I}(A_i = a)(Y_i - h(W_i, a, X_i)).$$

Similarly, we can also consider estimating the action bridge function $q_0(Z, a, X)$ based on a bridge class $\mathbb{Q}$ over $\mathcal{Z} \times \mathcal{A} \times \mathcal{X}$. For the critic class $\mathbb{H}' : \mathcal{W} \times \mathcal{A} \times \mathcal{X} \mapsto \mathbb{R}$, we can first specify a kernel function $\bar{k}_w : (\mathcal{W}, \mathcal{X}) \times (\mathcal{W}, \mathcal{X}) \to \mathbb{R}$ and then use the RKHS induced by the following product kernel:

$$k_w((W_i, A_i, X_i), (W_j, A_j, X_j)) = \mathbb{I}[A_i = A_j] \bar{k}_w((W_i, X_i), (W_j, X_j)).$$

Again, we can apply Equation (25) in Lemma 6 to construct the estimator $\hat{q}$ with

$$\{K_{w1,n}\}_{i,j} = \{K_{w2,n}\}_{i,j} = \mathbb{I}(A_i = a)\text{I}(A_j = a)\bar{k}_w((W_i, X_i), (W_j, X_j)), \quad \{\phi_n(q)\}_i = \mathbb{I}(A_i = a) q(Z_i, a, X_i).$$

The resulting estimator can be equivalently written as

$$\arg\min_{q \in \mathbb{Q}} (\{\phi_n(q) - \mathbf{1}_n\})^\top K_{w1,n} \{\phi_n(q) - \mathbf{1}_n\}.$$

E.2 Minimax Estimators with Stabilizers

In Section 4.3, we discuss the explicit form of minimax estimators without stabilizers when specialized to some common function classes. Here, we do the same for minimax estimators with stabilizers given in Equations (16) and (17).

Linear Classes. We can again consider the linear critic classes $\mathbb{Q}', \mathbb{H}'$ in Equation (18). However, with the norm constraints in $\mathbb{Q}', \mathbb{H}'$, the inner maximization in Equations (16) and (17) may no longer admit closed-form solutions. To circumvent this issue, we relax the hard norm constraints by setting $c_1 = c_2 = \infty$ in Equation (18), and regularizing coefficient norms in the inner maximization objectives.

Lemma 34. *Fix $c_1 = c_2 = +\infty$ in the critic classes $\mathbb{Q}', \mathbb{H}'$ in Equations* (18) *and* (19)*. Consider the following estimators adapted from Equations* (16) *and* (17)*:*

$$\hat{h} = \arg\min_{h\in\mathbb{H}} \max_{q\in\mathbb{Q}'} 2\mathbb{E}_n[\{h(W,A,X) - Y\}q\left(Z,A,X\right)] - \lambda\mathbb{E}_n[q^2\left(Z,A,X\right)] - \gamma_1\|\alpha_1\|^2,$$
$$\hat{q} = \arg\min_{q\in\mathbb{Q}} \max_{h\in\mathbb{H}'} 2\mathbb{E}_n[q(Z,A,X)\pi(A \mid X)h(W,A,X) - \mathcal{T}h(W,X)] - \lambda\mathbb{E}_n[h^2\left(W,A,X\right)] - \gamma_2\|\alpha_2\|^2$$

Then we have

$$\hat{h} = \arg\min_{h\in\mathbb{H}} \mathbb{E}_n[\{Y-h\}\phi]^\top\{\gamma_1 I + \lambda\mathbb{E}_n[\phi\phi^\top]\}^{-1}\mathbb{E}_n[\{Y-h\}\phi], \tag{37}$$
$$\hat{q} = \arg\min_{q\in\mathbb{Q}} \mathbb{E}_n[q\pi\psi - \mathcal{T}\psi]^\top\{\gamma_2 I + \lambda\mathbb{E}_n[\psi\psi^\top]\}^{-1}\mathbb{E}_n[\{q\pi\psi - \mathcal{T}\psi\}]. \tag{38}$$

If further $\mathbb{Q}$ and $\mathbb{H}$ are linear classes as in Equation (21), by adding Tikhonov regularization on the coefficients of $\mathbb{H}, \mathbb{Q}$ as

$$\hat{h} = \arg\min_{h\in\mathbb{H}} \mathbb{E}_n[\{Y-h\}\phi]^\top\{\gamma_1 I + \lambda\mathbb{E}_n[\phi\phi^\top]\}^{-1}\mathbb{E}_n[\{Y-h\}\phi] + \gamma_1'\|\alpha_2\|^2,$$
$$\hat{q} = \arg\min_{q\in\mathbb{Q}} \mathbb{E}_n[q\pi\psi - \mathcal{T}\psi]^\top\{\gamma_2 I + \lambda\mathbb{E}_n[\psi\psi^\top]\}^{-1}\mathbb{E}_n[\{q\pi\psi - \mathcal{T}\psi\}] + \gamma_2'\|\alpha_1\|^2,$$

we can obtain the closed-form solution:

$$\hat{h} = \tilde{\psi}^\top\hat{\alpha}_2,\ \hat{\alpha}_2 = \{\mathbb{E}_n[\tilde{\psi}\phi^\top]\{\gamma_1 I + \lambda\mathbb{E}_n[\phi\phi^\top]\}^{-1}\mathbb{E}_n[\phi\tilde{\psi}^\top] + \gamma_1' I\}^{-1}\mathbb{E}_n[\tilde{\psi}\phi^\top]\{\gamma_1 I + \lambda\mathbb{E}_n[\phi\phi^\top]\}^{-1}\mathbb{E}_n[Y\phi],$$
$$\pi\hat{q} = \tilde{\phi}^\top\hat{\alpha}_1,\ \hat{\alpha}_1 = \{\mathbb{E}_n[\tilde{\phi}\psi^\top]\{\gamma_2 I + \lambda\mathbb{E}_n[\psi\psi^\top]\}^{-1}\mathbb{E}_n[\psi\tilde{\phi}^\top] + \gamma_2' I\}^{-1}\mathbb{E}_n[\tilde{\phi}\psi^\top]\{\gamma_2 I + \lambda\mathbb{E}_n[\psi\psi^\top]\}^{-1}\mathbb{E}_n[\mathcal{T}\psi].$$

Finally,

$$\hat{J}_{\mathrm{REG}} = \mathbb{E}_n[\mathcal{T}\psi]^\top\{\mathbb{E}_n[\tilde{\psi}\phi^\top]\{\gamma_1 I + \lambda\mathbb{E}_n[\phi\phi^\top]\}^{-1}\mathbb{E}_n[\phi\tilde{\psi}^\top] + \gamma_1' I\}^{-1}\mathbb{E}_n[\tilde{\psi}\phi^\top]\{\gamma_1 I + \lambda\mathbb{E}_n[\phi\phi^\top]\}^{-1}\mathbb{E}_n[Y\phi],$$
$$\hat{J}_{\mathrm{IPW}} = \mathbb{E}_n[Y\phi]^\top\{\mathbb{E}_n[\tilde{\phi}\psi^\top]\{\gamma_2 I + \lambda\mathbb{E}_n[\psi\psi^\top]\}^{-1}\mathbb{E}_n[\psi\tilde{\phi}^\top] + \gamma_2' I\}^{-1}\mathbb{E}_n[\tilde{\phi}\psi^\top]\{\gamma_2 I + \lambda\mathbb{E}_n[\psi\psi^\top]\}^{-1}\mathbb{E}_n[\mathcal{T}\psi].$$

RKHS. We can also consider the RKHS critic classes $\mathbb{Q}', \mathbb{H}'$ in Equation (23). But just like the linear classes above, the inner maximization in Equations (16) and (17) may no longer have closed-form solutions with the RKHS norm constraints in Equation (23) either. We can again relax these norm constraints by setting $c_1 = c_2 = \infty$ in Equation (23), and instead regularize the inner maximization objectives.

Lemma 35. *Let $\lambda \geq 0$ and $\gamma_1, \gamma_2 > 0$, and assume that $h \mapsto \mathbb{E}_n[(\mathcal{T}h)(W,X)]$ is a bounded linear functional on $\mathcal{L}_w$. Consider the following estimators adapted from Equations* (16) *and* (17)*:*

$$\hat{h} = \arg\min_{h\in\mathbb{H}} \max_{q\in\mathcal{L}_z} 2\mathbb{E}_n[\{h(W,A,X) - Y\}q\left(Z,A,X\right)] - \lambda\mathbb{E}_n[q^2\left(Z,A,X\right)] - \gamma_1\|q\|_{\mathcal{L}_z}^2, \tag{39}$$

$$\hat{q} = \arg\min_{q \in \mathbb{Q}} \max_{h \in \mathcal{L}_w} 2\mathbb{E}_n[q(Z,A,X)\pi(A \mid X)h(W,A,X) - \mathcal{T}h(W,X)] - \lambda \mathbb{E}_n[h^2(W,A,X)] - \gamma_2 \|h\|^2_{\mathcal{L}_w}. \tag{40}$$

Then we have

$$\hat{h} = \arg\min_{h \in \mathbb{H}} \psi_n(h)^\top K_{z,n}^{1/2} \{n\gamma_1 I + \lambda K_{z,n}\}^{-1} K_{z,n}^{1/2} \psi_n(h), \tag{41}$$

$$\begin{aligned}\hat{q} = \arg\min_{q \in \mathbb{Q}} \Big[& \phi_n(q)^\top K_{w1,n}^{1/2} (\lambda K_{w1,n} + n\gamma_2 I)^{-1} K_{w1,n}^{1/2} \phi_n(q) \\ & - 2\phi_n(q)^\top (\lambda K_{w1,n} + n\gamma_2 I)^{-1} K_{w2,n} \mathbf{1}_n \Big].\end{aligned} \tag{42}$$

We remark that the computational time for RKHS estimators with stabilizers may be larger than the counterparts without stabilizers. Indeed, for RKHS estimators with stabilizers, evaluating the objective functions in Equations (41) and (42) requires matrix inverse that generally takes $O(n^3)$ time but evaluating the counterparts in Equations (24) and (25) only takes $O(n^2)$ time.

If further $\mathbb{H} = \mathcal{L}_w$ and $\mathbb{Q} = \mathcal{L}_z$, we can add the penalties $\gamma_1' \|h\|^2_{\mathcal{L}_w}$ and $\gamma_2' \|q\|^2_{\mathcal{L}_z}$ to the objectives in Equations (41) and (42), respectively. Take $\gamma_1, \gamma_2, \gamma_1', \gamma_2' > 0$ and $\lambda \geq 0$, and write

$$\begin{aligned} D_\pi &= \operatorname{diag}(\pi(A_1 \mid X_1), \ldots, \pi(A_n \mid X_n)), \\ \bar{Y} &= (Y_1, \ldots, Y_n)^\top, \\ M_z &= K_{z,n}^{1/2} (n\gamma_1 I + \lambda K_{z,n})^{-1} K_{z,n}^{1/2}, \\ M_w &= K_{w1,n}^{1/2} (n\gamma_2 I + \lambda K_{w1,n})^{-1} K_{w1,n}^{1/2}, \\ b_w &= (n\gamma_2 I + \lambda K_{w1,n})^{-1} K_{w2,n} \mathbf{1}_n. \end{aligned}$$

The representer theorem gives

$$\begin{aligned} \hat{h}(\cdot) &= \sum_{i=1}^n \hat{\alpha}_i k_w((W_i, A_i, X_i), \cdot), \\ \hat{q}(\cdot) &= \sum_{i=1}^n \hat{\beta}_i k_z((Z_i, A_i, X_i), \cdot). \end{aligned}$$

The evaluation vectors are therefore $\psi_n(\hat{h}) = \bar{Y} - K_{w1,n}\hat{\alpha}$ and $\phi_n(\hat{q}) = D_\pi K_{z,n} \hat{\beta}$. Substituting these expressions into the regularized objectives and differentiating gives the normal equations

$$\begin{aligned} (K_{w1,n} M_z K_{w1,n} + \gamma_1' K_{w1,n}) \hat{\alpha} &= K_{w1,n} M_z \bar{Y}, \\ (K_{z,n} D_\pi M_w D_\pi K_{z,n} + \gamma_2' K_{z,n}) \hat{\beta} &= K_{z,n} D_\pi b_w. \end{aligned}$$

These systems are consistent even when the Gram matrices are singular; one solution is obtained using the Moore–Penrose inverse. Finally, with $c_w = n^{-1} K_{w2,n} \mathbf{1}_n$, the corresponding GACE estimators are

$$\hat{J}_{\mathrm{REG}} = c_w^\top \hat{\alpha}, \qquad \hat{J}_{\mathrm{IPW}} = n^{-1} \bar{Y}^\top D_\pi K_{z,n} \hat{\beta}.$$

Appendix F. Error Bounds for Linear Models under Single Realizability

In Corollary 8, we establish consistency of estimators when realizability $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset, \mathbb{H}' \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset$ or $\mathbb{H} \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset, \mathbb{Q}' \cap \pi\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ holds. These conditions essentially require realizability for both bridge classes and critic classes. In this section, we show that when using linear models in Section 4.3.1, under mild matrix invertibility conditions, we need realizability condition on only a single but not both function classes.

Corollary 36 (IPW and REG under single realizability)**.** *Let $\mathbb{Q}, \mathbb{Q}', \mathbb{H}, \mathbb{H}'$ be the linear classes given in Section 4.3.1.*

1. *If $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{E}[\tilde{\psi}(Z,A,X)\psi^\top(W,A,X)]$ has full column rank, then*

$$|\hat{J}_{\text{IPW}} - J| \leq \sup_{q\in\mathbb{Q}} |(\mathbb{E}_n - \mathbb{E})\left[q\pi Y\right]| + 2\sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |(\mathbb{E}_n - \mathbb{E})\left[-q\pi h + \mathcal{T}h\right]|.$$

   This statement also holds if instead $\mathbb{H} \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{E}[\tilde{\psi}(Z,A,X)\psi^\top(W,A,X)]$ is full column rank.

2. *If $\mathbb{H} \cap \mathbb{H}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{E}[\psi(Z,A,X)\tilde{\psi}^\top(W,A,X)]$ has full row rank, then*

$$|\hat{J}_{\text{REG}} - J| \leq \sup_{h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})\left[\mathcal{T}h\right]| + 2\sup_{q\in\mathbb{Q}',h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})\left[q(Y-h)\right]|.$$

   This statement also holds if instead $\mathbb{Q}' \cap \pi\mathbb{Q}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{E}[\psi(Z,A,X)\tilde{\psi}^\top(W,A,X)]$ has full column rank.

Corollary 36 implies that the IPW and REG estimators are consistent as long as a single realizability condition holds and certain matrices are full column or row rank. Compared to Corollary 8, we do not need realizability for both the bridge function class and critic function class.

Corollary 36 is immediately obtained from Theorem 7 by showing bias terms are 0. For instance, suppose $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset$ and $\mathbb{E}[\tilde{\psi}(Z,A,X)\psi^\top(W,A,X)]$ is full row rank. Then, the second term (27) is 0 since $\mathbb{Q} \cap \mathbb{Q}_0^{\text{obs}} \neq \emptyset$. The first term is also 0 since we can take a linear projection of h_0 in $\mathbb{H}'$ even if we cannot choose $h = h_0$.

Appendix G. Linear Sieves and Neural Networks

In Example 8, we consider nonparametric classes such as Hölder balls and Sobolev balls. However, it is difficult to optimize over these function classes since they are infinite-dimensional. Instead, we may use a sequence of more tractable function classes with increasing complexity to approximate those nonparametric classes.

G.1 Linear Sieves

We first consider using linear sieve classes for $\mathbb{H}$ and $\mathbb{Q}$ that consist of linear combinations of some basis functions such as splines, polynomials, wavelets, and so on (Chen, 2007).

To give an example, we consider a simple setting where $\mathcal{W} \times \mathcal{Z} \times \mathcal{A} \times \mathcal{X} = [0,1]^{d_W+d_Z+d_A+d_X}$ with $d_W = d_Z$. We consider generic basis functions $\{\psi_j(w,a,x)\}_{j=1}^{k_n}$ and $\{\phi_j(z,a,x)\}_{j=1}^{k_n}$ (which for simplicity have the same size k_n) where $\psi_j : [0,1]^d \mapsto \mathbb{R}, \phi_j :$

$[0,1]^d \mapsto \mathbb{R}$ for $d = d_W + d_A + d_X = d_Z + d_A + d_X$. These basis functions induce the following linear sieve classes

$$\mathcal{S}_{1,n} = \left\{ (w,a,x) \mapsto \sum_{j=1}^{k_n} \omega_j \psi_j(w,a,x) : \omega \in \mathbb{R}^{k_n} \right\}, \quad \mathcal{S}_{2,n} = \left\{ (z,a,x) \mapsto \sum_{j=1}^{k_n} \omega_j \phi_j(z,a,x) : \omega \in \mathbb{R}^{k_n} \right\}. \tag{43}$$

We assume these two classes can approximate Hölder balls over $[0,1]^d$ whose Hölder smoothness parameters and radius are assumed to be the same for simplicity. We denote the common smoothness level as α and denote the corresponding Hölder balls as $\Lambda^\alpha([0,1]^d)$. Following Chen (2007), we assume standard approximation errors as follows:

$$\forall h \in \Lambda^\alpha([0,1]^d), \exists h_n \in S_{1,n}, \text{ s.t. } \|h - h_n\| = O\left(k_n^{-\alpha/d}\right), \tag{44}$$

$$\forall q \in \Lambda^\alpha([0,1]^d), \exists q_n \in S_{2,n}, \text{ s.t. } \|q - q_n\| = O\left(k_n^{-\alpha/d}\right). \tag{45}$$

Estimators without Stabilizers. In the following Corollary 37, we first bound the errors of IPW and REG estimators based on sieve minimax bridge function estimators *without* stabilizers.

Corollary 37. *Assume sieve classes $\mathcal{S}_{1,n}, \mathcal{S}_{2,n}$ satisfy the approximation error conditions in Equations (44) and (45) with $k_n = O(n^{\frac{d}{d+2\alpha}})$ and $\|\mathcal{S}_{1,n}\|_\infty < \infty, \|\mathcal{S}_{2,n}\|_\infty < \infty$. Also suppose $\|\pi(A \mid X)\|_\infty < \infty, \|Y\|_\infty < \infty, \|\pi(A \mid X)/f(A \mid X, W)\|_2 < \infty$, and $\mathbb{H}_0^{\text{obs}} \cap \Lambda^\alpha([0,1]^d) \neq \emptyset$, $\pi\mathbb{Q}_0^{\text{obs}} \cap \Lambda^\alpha([0,1]^d) \neq \emptyset$. Let $\hat{J}_{\text{REG}}$ be the REG estimator based on the $\hat{h}$ given in Equation (12) with $\mathbb{H} = \mathcal{S}_{1,n}$ and $\mathbb{Q}' = \mathcal{S}_{2,n}$, and $\hat{J}_{\text{IPW}}$ be the IPW estimator based on $\hat{q}$ given in Equation (13) with $\mathbb{Q} = \mathcal{S}_{2,n}$ and $\mathbb{H}' = \mathcal{S}_{1,n}$. Then with probability $1-\delta$, we have*

$$\max\left\{ |\hat{J}_{\text{REG}} - J|, \ |\hat{J}_{\text{IPW}} - J| \right\} = O\left(n^{-\alpha/(2\alpha+d)} + \sqrt{\log(1/\delta)/n} \right).$$

Compared to Corollaries 11 and 12, Corollary 37 assumes realizability on $\Lambda^\alpha([0,1]^d)$ rather than on $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$. Therefore, we cannot prove Corollary 37 by directly applying Corollary 8, as the latter requires realizability assumptions for $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$. Instead, we need to use the more general Theorem 7, which do not require such realizability assumptions and characterize the bias due to violations of these assumptions. These bias terms can be upper-bounded by sieve approximation errors in Equations (44) and (45), and the variance terms can be bounded in terms of VC dimensions that depend on k_n. Choosing the sieve size $k_n = O(n^{\frac{d}{d+2\alpha}})$ to balance the bias upper bounds and the variance upper bounds in Theorem 7, we can obtain the convergence rate in Corollary 37.

Estimators with Stabilizers. Below we also bound the projected MSEs of sieve minimax bridge function estimators with stabilizers. These bounds can easily translate into error bounds of the corresponding GACE estimators according to Theorems 14 and 15. Alternatively, they can also translate into more refined error bounds for the DR estimator by leveraging the ill-posedness measures in Definition 4 and following Theorems 22 and 23.

Corollary 38. *Assume sieve classes $\mathcal{S}_{1,n}, \mathcal{S}_{2,n}$ satisfy the approximation error conditions in Equations (44) and (45) with $k_n = O(n^{\frac{d}{d+2\alpha}})$ and $\|\mathcal{S}_{1,n}\|_\infty < \infty, \|\mathcal{S}_{2,n}\|_\infty < \infty$. Also suppose $\|\pi(A \mid X)\|_\infty < \infty, \|Y\|_\infty < \infty, \|\pi(A \mid X)/f(A \mid X, W)\|_2 < \infty$.*

1. *Let $\hat{h}$ be the outcome bridge function estimator given in Equation* (16) *with $\mathbb{H} = \mathcal{S}_{1,n}, \mathbb{Q}' = \mathcal{S}_{2,n}$ and $\lambda = 1$. If $\mathbb{H}_0^{\text{obs}} \cap \Lambda^\alpha([0,1]^d) \neq \emptyset$ and there exists $h_0 \in \mathbb{H}_0^{\text{obs}} \cap \Lambda^\alpha([0,1]^d)$ such that $P_z(\Lambda^\alpha([0,1]^d) - h_0) \subset \Lambda^\alpha([0,1]^d)$, then with probability at least $1-\delta$,*

$$\|P_z(\hat{h} - h_0)\|_2 = \tilde{O}(n^{-\alpha/(2\alpha+d)} + \sqrt{\log(1/\delta)/n}).$$

2. *Let $\hat{q}$ be the action bridge function estimator given in Equation* (17) *with $\mathbb{Q} = \mathcal{S}_{2,n}, \mathbb{H}' = \mathcal{S}_{1,n}$ and $\lambda = 1$. If $\mathbb{Q}_0^{\text{obs}} \cap \Lambda^\alpha([0,1]^d) \neq \emptyset$ and there exists $q_0 \in \mathbb{Q}_0^{\text{obs}} \cap \Lambda^\alpha([0,1]^d)$ such that $\pi P_w(\Lambda^\alpha([0,1]^d) - q_0) \subset \Lambda^\alpha([0,1]^d)$, then with probability at least $1-\delta$,*

$$\|P_w(\hat{q} - q_0)\|_2 = \tilde{O}(n^{-\alpha/(2\alpha+d)} + \sqrt{\log(1/\delta)/n}).$$

Similarly, Corollary 38 assumes realizability on $\Lambda^\alpha([0,1]^d)$ rather than on $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$. So we cannot prove Corollary 38 by directly applying Corollaries 17 and 19, as the latter requires realizability assumptions for $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$. Instead, we need to use the more general Theorems 16 and 18 which do not require such realizability assumptions and characterize the bias due to violations of these assumptions. By appropriately balancing the bias and variance terms, we can obtain the projected MSE bounds above.

G.2 Neural Networks

Alternatively, we can consider neural networks in the minimax estimation of bridge functions. For example, we again consider $\mathcal{W} \times \mathcal{Z} \times \mathcal{A} \times \mathcal{X} = [0,1]^{d_W + d_Z + d_A + d_X}$ with $d_W = d_Z$ and denote $d = d_W + d_A + d_X = d_Z + d_A + d_X$. We consider $\mathcal{F}_{1,n}$ and $\mathcal{F}_{2,n}$ as classes of neural networks over $\mathcal{W} \times \mathcal{A} \times \mathcal{X}$ and $\mathcal{Z} \times \mathcal{A} \times \mathcal{X}$ respectively, both with L layers, Ω parameters and $1-$Lipschitz continuous activation function (*e.g.*, ReLU activation functions).

It is known from the universal approximation property of neural networks (Yarotsky, 2017) that these neural network classes can approximate Sobolev balls over $[0,1]^d$. For simplicity, we assume that the two limiting Sobolev balls have the same smoothness parameter α, which we denote as $\mathcal{S}^\alpha([0,1]^d)$.

Lemma 39 (Theorem 1 in Yarotsky (2017))**.** *Suppose $\mathcal{F}_{1,n}, \mathcal{F}_{2,n}$ have at most $L = O(\log \Omega)$ layers and Ω parameters. Then for any $h, q \in \mathcal{S}^\alpha([0,1]^d)$, there exists $h_n \in \mathcal{F}_{1,n}$ and $q_n \in \mathcal{F}_{2,n}$ such that*

$$\|h - h_n\|_\infty = \tilde{O}(\Omega^{-\alpha/d}), \quad \|q - q_n\|_\infty = \tilde{O}(\Omega^{-\alpha/d}).$$

Estimators without Stabilizers. We first regard neural networks as sieve estimators which approximate Sobolev balls. Thus similar to the analysis for Corollary 38, we can also bound the errors of IPW and REG estimators based on neural network minimax bridge function estimators *without* stabilizers.

Corollary 40. *Assume neural network classes $\mathcal{F}_{1,n}, \mathcal{F}_{2,n}$ both have layers $L = \Theta(\log(n))$ and $\Omega = \Theta(n^{d/(2\alpha+d)})$ weight parameters bounded in $[-B, B]$ for a positive constant B and $\|\mathcal{F}_{1,n}\|_\infty, \|\mathcal{F}_{2,n}\|_\infty < \infty$. Also suppose $\|\pi(A \mid X)\|_\infty < \infty, \|Y\|_\infty < \infty, \|\pi(A \mid X)/f(A \mid X, W)\|_2 < \infty$, and $\mathbb{H}_0^{\text{obs}} \cap \mathcal{S}^\alpha([0,1]^d) \neq \emptyset$, $\pi\mathbb{Q}_0^{\text{obs}} \cap \mathcal{S}^\alpha([0,1]^d) \neq \emptyset$. Let $\hat{J}_{\text{REG}}$ be the REG*

estimator based on the $\hat{h}$ given in Equation (12) *with $\mathbb{H} = \mathcal{F}_{1,n}$ and $\mathbb{Q}' = \mathcal{F}_{2,n}$, and $\hat{J}_{\mathrm{IPW}}$ be the IPW estimator based on $\hat{q}$ given in Equation* (13) *with $\mathbb{Q} = \mathcal{F}_{2,n}$ and $\mathbb{H}' = \mathcal{F}_{1,n}$. Then with probability $1-\delta$, we have*

$$\max\left\{|\hat{J}_{\mathrm{REG}} - J|,\ |\hat{J}_{\mathrm{IPW}} - J|\right\} = O\left(n^{-\alpha/(2\alpha+d)} + \sqrt{\log(1/\delta)/n}\right). \tag{46}$$

Note that in Corollary 40, we assume realizability on the Sobolev balls $\mathcal{S}^\alpha([0,1]^d)$ rather than on $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$. To prove Corollary 40, we bound the bias terms in Theorem 7 by the neural network approximation errors given in Yarotsky (2017). The variance terms therein can be upper bounded by calculating Rademacher complexities of some function classes based on neural networks (see Lemmas 42 and 48). By choosing the neural network size appropriately to balance these terms, we can arrive at the final convergence rate in Equation (46).

Note that the analysis in Corollary 40 requires the number of parameters in neural networks to be smaller than the sample size n. In practice, neural networks are sometimes over-parameterized to achieve better performance so that the number of parameters can indeed exceed the sample size. In the over-parameterized case, we may leverage the results in Golowich et al. (2018) to compute Rademacher complexities of neural networks and obtain similar conclusions.

Estimators with Stabilizers. Below we also bound the projected MSEs of neural network minimax bridge function estimators with stabilizers. These bounds again can be easily translated into error bounds of the corresponding GACE estimators according to Theorems 14, 15, 22 and 23.

Corollary 41. *Assume neural network classes $\mathcal{F}_{1,n}, \mathcal{F}_{2,n}$ both have layers $L = \Theta(\log(n))$ and $\Omega = \Theta(n^{d/(2\alpha+d)})$ weight parameters bounded in $[-B, B]$ for a positive constant B and $\|\mathcal{F}_{1,n}\|_\infty, \|\mathcal{F}_{2,n}\|_\infty < \infty$. Also suppose $\|\pi(A \mid X)\|_\infty < \infty, \|Y\|_\infty < \infty, \|\pi(A \mid X)/f(A \mid X, W)\|_2 < \infty$.*

1. *Let $\hat{h}$ be the outcome bridge function estimator given in Equation* (16) *with $\mathbb{H} = \mathcal{F}_{1,n}$ and $\mathbb{Q}' = \mathcal{F}_{2,n}$. If $\mathbb{H}_0^{\mathrm{obs}} \cap \mathcal{S}^\alpha([0,1]^d) \neq \emptyset$ and there exists $h_0 \in \mathbb{H}_0^{\mathrm{obs}} \cap \mathcal{S}^\alpha([0,1]^d)$ such that $P_z(\mathcal{S}^\alpha([0,1]^d) - h_0) \subset \mathcal{S}^\alpha([0,1]^d)$, then with probability at least $1-\delta$,*

$$\|P_z(\hat{h} - h_0)\|_2 = \tilde{O}(n^{-\alpha/(2\alpha+d)} + \sqrt{\log(1/\delta)/n}).$$

2. *Let $\hat{q}$ be the action bridge function estimator given in Equation* (17) *with $\mathbb{Q} = \mathcal{F}_{2,n}$ and $\mathbb{H}' = \mathcal{F}_{1,n}$. If $\mathbb{Q}_0^{\mathrm{obs}} \cap \mathcal{S}^\alpha([0,1]^d) \neq \emptyset$ and there exists $q_0 \in \mathbb{Q}_0^{\mathrm{obs}} \cap \mathcal{S}^\alpha([0,1]^d)$ such that $\pi P_w(\mathcal{S}^\alpha([0,1]^d) - q_0) \subset \mathcal{S}^\alpha([0,1]^d)$, then with probability at least $1-\delta$,*

$$\|P_w(\hat{q} - q_0)\|_2 = \tilde{O}(n^{-\alpha/(2\alpha+d)} + \sqrt{\log(1/\delta)/n}).$$

Again, we assume realizability on the Sobolev balls $\mathcal{S}^\alpha([0,1]^d)$ rather than on $\mathbb{H}, \mathbb{Q}, \mathbb{H}', \mathbb{Q}'$. So we use the more general Theorems 16 and 18 and choose the size of neural networks appropriately to get the bounds above.

Appendix H. Additional Experiment Details

This appendix provides implementation details for the supplementary numerical and real-data analyses. We first describe the data-generating process and estimation settings for the simulation studies. We then report two complementary robustness analyses for the RHC study: sensitivity to the allocation of candidate proxy variables to Z and W, and sensitivity to the omission of observed baseline covariates. These analyses probe distinct sources of specification uncertainty in the real-data application.

H.1 Simulation Study

Tailoring the data generating process in Cui et al. (2023) to multi-dimensional variables, we generate $U, X', Z', W' \in \mathbb{R}^d$ with $d = 60$ and $A \in \{0, 1\}$ as follows:

1. $X' \sim \mathcal{N}(0, 0.5\mathbb{I}_d)$ where $\mathbb{I}_d$ is a $d \times d$ identity matrix.

2. $A \mid X' \sim \text{Ber}\left(p(X')\right)$ where

$$p\left(X'\right) = \frac{1}{1 + \exp(0.5 - 0.05\mathbf{e}_d{}^\top X')},$$

   where $\mathbf{e}_d \in \mathbb{R}^d$ is an all-one vector.

3. Draw W', Z', U from

$$W', Z', U \mid A, X' \sim \mathcal{N}\left(\begin{bmatrix}\mu_0 + \mu_a A + \mu_x X' \\ \alpha_0 + \alpha_a A + \alpha_x X' \\ \kappa_0 + \kappa_a A + \kappa_x X'\end{bmatrix}, \begin{bmatrix}\sigma_w^2, \sigma_{wz}^2, \sigma_{wu}^2 \\ \sigma_{wz}^2, \sigma_z^2, \sigma_{zu}^2 \\ \sigma_{wu}^2, \sigma_{zu}^2, \sigma_u^2\end{bmatrix}\right).$$

   Here we set the parameters above as $\mu_0 = \alpha_0 = \kappa_0 = 0.2\mathbf{e}_d$, $\alpha_a = \kappa_a = \mu_x = \alpha_x = \kappa_x = \mathbb{I}_d$, $\sigma_z^2 = \sigma_u^2 = \sigma_w^2 = 0.1(\mathbb{I}_d + \mathbf{e}_d\mathbf{e}_d^\top), \sigma_{wu}^2 = \sigma_{zu}^2 = 0.1\mathbf{e}_d\mathbf{e}_d^\top$. Finally, we choose σ_{wz}^2 and μ_a to ensure that $W' \perp (A, Z') \mid U, X'$, which corresponds to condition 4 in Assumption 1. To achieve this, note that

$$\mathbb{E}[W' \mid U, X', A, Z'] = \mu_0 + \mu_a A + \mu_x X' + \Sigma_{w(z,u)}\Sigma_{z,u}^{-1}\begin{bmatrix}Z' - \alpha_0 - \alpha_a A - \alpha_x X' \\ U - \kappa_0 - \kappa_a A - \kappa_x X'\end{bmatrix} \quad (47)$$

   where

$$\Sigma_{w(z,u)} = (\sigma_{wz}^2, \sigma_{wu}^2), \quad \Sigma_{z,u} = \begin{bmatrix}\sigma_z^2, \sigma_{zu}^2 \\ \sigma_{zu}^2, \sigma_u^2\end{bmatrix}.$$

   We simply select σ_{wz}^2 and μ_a so that Equation (47) does not depend on A and Z'.

4. Draw Y from

$$Y \mid X', U, W' \sim \mathcal{N}\left(A + \mathbf{e}_d^\top X' + \mathbf{e}_d^\top U + \mathbf{e}_d^\top W', 1\right).$$

5. Transform (W', Z', X') into (W, Z, X) via $X = g\left(GX'\right)$, $Z = g\left(GZ'\right)$, $W = g\left(GW'\right)$ where $G \in \mathbb{R}^{d\times d}$ is an invertible matrix and $g\left(\cdot\right)$ is a nonlinear function applied elementwise to GX', GZ', GW' respectively.

6. Make only (W, Z, X, A, Y) observable in the final data.

Following Cui et al. (2023), we can show that there exist a linear function $\tilde{h}(W', a, X')$ and a generalized linear function $\tilde{q}(Z', a, X')$ satisfying

$$\begin{aligned}\mathbb{E}[Y \mid U, a, X'] &= \mathbb{E}[\tilde{h}(W', a, X') \mid U, A = a, X'],\\ 1/\Pr(A = a \mid U, X') &= \mathbb{E}[\tilde{q}(Z', a, X') \mid U, A = a, X'], \qquad a = 0, 1.\end{aligned}$$

These induce bridge functions (when $g()$ is invertible)

$$h_0(W, A, X) = \tilde{h}\left(G^{-1}g^{-1}(W), a, G^{-1}g^{-1}(X)\right), \quad q_0(Z, A, X) = \tilde{q}\left(G^{-1}g^{-1}(Z), a, G^{-1}g^{-1}(X)\right).$$

For the experiments in Section 9.1, we use three-layer neural networks of the following form as the bridge class $\mathbb{H}$:

$$x_1 \leftarrow \mathrm{ReLU}(\Omega_1 x_0 + B_1), \Omega_1 \in \mathbb{R}^{2d\times 2d}, B_1 \in \mathbb{R}^{2d},\ x_2 \leftarrow \mathrm{ReLU}(\Omega_2 x_1 + B_2), \Omega_2 \in \mathbb{R}^{d\times d}, B_2 \in \mathbb{R}^d \tag{48}$$

$$x_3 \leftarrow (\Omega_3 x_2 + B_3), \Omega_3 \in \mathbb{R}^{1\times d}, B_3 \in \mathbb{R}.$$

The bridge class for $\mathbb{Q}$ is the same neural network class except that the last layer is replaced by $\mathrm{Softplus}(\Omega_3 x_2 + B_3)$. For the critic class $\mathbb{Q}'$, we use an RKHS with an RBF kernel $k(a, b) = \exp(-\|a-b\|^2/(2\iota))$ where the variance parameter ι is set to the median of pairwise distances in the data. When we use stabilizers, we set the parameters as $\gamma_1 = 5$ and $\lambda = 1$. Recall these hyperparameters appear in Equation (37) and Equation (38). For the critic class $\mathbb{H}'$, we use an RKHS with a linear kernel. Since the critic classes are both RKHS, computing the bridge function estimators only requires optimizing over $\mathbb{H}, \mathbb{Q}$ as we show in Lemma 6. We typically implement this using RMSprop (Tieleman and Hinton, 2012) with a learning rate 0.0002, momentum 0.95 and batch size $n/40$.

To evaluate the estimation performance, we run 1000 replications of each experiment. Thus for each estimator, we obtain 1000 estimated values $\hat{J}_k$ for $k = 1, \ldots, 1000$. We report the normalized mean squared errors (MSE) $\frac{1}{1000}\sum_{k=1}^{1000} M_k$ with $M_k = \frac{|\hat{J}_k - J|^2}{J^2}$, normalized squared bias $\left|\frac{1}{1000}\sum_{k=1}^{1000}\hat{J}_k - J\right|^2/J^2$, and the variance of $\left\{\hat{J}_k : k = 1, \ldots, 1000\right\}$ normalized by the squared true estimand J^2.

Ablation Study (No Unmeasured Confounding).

The goal of the ablation in Tables 6–7 is to remove unmeasured confounding. We keep a similar high-dimensional structure to that used above. The main change—and the reason this ablation has *no* unmeasured confounding—is how we decouple A and the latent noise U. In the main simulation, even after conditioning on X', the latent U still depends on A (its conditional mean includes an A-shift), and U also enters Y. Thus $A \not\perp U \mid X'$, so $A \not\perp (Y(1), Y(0)) \mid X$, and standard adjustment on X alone is biased.

In the ablation, we choose the conditional law so that, given X', the latent U no longer carries any A-specific shift (i.e., $\kappa_a = 0$). Equivalently, in the ablation we enforce $A \perp U \mid X'$ while keeping the rest of the model (the dimension, the nonlinear transformations g, the correlated Gaussian structure of (W', Z', U), and the outcome equation for Y) unchanged.

For $g(t) = \sin(t)$, the global non-invertibility of the transformation induces a small residual dependence after conditioning on the transformed covariates X. In the present design, this effect is negligible at the reported precision: both Plain REG and DR (no W, Z) have normalized squared bias reported as 0.0000 in Table 6, while the cubic transformation in Table 7 is one-to-one and satisfies the stated ignorability exactly.

Permutation		DR	DR(sta)	Linear(moment-based)
I	Treatment effects (SEs)	-0.0528(0.01194)	-0.0535 (0.01064)	-0.0828 (0.01977)
	95% CIs	[-0.0762, -0.0294]	[-0.0743, -0.0326]	[-0.1216, -0.0441]
II	Treatment effects (SEs)	-0.0515(0.01175)	-0.0619 (0.01065)	-0.0805 (0.02092)
	95% CIs	[-0.0745, -0.0284]	[-0.0828, -0.0411]	[-0.1215, -0.0395]
III	Treatment effects (SEs)	-0.0499 (0.01156)	-0.0630 (0.01037)	-0.0710 (0.02148)
	95% CIs	[-0.0725, -0.0272]	[-0.0833, -0.0427]	[-0.1132, -0.0289]
IV	Treatment effects (SEs)	-0.0505 (0.01232)	-0.0473 (0.01080)	-0.1048 (0.05429)
	95% CIs	[-0.0747, -0.0264]	[-0.0685, -0.0261]	[-0.2112, +0.0016]
V	Treatment effects (SEs)	-0.0454 (0.01176)	-0.0474 (0.01150)	+0.0435 (0.10965)
	95% CIs	[-0.0684, -0.0223]	[-0.0699, -0.0248]	[-0.1714, +0.2584]
VI	Treatment effects (SEs)	-0.0429 (0.01106)	-0.0590 (0.01084)	+0.0113 (0.07680)
	95% CIs	[-0.0646, -0.0213]	[-0.0802, -0.0377]	[-0.1393, +0.1618]

Permutation		Linear(Interaction)	Linear(Quadratic)	Linear(Both)
I	Treatment effects (SEs)	-0.0372 (0.07466)	-0.0859 (0.02427)	+0.1053 (0.97969)
	95% CIs	[-0.1835, +0.1091]	[-0.1334, -0.0383]	[-1.8149, +2.0255]
II	Treatment effects (SEs)	-0.0931 (0.11916)	-0.0672 (0.02117)	-0.0488 (0.04447)
	95% CIs	[-0.3267, +0.1404]	[-0.1087, -0.0258]	[-0.1359, +0.0384]
III	Treatment effects (SEs)	-0.0776 (0.01704)	+0.0325 (0.13767)	+0.1296 (0.53953)
	95% CIs	[-0.1110, -0.0442]	[-0.2373, +0.3023]	[-0.9279, +1.1870]
IV	Treatment effects (SEs)	-0.0716 (0.03930)	-0.0885 (0.03415)	-0.0897 (0.05779)
	95% CIs	[-0.1486, +0.0054]	[-0.1554, -0.0216]	[-0.2030, +0.0235]
V	Treatment effects (SEs)	+0.1496 (0.72227)	+0.1067 (0.09530)	+0.0828 (0.20450)
	95% CIs	[-1.2660, +1.5653]	[-0.0801, +0.2935]	[-0.3180, +0.4836]
VI	Treatment effects (SEs)	-0.1554 (0.24149)	+0.0964 (0.09538)	+0.6260 (3.24644)
	95% CIs	[-0.6287, +0.3179]	[-0.0905, +0.2833]	[-5.7370, +6.9890]

Table 13: The performance of different estimators across $\binom{4}{2} = 6$ permutations of the proxy variables Z and W. The linear moment-based baseline estimator is the same as the one in Table 8. Permutation I: Z=["pafi1", "paco21"],W=["ph1", "hema1"]; Permutation II: Z = ["pafi1", "ph1"], W = ["paco21", "hema1"] ; Permutation III: Z = ["pafi1", "hema1"], W = ["paco21", "ph1"]; Permutation IV: Z = ["paco21", "ph1"], W = ["pafi1", "hema1"]; Permutation V: Z = ["paco21", "hema1"], W = ["pafi1", "ph1"]; Permutation VI: Z = ["ph1", "hema1"], W = ["pafi1", "paco21"].

H.2 Real Data Analysis

Because the true treatment effect is unavailable in the RHC study, we assess robustness along two distinct dimensions. The first analysis varies the roles assigned to the candidate proxy variables, whereas the second progressively removes baseline covariates to assess sensitivity to unmeasured confounding.

H.2.1 Robustness to Proxy-Variable Allocation

In the primary RHC analysis in the main text, we follow Cui et al. (2023) and use $Z = (\texttt{pafi1}, \texttt{paco21}), W = (\texttt{ph1}, \texttt{hema1})$ as the proxy variables. These proxy variables are in fact selected according to the procedure proposed in Tchetgen et al. (2020).

Permutation		DR	DR(sta)	Linear(closed-form)
I	Treatment effects (SEs)	-0.0528 (0.01194)	-0.0535 (0.01064)	-0.0860 (0.02408)
	95% CIs	[-0.0762, -0.0294]	[-0.0743, -0.0326]	[-0.1332, -0.0388]
II	Treatment effects (SEs)	-0.0515 (0.01175)	-0.0619 (0.01065)	-0.0808 (0.02439)
	95% CIs	[-0.0745, -0.0284]	[-0.0828, -0.0411]	[-0.1286, -0.0330]
III	Treatment effects (SEs)	-0.0499 (0.01156)	-0.0630 (0.01037)	-0.1588 (1.95893)
	95% CIs	[-0.0725, -0.0272]	[-0.0833, -0.0427]	[-3.9983, +3.6807]
IV	Treatment effects (SEs)	-0.0505 (0.01232)	-0.0473 (0.01080)	+1.0435 (25.86638)
	95% CIs	[-0.0747, -0.0264]	[-0.0685, -0.0261]	[-49.6546, +51.7416]
V	Treatment effects (SEs)	-0.0454 (0.01176)	-0.0474 (0.01150)	+0.1193 (0.36014)
	95% CIs	[-0.0684, -0.0223]	[-0.0699, -0.0248]	[-0.5866, +0.8252]
VI	Treatment effects (SEs)	-0.0429 (0.01106)	-0.0590 (0.01084)	+0.0299 (0.08592)
	95% CIs	[-0.0646, -0.0213]	[-0.0802, -0.0377]	[-0.1385, +0.1983]

Permutation		Linear(Interaction)	Linear(Quadratic)	Linear(Both)
I	Treatment effects(SEs)	+0.0017(5.32580)	+0.0301 (0.28169)	-0.2281(2.21813)
	95% CIs	[-10.4369, +10.4402]	[-0.5220, +0.5822]	[-4.5756, +4.1194]
II	Treatment effects (SEs)	-0.9171 (36.43500)	-0.2175 (0.56159)	-0.0269 (0.14245)
	95% CIs	[-72.3297, +70.4955]	[-1.3182, +0.8832]	[-0.3061, +0.2523]
III	Treatment effects (SEs)	-0.0926 (0.07764)	-0.8109 (17.23161)	-0.4625 (2.08727)
	95% CIs	[-0.2448, +0.0596]	[-34.5849, +32.9630]	[-4.5535, +3.6286]
IV	Treatment effects (SEs)	-0.0820 (0.18875)	-0.2891 (4.99796)	-0.0683 (0.32093)
	95% CIs	[-0.4520, +0.2879]	[-10.0851, +9.5069]	[-0.6973, +0.5607]
V	Treatment effects (SEs)	+0.5472 (30.68336)	+0.1404 (0.31630)	+0.2647 (0.74140)
	95% CIs	[-59.5922, +60.6866]	[-0.4795, +0.7604]	[-1.1884, +1.7179]
VI	Treatment effects (SEs)	-4.4027 (408.41541)	+0.0538 (0.22008)	-1.4556 (20.80650)
	95% CIs	[-804.8970, +796.0915]	[-0.3776, +0.4851]	[-42.2364, +39.3251]

Table 14: The performance of different estimators across $\binom{4}{2} = 6$ permutations of the proxy variables Z and W as described in Table 13. The linear closed-form baseline estimator is the same as the one in Table 8.

Among the ten physiological measurements collected within the first 24 hours of ICU admission, Tchetgen et al. (2020) first identify a reduced set of candidate proxy variables. Specifically, each of the ten variables (`serum sodium, serum potassium, serum creatinine, bilirubin, albumin, PaO2/(.01*FiO2) ratio, PaCO2, serum PH (arterial), white blood cell count, and hematocrit`) is evaluated based on its association with the treatment A and the outcome Y: treatment associations are assessed via logistic regression of A on the candidate variable and baseline covariates, while outcome associations are assessed via linear regression of Y on A, the candidate variable, and baseline covariates. Based on the strength of these associations, Tchetgen et al. (2020) select the four variables (`pafi1, ph1, paco21, hema1`) that are strongly related to both A and Y as candidate proxies, while treating the remaining variables as covariates X. Tchetgen et al. (2020) state that they "hypothesize that the four selected variables are a priori equally likely to serve as treatment-inducing proxies or outcome-inducing proxies". They propose to decide the allocation of the four variables to the negative control actions and negative control outcome proxies by ranking the associations of these variables with the outcome and treatment, respectively. Specifically, they "select proxies in decreasing order of strength of association, first selecting the proxy with strongest association with the outcome as outcome-inducing proxy and likewise for the treatment". They break the ties in a particular way during the process and end up with the proxy variable combination $Z = (\texttt{pafi1}, \texttt{paco21}), W = (\texttt{ph1}, \texttt{hema1})$.

In this part, we assess the robustness of the estimation results to different allocations of the four proxy variables. From the four proxy variables, we consider all possible ways of allocating two variables to the negative control treatment set Z and the remaining two to the negative control outcome set W, yielding $\binom{4}{2} = 6$ distinct (Z, W) settings. For each setting, we apply the same proximal estimation methods and report the results in Table 13 and Table 14. We can observe that our proposed DR estimators based on flexible minimax bridge function estimation are overall stable across different permutations of the proxy variables. In contrast, the linear moment-based baseline estimator and the linear closed-form baseline estimator are much more sensitive. They can easily give estimates with opposite signs to the original estimates in Table 8. Moreover, they are still sensitive to the choice of model specification (e.g., incorporating interaction, quadratic or both terms or not), across all six permutations.

H.2.2 Covariate-Omission Stress Test

This stress test examines the empirical sensitivity of the RHC conclusions to losing baseline covariates. We keep the proxy-variable allocation fixed at $Z = (\texttt{pafi1}, \texttt{paco21})$ and $W = (\texttt{ph1}, \texttt{hema1})$, but deliberately withhold selected observed covariates from X and refit the complete analysis. Thus, the exercise does not constitute a formal sensitivity analysis for arbitrary unmeasured confounding; rather, it assesses whether the conclusions remain stable under structured and increasingly severe omissions of information that is observed in the original data.

The original X contains 49 clinically meaningful raw-variable blocks, represented by 68 columns after one-hot encoding. We always remove all encoded columns belonging to a raw variable together, so that each omission scenario has a clear clinical interpretation. Together

with the four proxy variables, these variables account for the 72 covariates used in the main RHC analysis.

We construct two deterministic omission paths to avoid relying on an arbitrarily chosen set of covariates. The first, denoted the *A/Y-predictive* path, ranks each covariate block by its joint predictive strength for treatment and outcome: we fit ℓ_2-regularized regressions for A given X and for Y given (A, X), and rank block j by $(s_{A,j} s_{Y,j})^{1/2}$, where $s_{A,j}$ and $s_{Y,j}$ are standardized coefficient-block norms. The second, denoted the *proxy-predictive* path, ranks blocks by their standardized coefficient-block norms from a multivariate ridge regression of the four proxies (Z, W) on X. These rankings are used only to define omission scenarios and are not used to tune any treatment-effect estimator.

Along each path, we remove the top $k \in \{1, 2, 3, 4, 5, 10, 15, 20, 25, 30, 35, 40, 45\}$ blocks. We additionally report the full-X analysis and a final union scenario that removes the union of the two top-45 sets. Because the two top-45 sets overlap in 41 blocks, their union contains all 49 blocks and leaves no observed X covariates. Hence, the analysis comprises $1 + 2 \times 13 + 1 = 28$ scenarios. For every scenario, we refit all nuisance functions and effect estimators from scratch using the same preprocessing, two-fold cross-fitting, training epochs, random seeds, and tuning parameters as in the main RHC analysis.

Figures 3 and 4 display point estimates and 95% confidence intervals along the two omission paths; the solid vertical line denotes zero and the dashed red line denotes the corresponding full-X estimate. Tables 15 and 16 provide the numerical results for the moment-based and closed-form linear baselines, respectively.

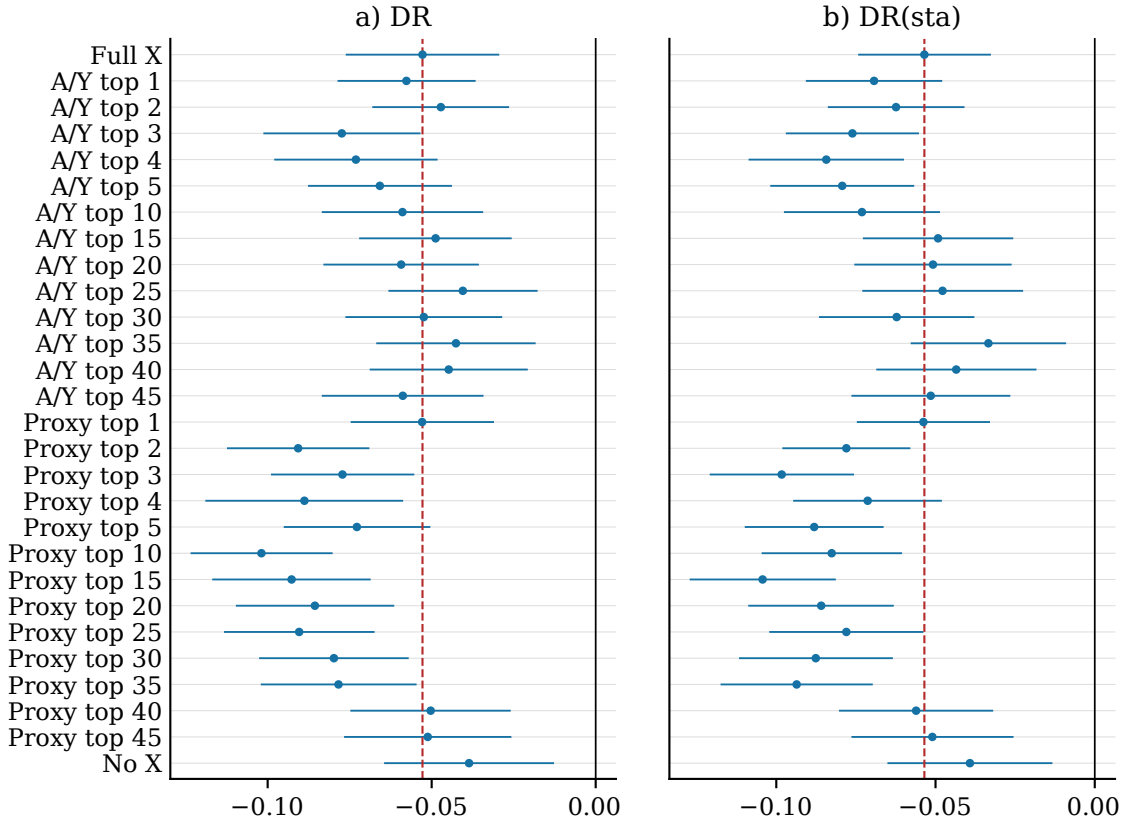


Figure 3: Covariate-omission stress test for DR, DR(sta), and moment-based linear specifications. Points and horizontal bars show treatment-effect estimates and 95% confidence intervals. "A/Y top k" and "Proxy top k" remove the k highest-ranked covariate blocks under the treatment–outcome and proxy-predictive rankings, respectively. The solid vertical line denotes zero; the dashed red line denotes the full-X estimate.

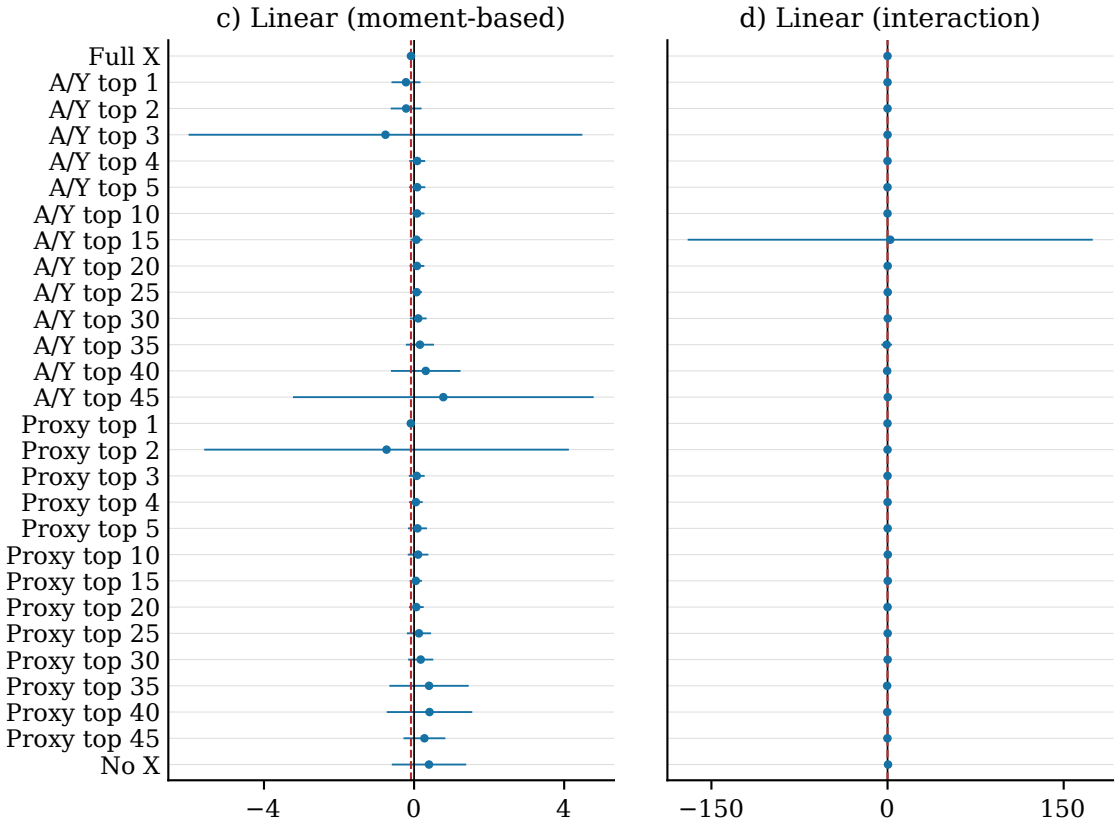


Figure 3: Covariate-omission stress test for moment-based linear specifications (continued): simple and interaction bases. The scenario labels, points, confidence intervals, and reference lines are as in the preceding panels.

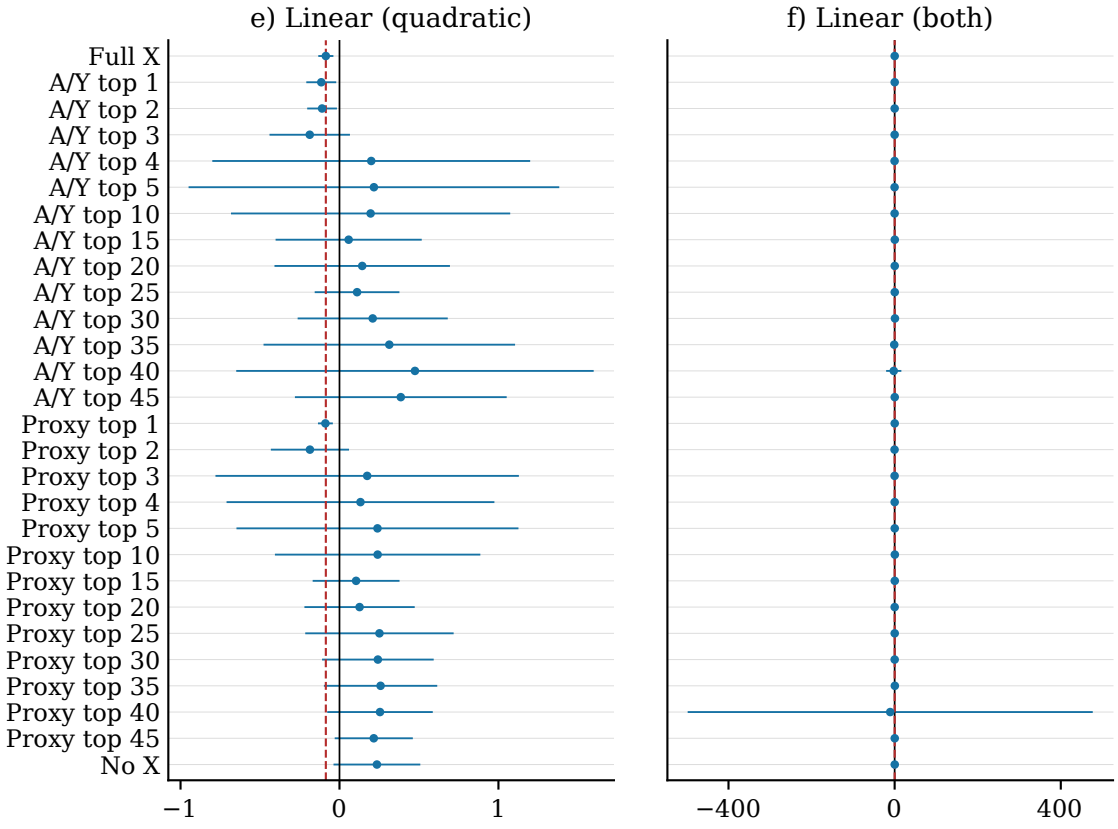


Figure 3: Covariate-omission stress test for moment-based linear specifications (continued): quadratic and combined bases. The scenario labels, points, confidence intervals, and reference lines are as in the preceding panels.

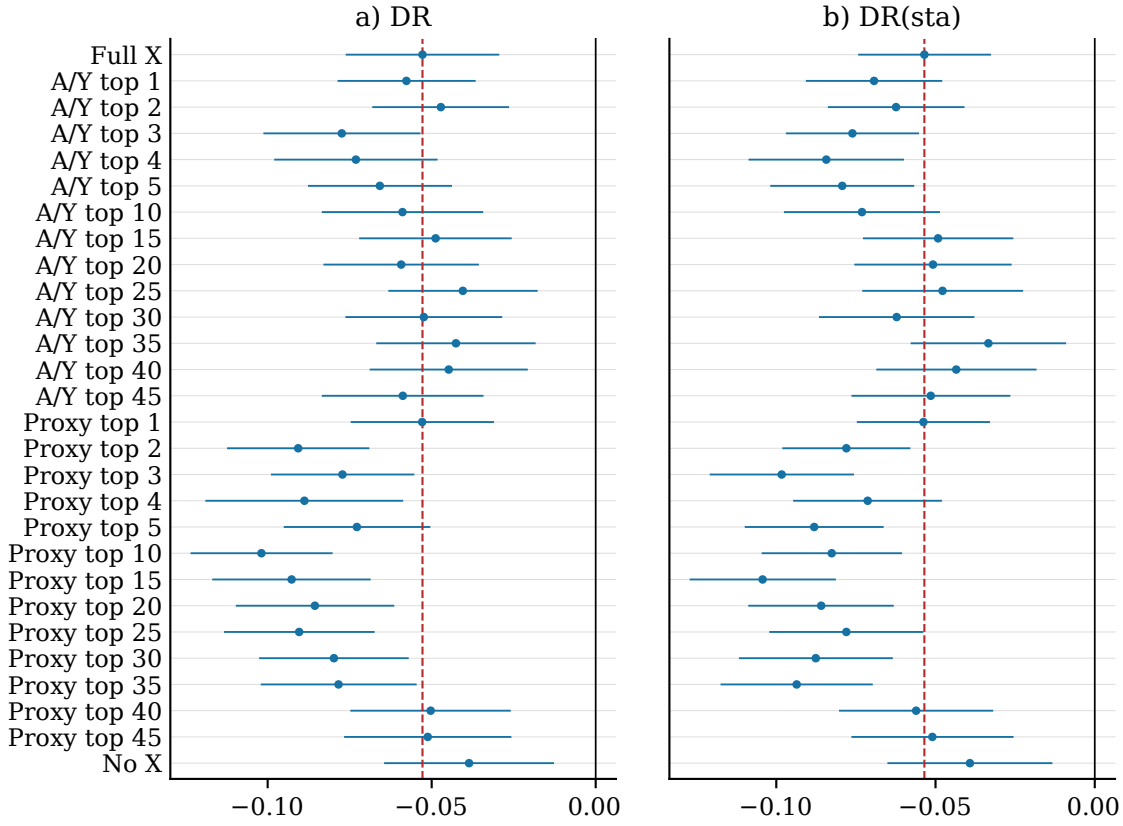


Figure 4: Covariate-omission stress test for DR, DR(sta), and closed-form linear specifications. Points and horizontal bars show treatment-effect estimates and 95% confidence intervals. "A/Y top k" and "Proxy top k" remove the k highest-ranked covariate blocks under the treatment–outcome and proxy-predictive rankings, respectively. The solid vertical line denotes zero; the dashed red line denotes the full-X estimate.

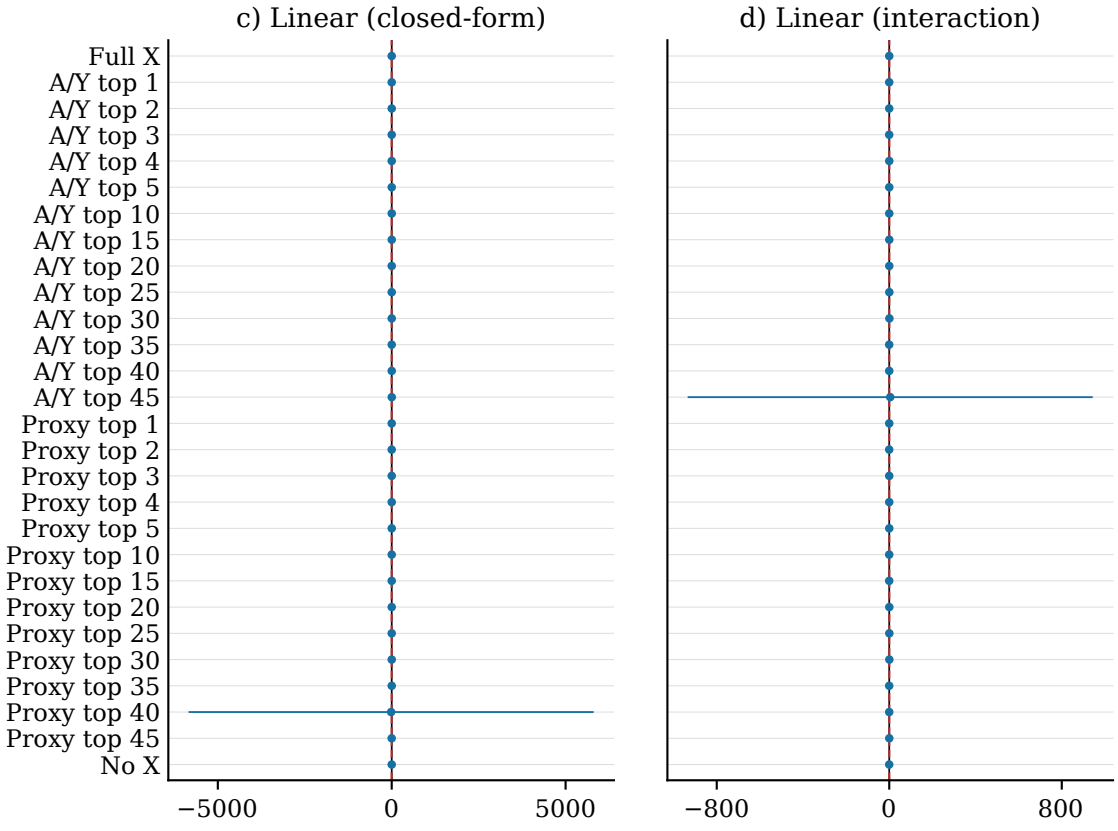


Figure 4: Covariate-omission stress test for closed-form linear specifications (continued): simple and interaction bases. The scenario labels, points, confidence intervals, and reference lines are as in the preceding panels.

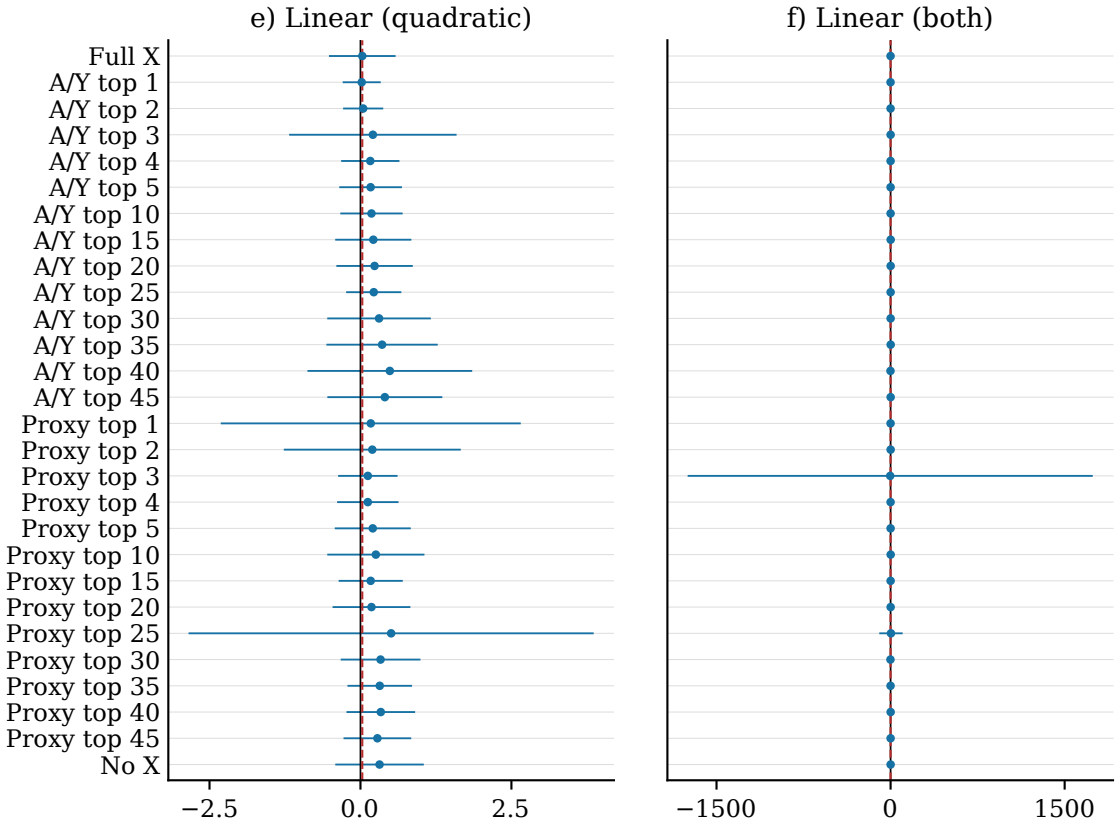


Figure 4: Covariate-omission stress test for closed-form linear specifications (continued): quadratic and combined bases. The scenario labels, points, confidence intervals, and reference lines are as in the preceding panels.

X omission	DR	DR(sta)	Linear (moment-based)
Full X	**−0.0528 (0.01194)** [−0.0762, −0.0294]	**−0.0535 (0.01064)** [−0.0743, −0.0326]	**−0.0828 (0.01977)** [−0.1216, −0.0441]
A/Y top 1	**−0.0577 (0.01073)** [−0.0787, −0.0366]	**−0.0693 (0.01091)** [−0.0907, −0.0479]	−0.2160 (0.19689) [−0.6019, +0.1699]
A/Y top 2	**−0.0472 (0.01064)** [−0.0681, −0.0264]	**−0.0624 (0.01096)** [−0.0838, −0.0409]	−0.2124 (0.20980) [−0.6236, +0.1989]
A/Y top 3	**−0.0774 (0.01223)** [−0.1013, −0.0534]	**−0.0761 (0.01065)** [−0.0970, −0.0552]	−0.7625 (2.67704) [−6.0095, +4.4845]
A/Y top 4	**−0.0731 (0.01272)** [−0.0980, −0.0482]	**−0.0843 (0.01243)** [−0.1087, −0.0599]	+0.0828 (0.10833) [−0.1296, +0.2951]
A/Y top 5	**−0.0658 (0.01120)** [−0.0877, −0.0438]	**−0.0793 (0.01155)** [−0.1019, −0.0567]	+0.0833 (0.10994) [−0.1322, +0.2988]
A/Y top 10	**−0.0589 (0.01255)** [−0.0835, −0.0343]	**−0.0731 (0.01250)** [−0.0976, −0.0486]	+0.0793 (0.10028) [−0.1172, +0.2759]
A/Y top 15	**−0.0488 (0.01187)** [−0.0721, −0.0256]	**−0.0492 (0.01205)** [−0.0728, −0.0256]	+0.0614 (0.08117) [−0.0976, +0.2205]
A/Y top 20	**−0.0593 (0.01210)** [−0.0830, −0.0356]	**−0.0508 (0.01260)** [−0.0755, −0.0261]	+0.0760 (0.10037) [−0.1207, +0.2727]
A/Y top 25	**−0.0405 (0.01159)** [−0.0632, −0.0177]	**−0.0478 (0.01289)** [−0.0730, −0.0225]	+0.0690 (0.07118) [−0.0705, +0.2085]
A/Y top 30	**−0.0524 (0.01220)** [−0.0763, −0.0285]	**−0.0622 (0.01246)** [−0.0866, −0.0378]	+0.1103 (0.11384) [−0.1129, +0.3334]
A/Y top 35	**−0.0426 (0.01239)** [−0.0669, −0.0183]	**−0.0334 (0.01245)** [−0.0578, −0.0090]	+0.1559 (0.19132) [−0.2191, +0.5309]
A/Y top 40	**−0.0448 (0.01230)** [−0.0689, −0.0207]	**−0.0435 (0.01285)** [−0.0686, −0.0183]	+0.3104 (0.47300) [−0.6167, +1.2375]
A/Y top 45	**−0.0588 (0.01258)** [−0.0835, −0.0342]	**−0.0515 (0.01272)** [−0.0764, −0.0265]	+0.7789 (2.04441) [−3.2282, +4.7859]
Proxy top 1	**−0.0529 (0.01114)** [−0.0747, −0.0310]	**−0.0538 (0.01067)** [−0.0747, −0.0329]	**−0.0882 (0.02086)** [−0.1291, −0.0474]
Proxy top 2	**−0.0907 (0.01108)** [−0.1124, −0.0690]	**−0.0780 (0.01026)** [−0.0981, −0.0579]	−0.7343 (2.47874) [−5.5926, +4.1241]
Proxy top 3	**−0.0772 (0.01115)** [−0.0990, −0.0553]	**−0.0983 (0.01155)** [−0.1209, −0.0756]	+0.0708 (0.10719) [−0.1393, +0.2809]
Proxy top 4	**−0.0888 (0.01536)** [−0.1190, −0.0587]	**−0.0713 (0.01191)** [−0.0947, −0.0480]	+0.0485 (0.09224) [−0.1323, +0.2293]
Proxy top 5	**−0.0728 (0.01138)** [−0.0951, −0.0504]	**−0.0881 (0.01112)** [−0.1099, −0.0663]	+0.0902 (0.12862) [−0.1619, +0.3423]
Proxy top 10	**−0.1019 (0.01105)** [−0.1235, −0.0802]	**−0.0826 (0.01125)** [−0.1046, −0.0605]	+0.1062 (0.14018) [−0.1685, +0.3810]
Proxy top 15	**−0.0927 (0.01233)** [−0.1169, −0.0686]	**−0.1043 (0.01172)** [−0.1272, −0.0813]	+0.0470 (0.08288) [−0.1155, +0.2094]
Proxy top 20	**−0.0856 (0.01232)** [−0.1097, −0.0614]	**−0.0859 (0.01166)** [−0.1088, −0.0631]	+0.0594 (0.10014) [−0.1369, +0.2556]
Proxy top 25	**−0.0904 (0.01172)** [−0.1133, −0.0674]	**−0.0780 (0.01235)** [−0.1022, −0.0538]	+0.1295 (0.16461) [−0.1931, +0.4522]
Proxy top 30	**−0.0798 (0.01163)** [−0.1026, −0.0570]	**−0.0876 (0.01232)** [−0.1117, −0.0634]	+0.1757 (0.17122) [−0.1599, +0.5113]
Proxy top 35	**−0.0784 (0.01212)** [−0.1021, −0.0546]	**−0.0936 (0.01219)** [−0.1175, −0.0697]	+0.3989 (0.53878) [−0.6572, +1.4549]
Proxy top 40	**−0.0503 (0.01246)** [−0.0748, −0.0259]	**−0.0561 (0.01236)** [−0.0803, −0.0319]	+0.4115 (0.58097) [−0.7272, +1.5502]
Proxy top 45	**−0.0512 (0.01301)** [−0.0767, −0.0257]	**−0.0510 (0.01299)** [−0.0764, −0.0255]	+0.2756 (0.28496) [−0.2829, +0.8341]
Top-45 union (no X)	**−0.0386 (0.01322)** [−0.0645, −0.0127]	**−0.0392 (0.01322)** [−0.0651, −0.0133]	+0.3976 (0.50583) [−0.5938, +1.3890]

Table 15: Covariate-omission analogue with moment-based linear specifications, panel (a): DR, DR(sta), and simple linear. Each cell reports the effect estimate (standard error) above its 95% confidence interval; bold cells have upper confidence limits below zero.

X omission	Linear (Interaction)	Linear (Quadratic)	Linear (Both)
Full X	−0.0372 (0.07466) [−0.1835, +0.1091]	**−0.0859 (0.02427)** [−0.1334, −0.0383]	+0.1053 (0.97969) [−1.8149, +2.0255]
A/Y top 1	−0.0237 (0.06570) [−0.1524, +0.1051]	**−0.1146 (0.04802)** [−0.2088, −0.0205]	+0.1316 (2.64882) [−5.0601, +5.3233]
A/Y top 2	−0.0122 (0.06684) [−0.1432, +0.1188]	**−0.1096 (0.04773)** [−0.2032, −0.0161]	+0.0787 (2.06838) [−3.9753, +4.1328]
A/Y top 3	−0.0007 (0.10310) [−0.2028, +0.2014]	−0.1872 (0.12910) [−0.4403, +0.0658]	−0.0118 (0.44009) [−0.8743, +0.8508]
A/Y top 4	−0.0271 (0.04441) [−0.1142, +0.0599]	+0.1997 (0.51015) [−0.8002, +1.1996]	−0.1419 (0.14219) [−0.4206, +0.1368]
A/Y top 5	−0.0270 (0.04419) [−0.1137, +0.0596]	+0.2166 (0.59490) [−0.9494, +1.3826]	−0.1459 (0.14819) [−0.4363, +0.1446]
A/Y top 10	−0.0054 (0.05192) [−0.1072, +0.0963]	+0.1958 (0.44810) [−0.6825, +1.0741]	−0.0941 (0.14009) [−0.3687, +0.1804]
A/Y top 15	+2.2670 (88.00725) [−170.2272, +174.7612]	+0.0577 (0.23454) [−0.4020, +0.5174]	+0.3397 (2.16806) [−3.9097, +4.5891]
A/Y top 20	+0.0982 (0.22882) [−0.3503, +0.5467]	+0.1429 (0.28161) [−0.4090, +0.6949]	+0.2442 (1.15575) [−2.0211, +2.5094]
A/Y top 25	+0.1427 (0.44011) [−0.7200, +1.0053]	+0.1106 (0.13600) [−0.1560, +0.3771]	+0.2393 (0.84064) [−1.4084, +1.8869]
A/Y top 30	+0.1804 (0.24856) [−0.3068, +0.6676]	+0.2089 (0.24092) [−0.2633, +0.6811]	+0.7489 (2.93033) [−4.9945, +6.4924]
A/Y top 35	−0.7394 (2.31807) [−5.2828, +3.8040]	+0.3133 (0.40369) [−0.4779, +1.1046]	−0.9834 (2.50975) [−5.9025, +3.9357]
A/Y top 40	−0.3323 (0.52529) [−1.3619, +0.6973]	+0.4748 (0.57366) [−0.6496, +1.5992]	−2.1529 (9.35104) [−20.4809, +16.1751]
A/Y top 45	+0.1732 (0.35572) [−0.5240, +0.8704]	+0.3857 (0.33988) [−0.2804, +1.0519]	+0.1987 (0.30957) [−0.4080, +0.8055]
Proxy top 1	−0.0338 (0.13887) [−0.3060, +0.2384]	**−0.0886 (0.02385)** [−0.1353, −0.0418]	+0.1565 (1.64385) [−3.0655, +3.3784]
Proxy top 2	+0.0032 (0.11144) [−0.2152, +0.2216]	−0.1860 (0.12546) [−0.4319, +0.0599]	−0.2244 (1.72596) [−3.6073, +3.1585]
Proxy top 3	−0.0184 (0.04342) [−0.1035, +0.0667]	+0.1742 (0.48693) [−0.7802, +1.1286]	−0.1019 (0.11032) [−0.3182, +0.1143]
Proxy top 4	+0.0465 (0.08656) [−0.1232, +0.2162]	+0.1320 (0.42988) [−0.7105, +0.9746]	+0.0899 (0.47028) [−0.8318, +1.0117]
Proxy top 5	+0.0832 (0.11481) [−0.1418, +0.3083]	+0.2390 (0.45256) [−0.6481, +1.1260]	+0.2179 (0.46996) [−0.7032, +1.1391]
Proxy top 10	+0.1977 (0.35576) [−0.4996, +0.8950]	+0.2400 (0.32968) [−0.4062, +0.8862]	+0.4553 (1.24366) [−1.9823, +2.8928]
Proxy top 15	+0.1633 (0.33528) [−0.4938, +0.8205]	+0.1044 (0.13940) [−0.1688, +0.3777]	+0.3236 (0.85749) [−1.3570, +2.0043]
Proxy top 20	+0.0192 (0.23968) [−0.4506, +0.4889]	+0.1264 (0.17701) [−0.2206, +0.4733]	+0.0684 (0.38524) [−0.6867, +0.8234]
Proxy top 25	+0.1000 (0.22317) [−0.3374, +0.5375]	+0.2514 (0.23815) [−0.2154, +0.7182]	+0.1475 (0.47312) [−0.7798, +1.0748]
Proxy top 30	+0.0635 (0.33287) [−0.5889, +0.7159]	+0.2415 (0.17917) [−0.1096, +0.5927]	+0.0673 (0.60075) [−1.1102, +1.2447]
Proxy top 35	−0.3444 (0.39895) [−1.1263, +0.4375]	+0.2585 (0.18191) [−0.0981, +0.6150]	+0.4921 (0.59737) [−0.6788, +1.6629]
Proxy top 40	−0.2305 (0.13512) [−0.4954, +0.0343]	+0.2549 (0.16939) [−0.0771, +0.5869]	−10.6770 (248.79711) [−498.3193, +476.9653]
Proxy top 45	−0.1269 (1.65824) [−3.3770, +3.1233]	+0.2157 (0.12537) [−0.0301, +0.4614]	+0.2299 (0.20345) [−0.1689, +0.6286]
Top-45 union (no X)	+0.3976 (0.50583) [−0.5938, +1.3890]	+0.2355 (0.13931) [−0.0376, +0.5085]	+0.2355 (0.13931) [−0.0376, +0.5085]

Table 15: Covariate-omission analogue with moment-based linear specifications, panel (b), continued: interaction, quadratic, and combined linear. Each cell reports the effect estimate (standard error) above its 95% confidence interval; bold cells have upper confidence limits below zero.

X omission	DR	DR(sta)	Linear (closed-form)
Full X	**−0.0528 (0.01194)** [−0.0762, −0.0294]	**−0.0535 (0.01064)** [−0.0743, −0.0326]	**−0.0860 (0.02408)** [−0.1332, −0.0388]
A/Y top 1	**−0.0577 (0.01073)** [−0.0787, −0.0366]	**−0.0693 (0.01091)** [−0.0907, −0.0479]	−0.0455 (0.21887) [−0.4745, +0.3835]
A/Y top 2	**−0.0472 (0.01064)** [−0.0681, −0.0264]	**−0.0624 (0.01096)** [−0.0838, −0.0409]	−0.0313 (0.20944) [−0.4418, +0.3792]
A/Y top 3	**−0.0774 (0.01223)** [−0.1013, −0.0534]	**−0.0761 (0.01065)** [−0.0970, −0.0552]	−0.1187 (0.16266) [−0.4375, +0.2001]
A/Y top 4	**−0.0731 (0.01272)** [−0.0980, −0.0482]	**−0.0843 (0.01243)** [−0.1087, −0.0599]	+0.3822 (0.82374) [−1.2323, +1.9967]
A/Y top 5	**−0.0658 (0.01120)** [−0.0877, −0.0438]	**−0.0793 (0.01155)** [−0.1019, −0.0567]	+0.3925 (0.87208) [−1.3167, +2.1018]
A/Y top 10	**−0.0589 (0.01255)** [−0.0835, −0.0343]	**−0.0731 (0.01250)** [−0.0976, −0.0486]	+0.5083 (1.31866) [−2.0763, +3.0928]
A/Y top 15	**−0.0488 (0.01187)** [−0.0721, −0.0256]	**−0.0492 (0.01205)** [−0.0728, −0.0256]	+0.4872 (1.27862) [−2.0189, +2.9933]
A/Y top 20	**−0.0593 (0.01210)** [−0.0830, −0.0356]	**−0.0508 (0.01260)** [−0.0755, −0.0261]	+0.4307 (1.03634) [−1.6005, +2.4619]
A/Y top 25	**−0.0405 (0.01159)** [−0.0632, −0.0177]	**−0.0478 (0.01289)** [−0.0730, −0.0225]	+0.2415 (0.30545) [−0.3572, +0.8402]
A/Y top 30	**−0.0524 (0.01220)** [−0.0763, −0.0285]	**−0.0622 (0.01246)** [−0.0866, −0.0378]	+0.3082 (0.47931) [−0.6313, +1.2476]
A/Y top 35	**−0.0426 (0.01239)** [−0.0669, −0.0183]	**−0.0334 (0.01245)** [−0.0578, −0.0090]	+0.3763 (0.71757) [−1.0301, +1.7828]
A/Y top 40	**−0.0448 (0.01230)** [−0.0689, −0.0207]	**−0.0435 (0.01285)** [−0.0686, −0.0183]	+0.7328 (2.59698) [−4.3573, +5.8228]
A/Y top 45	**−0.0588 (0.01258)** [−0.0835, −0.0342]	**−0.0515 (0.01272)** [−0.0764, −0.0265]	+1.2266 (9.11972) [−16.6481, +19.1012]
Proxy top 1	**−0.0529 (0.01114)** [−0.0747, −0.0310]	**−0.0538 (0.01067)** [−0.0747, −0.0329]	**−0.1026 (0.02854)** [−0.1585, −0.0466]
Proxy top 2	**−0.0907 (0.01108)** [−0.1124, −0.0690]	**−0.0780 (0.01026)** [−0.0981, −0.0579]	−0.1310 (0.15201) [−0.4289, +0.1670]
Proxy top 3	**−0.0772 (0.01115)** [−0.0990, −0.0553]	**−0.0983 (0.01155)** [−0.1209, −0.0756]	+0.3739 (0.87271) [−1.3366, +2.0844]
Proxy top 4	**−0.0888 (0.01536)** [−0.1190, −0.0587]	**−0.0713 (0.01191)** [−0.0947, −0.0480]	+0.3076 (0.75082) [−1.1640, +1.7792]
Proxy top 5	**−0.0728 (0.01138)** [−0.0951, −0.0504]	**−0.0881 (0.01112)** [−0.1099, −0.0663]	+0.2866 (0.55437) [−0.8000, +1.3732]
Proxy top 10	**−0.1019 (0.01105)** [−0.1235, −0.0802]	**−0.0826 (0.01125)** [−0.1046, −0.0605]	+0.2916 (0.51606) [−0.7198, +1.3031]
Proxy top 15	**−0.0927 (0.01233)** [−0.1169, −0.0686]	**−0.1043 (0.01172)** [−0.1272, −0.0813]	+0.2738 (0.54692) [−0.7981, +1.3458]
Proxy top 20	**−0.0856 (0.01232)** [−0.1097, −0.0614]	**−0.0859 (0.01166)** [−0.1088, −0.0631]	+0.2803 (0.61124) [−0.9177, +1.4784]
Proxy top 25	**−0.0904 (0.01172)** [−0.1133, −0.0674]	**−0.0780 (0.01235)** [−0.1022, −0.0538]	+0.1919 (0.24210) [−0.2826, +0.6665]
Proxy top 30	**−0.0798 (0.01163)** [−0.1026, −0.0570]	**−0.0876 (0.01232)** [−0.1117, −0.0634]	+0.2485 (0.26836) [−0.2775, +0.7745]
Proxy top 35	**−0.0784 (0.01212)** [−0.1021, −0.0546]	**−0.0936 (0.01219)** [−0.1175, −0.0697]	+0.5998 (1.77453) [−2.8783, +4.0779]
Proxy top 40	**−0.0503 (0.01246)** [−0.0748, −0.0259]	**−0.0561 (0.01236)** [−0.0803, −0.0319]	−16.6860 (2971.51473) [−5840.8549, +5807.4828]
Proxy top 45	**−0.0512 (0.01301)** [−0.0767, −0.0257]	**−0.0510 (0.01299)** [−0.0764, −0.0255]	−1.1029 (14.56344) [−29.6473, +27.4414]
Top-45 union (no X)	**−0.0386 (0.01322)** [−0.0645, −0.0127]	**−0.0392 (0.01322)** [−0.0651, −0.0133]	−0.2408 (0.87537) [−1.9565, +1.4749]

Table 16: Covariate-omission analogue with closed-form linear specifications, panel (a): DR, DR(sta), and simple linear. Each cell reports the effect estimate (standard error) above its 95% confidence interval; bold cells have upper confidence limits below zero.

X omission	Linear (Interaction)	Linear (Quadratic)	Linear (Both)
Full X	+0.0017 (5.32580) [−10.4369, +10.4402]	+0.0301 (0.28169) [−0.5220, +0.5822]	−0.2281 (2.21813) [−4.5756, +4.1194]
A/Y top 1	**−0.1369 (0.06844)** [−0.2710, −0.0027]	+0.0207 (0.16058) [−0.2941, +0.3354]	−0.3788 (3.96309) [−8.1464, +7.3889]
A/Y top 2	−0.1212 (0.06817) [−0.2548, +0.0124]	+0.0450 (0.16952) [−0.2873, +0.3773]	−0.2390 (1.93839) [−4.0382, +3.5602]
A/Y top 3	−0.2382 (0.29158) [−0.8097, +0.3333]	+0.2054 (0.70717) [−1.1806, +1.5915]	+0.2710 (2.09452) [−3.8342, +4.3763]
A/Y top 4	+0.0320 (0.05937) [−0.0843, +0.1484]	+0.1626 (0.24639) [−0.3203, +0.6455]	+0.0048 (0.09888) [−0.1890, +0.1986]
A/Y top 5	+0.0325 (0.05943) [−0.0840, +0.1490]	+0.1681 (0.26488) [−0.3511, +0.6873]	+0.0040 (0.10115) [−0.1942, +0.2023]
A/Y top 10	+0.0465 (0.05949) [−0.0701, +0.1631]	+0.1822 (0.26321) [−0.3337, +0.6981]	+0.0063 (0.10715) [−0.2038, +0.2163]
A/Y top 15	+0.1881 (0.62168) [−1.0304, +1.4066]	+0.2125 (0.32162) [−0.4179, +0.8428]	+0.7678 (7.80289) [−14.5259, +16.0615]
A/Y top 20	+0.1980 (0.55698) [−0.8936, +1.2897]	+0.2337 (0.32290) [−0.3992, +0.8666]	+0.5704 (1.50678) [−2.3828, +3.5237]
A/Y top 25	+0.3289 (0.61259) [−0.8718, +1.5296]	+0.2201 (0.23303) [−0.2367, +0.6768]	+0.3053 (0.56427) [−0.8006, +1.4113]
A/Y top 30	+0.5785 (4.67175) [−8.5781, +9.7351]	+0.3078 (0.43772) [−0.5501, +1.1657]	+0.3648 (0.78970) [−1.1830, +1.9126]
A/Y top 35	−0.2261 (0.43047) [−1.0698, +0.6176]	+0.3580 (0.47028) [−0.5638, +1.2797]	+0.7547 (8.04350) [−15.0105, +16.5200]
A/Y top 40	−0.0929 (0.17329) [−0.4326, +0.2468]	+0.4864 (0.69580) [−0.8773, +1.8502]	−0.9167 (8.91968) [−18.3993, +16.5659]
A/Y top 45	+4.4303 (479.22105) [−934.8430, +943.7035]	+0.4038 (0.48580) [−0.5484, +1.3559]	+0.5914 (1.74717) [−2.8331, +4.0159]
Proxy top 1	−0.0281 (0.07174) [−0.1687, +0.1125]	+0.1708 (1.26729) [−2.3131, +2.6546]	−0.1116 (0.88010) [−1.8366, +1.6134]
Proxy top 2	−0.2657 (0.22036) [−0.6976, +0.1662]	+0.1959 (0.74761) [−1.2694, +1.6612]	+0.4347 (9.99059) [−19.1468, +20.0163]
Proxy top 3	−0.0823 (0.33441) [−0.7378, +0.5731]	+0.1220 (0.25107) [−0.3701, +0.6141]	−3.2327 (890.53077) [−1748.6731, +1742.2076]
Proxy top 4	−0.0687 (0.51351) [−1.0751, +0.9378]	+0.1220 (0.25938) [−0.3864, +0.6304]	+0.0528 (1.32697) [−2.5480, +2.6537]
Proxy top 5	−0.0094 (0.72909) [−1.4384, +1.4196]	+0.2036 (0.32108) [−0.4257, +0.8329]	+0.0415 (5.00662) [−9.7715, +9.8545]
Proxy top 10	−0.1830 (0.65134) [−1.4596, +1.0937]	+0.2547 (0.41040) [−0.5496, +1.0591]	+0.7833 (3.97540) [−7.0085, +8.5751]
Proxy top 15	−0.2750 (0.71043) [−1.6675, +1.1174]	+0.1694 (0.27099) [−0.3618, +0.7005]	+0.2775 (0.74544) [−1.1836, +1.7385]
Proxy top 20	+0.1843 (0.57232) [−0.9375, +1.3060]	+0.1824 (0.32865) [−0.4618, +0.8265]	+0.5201 (1.42461) [−2.2721, +3.3123]
Proxy top 25	+0.3316 (2.52698) [−4.6213, +5.2845]	+0.5078 (1.71145) [−2.8466, +3.8623]	+3.2718 (51.40849) [−97.4888, +104.0325]
Proxy top 30	+0.3808 (3.13343) [−5.7607, +6.5223]	+0.3328 (0.33683) [−0.3274, +0.9930]	−1.2052 (9.89994) [−20.6091, +18.1987]
Proxy top 35	−0.2672 (0.23840) [−0.7344, +0.2001]	+0.3193 (0.27267) [−0.2151, +0.8537]	−0.2187 (1.20054) [−2.5717, +2.1344]
Proxy top 40	−0.1922 (0.21057) [−0.6049, +0.2205]	+0.3367 (0.28999) [−0.2317, +0.9051]	+0.5845 (2.25366) [−3.8327, +5.0017]
Proxy top 45	+0.1222 (0.11813) [−0.1093, +0.3538]	+0.2804 (0.28593) [−0.2801, +0.8408]	+0.1626 (0.15495) [−0.1411, +0.4663]
Top-45 union (no X)	−0.2408 (0.87537) [−1.9565, +1.4749]	+0.3161 (0.37408) [−0.4170, +1.0493]	+0.3161 (0.37408) [−0.4170, +1.0493]

Table 16: Covariate-omission analogue with closed-form linear specifications, panel (b), continued: interaction, quadratic, and combined linear. Each cell reports the effect estimate (standard error) above its 95% confidence interval; bold cells have upper confidence limits below zero.

Across the 28 constructed scenarios, both DR and DR(sta) retain negative point estimates and 95% confidence intervals below zero. Their estimates remain close to the full-X reference even under the most severe omissions. In contrast, the linear estimators show substantially greater sensitivity to both covariate omission and basis specification, with several scenarios yielding sign changes or very wide confidence intervals. Thus, within the structured omission designs considered here, the flexible minimax estimators exhibit markedly greater empirical stability than the parametric linear alternatives.

Appendix I. Proofs

In all of the proofs, when we assume

$$\|Y\|_\infty, \|\mathbb{H}\|_\infty, \|\mathbb{Q}\|_\infty, \|\mathbb{H}'\|_\infty, \|\mathbb{Q}'\|_\infty, \|\pi\|_\infty, \|\pi/f(a|x,w)\|_2$$

are finite, we denote their upper bounds as $C_Y, C_\mathbb{H}, C_\mathbb{Q}, C_{\mathbb{H}'}, C_{\mathbb{Q}'}, C_\pi, C_w$ respectively.

I.1 Supporting Lemmas

Lemma 42 (Dudley integral)**.**

$$\mathcal{R}(\mathcal{F}) \lesssim \inf_{\tau \geq 0} \left\{ \tau + \int_\tau^{\sup_{f\in\mathcal{F}} \sqrt{\mathbb{E}_n[f^2]}} \sqrt{\frac{\log \mathcal{N}(\beta, \mathcal{F}, \|\cdot\|_n)}{n}} d\beta \right\}.$$

Note $\sup_{f\in\mathcal{F}} \sqrt{\mathbb{P}_n[f^2]}$ is upper bounded by the envelope $\|\mathcal{F}\|_\infty$.

Lemma 43. *Covering number of VC-subgraph classes (van der Vaart, 1998, Lemma 19.15) For a VC class of functions $\mathcal{F}$ with measurable envelope function F and $r \geq 1$, one has for any probability measure Q with $\|F\|_{Q,r} > 0$,*

$$\mathcal{N}(\epsilon \|F\|_{Q,r}, \mathcal{F}, L_r(Q)) \lesssim V(\mathcal{F})(4e)^{V(\mathcal{F})} \left(\frac{2}{\epsilon}\right)^{rV(\mathcal{F})}.$$

Lemma 44 (Theorem 1 in Golowich et al. (2018))**.** *Let $\mathcal{F}_{NN}$ be a class of neural networks with L layers and activation functions that are 1-Lipschitz continuous and positive-homogeneous, that is, $\sigma(az) = a\sigma(z)$ holds with $a > 0$. Further, let A_ℓ be a weight matrix of an ℓ-th layers for $\ell = 1, ..., L$ and assume that $\|A_\ell\|_F \leq M_F(\ell)$ with a bound $M_F(\ell)$. Then, we have $\mathcal{R}_n(\infty; \mathcal{F}_{NN}) \leq \max_x \|x\|_2 (2\sqrt{\log L} + 1) \prod_{\ell=1}^L M_F(\ell)/\sqrt{n}$.*

Lemma 45 (Corollary 14.3 in Wainwright (2019))**.** *Let $\mathcal{N}(\tau; \mathcal{G}^{|\eta}, \|\cdot\|_n)$ denote the τ-covering number of $\mathcal{G}^{|\eta} = \{g \in \mathcal{G} \mid \|g\|_n \leq \delta\}$ under the $\|\cdot\|_n$ norm, where $\mathcal{G}$ is a uniformly bounded function class. Then, the empirical version of critical inequality $\mathcal{R}_n\left(\mathcal{G}^{|\eta}\right) \leq \eta^2/\|\mathcal{G}\|_\infty$ is satisfied for any $\eta > 0$ such that*

$$\frac{1}{\sqrt{n}} \int_{\eta^2/(2\|\mathcal{G}\|_\infty)}^{\eta} \sqrt{\log \mathcal{N}(t, \mathcal{G}^{|\eta}, \|\cdot\|_n)} \mathrm{d}t \leq \frac{\eta^2}{\|\mathcal{G}\|_\infty}.$$

Lemma 46 (Theorem 14.1 in Wainwright (2019)). *Given a star-shaped and b-uniformly bounded function class $\mathcal{G}$, let η_n be any positive solution of the inequality $\mathcal{R}_n(\mathcal{G}^{|\eta}) \leq \eta^2/b$. Then, for any $t \geq \eta_n$, we have*

$$\left|\|g\|_n^2 - \|g\|_2^2\right| \leq \frac{1}{2}\|g\|_2^2 + \frac{1}{2}t^2, \quad \forall g \in \mathcal{G},$$

with probability at least $1 - c_1 \exp\left(-c_2 \frac{nt^2}{b^2}\right)$.

Lemma 47 (Lemma 7 in (Foster and Syrgkanis, 2019)). *Consider a uniformly bounded function class $\mathcal{F}$ and pick an arbitrary $f^* \in \mathcal{F}$. Let η_n be solution to*

$$\mathcal{R}(\mathrm{star}^{|\eta}(\mathcal{F} - f^*)) \leq \frac{\eta^2}{\|\mathcal{F}\|_\infty}, \text{ where } \mathrm{star}^{|\eta}(\mathcal{F} - f^*) \coloneqq \{\alpha\,(f - f^*) : f \in \mathcal{F}, \alpha \in [0,1], \|\alpha\,(f - f^*)\,\|_n \leq \eta\}\,.$$

Moreover, assume that the loss $l(\cdot,\cdot)$ is L-Lipschitz in the first argument. Then, for $\tilde{\eta}_n = \eta_n + \sqrt{c_0 \log(c_1/\delta)/n}$ with some universal constants c_0, c_1, the following holds with probability at least $1 - \delta$,

$$|(\mathbb{E}_n[l(f(x),z)] - \mathbb{E}_n[l(f^*(x),z)]) - (\mathbb{E}[l(f(x),z)] - \mathbb{E}[l(f^*(x),z)])| \lesssim L\tilde{\eta}_n(\|f - f^*\|_2 + \tilde{\eta}_n).$$

Lemma 48 (Neural network; Lemma 21 in Nakada and Imaizumi (2020)). *Let $\mathcal{F}_{NN}$ be a set of functions by a neural network with L layers, Ω weights in $[-B, B]$, and $1-$Lipschitz continuous activation function. Then, for $\tau \in (0,1]$, we have*

$$\log \mathcal{N}(\tau, \mathcal{F}_{NN}, \|\cdot\|_\infty) \leq \Omega \log\left(\frac{2LB^L(\Omega+1)^L}{\tau}\right).$$

Lemma 49. *Suppose that Assumption 1 holds.*

1. *If $\mathbb{Q}_0 \neq \emptyset$ and $\mathbb{H}_0^{\mathrm{obs}} \neq \emptyset$, then for any $h_0 \in \mathbb{H}_0^{\mathrm{obs}}$ and $q \in L_2(Z,A,X)$,*

$$\mathbb{E}\left[\tilde{\phi}_{\mathrm{IPW}}\,(O;q)\right] - J = \mathbb{E}\left[h_0(W,A,X)\mathbb{E}\left[\pi(A \mid X)\,(q(Z,A,X) - 1/f(A \mid W,X)) \mid W,A,X\right]\right].$$

2. *If $\mathbb{H}_0 \neq \emptyset$ and $\mathbb{Q}_0^{\mathrm{obs}} \neq \emptyset$, then for any $q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$ and $h \in L_2(W,A,X)$,*

$$\mathbb{E}\left[\tilde{\phi}_{\mathrm{REG}}\,(O;h)\right] - J = \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)\mathbb{E}\left[h(W,A,X) - Y \mid Z,A,X\right]\right].$$

Proof for Lemma 49. Before proving the conclusions, note that for any $h \in L_2(W,A,X)$ and $q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$,

$$\begin{aligned}
\mathbb{E}\left[q_0(Z,A,X)\pi(A \mid X)h(W,A,X)\right] &= \mathbb{E}\left[\mathbb{E}\left[\mathbb{E}\left[q_0(Z,A,X)\pi(A|X) \mid W,A,X\right]h(W,A,X)\right]\right] \\
&= \mathbb{E}\left[\frac{\pi(A \mid X)}{f(A \mid W,X)}h(W,A,X)\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\frac{\pi(A \mid X)}{f(A \mid W,X)}h(W,A,X) \mid W,X\right]\right] \\
&= \mathbb{E}\left[\int \frac{\pi(a \mid X)}{f(a \mid W,X)}h(W,a,X)f(a \mid W,X)\mathrm{d}\mu(a)\right]
\end{aligned}$$

$$= \mathbb{E}\left[(\mathcal{T}h)(W,X)\right]. \tag{49}$$

Here the second equality follows because q_0 satisfies Equation (9).

IPW. First, by taking some element h_0 in $\mathbb{H}_0^{\text{obs}}$, we have

$$\begin{aligned}
\mathbb{E}\left[\tilde{\phi}_{\text{IPW}}(O;q)\right] &= \mathbb{E}\left[\pi(A \mid X)q(Z,A,X)Y\right] \\
&= \mathbb{E}\left[\pi(A \mid X)q(Z,A,X)\mathbb{E}\left[Y \mid Z,A,X\right]\right] \\
&= \mathbb{E}\left[\pi(A \mid X)q(Z,A,X)h_0(W,A,X)\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\pi(A \mid X)q(Z,A,X) \mid W,A,X\right]h_0(W,A,X)\right].
\end{aligned}$$

Moreover, by taking some element $q_0' \in \mathbb{Q}_0$,

$$\begin{aligned}
J &= \mathbb{E}\left[\pi(A \mid X)q_0'(Z,A,X)Y\right] && \text{(Lemma 2)} \\
&= \mathbb{E}\left[\pi(A \mid X)q_0'(Z,A,X)\mathbb{E}[Y|A,X,W]\right] && (Y \perp Z \mid A,X,W) \\
&= \mathbb{E}\left[\mathbb{E}\left[\pi(A \mid X)q_0'(Z,A,X) \mid W,A,X\right]h_0(W,A,X)\right] \\
&= \mathbb{E}\left[\frac{\pi(A \mid X)}{f(A \mid W,X)}h_0(W,A,X)\right].
\end{aligned}$$

It follows that

$$\mathbb{E}\left[\tilde{\phi}_{\text{IPW}}(O;q)\right] - J = \mathbb{E}\left[\{\mathbb{E}\left[\pi(A \mid X)q(Z,A,X) \mid W,A,X\right] - \frac{\pi(A \mid X)}{f(A \mid W,X)}\}h_0(W,A,X)\right].$$

REG. For any $q_0 \in \mathbb{Q}_0^{\text{obs}}$, note that

$$\begin{aligned}
\mathbb{E}\left[\tilde{\phi}_{\text{REG}}(O;h)\right] &= \mathbb{E}\left[(\mathcal{T}h)(W,X)\right] \\
&= \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)h(W,A,X)\right] && \text{(Equation (49))} \\
&= \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)\mathbb{E}\left[h(W,A,X) \mid Z,A,X\right]\right].
\end{aligned}$$

Moreover, by taking some element $h_0' \in \mathbb{H}_0$,

$$\begin{aligned}
J &= \mathbb{E}\left[(\mathcal{T}h_0')(W,X)\right] && \text{(Lemma 2)} \\
&= \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)h_0'(W,A,X)\right] && \text{(Equation (49))} \\
&= \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)\mathbb{E}\left[h_0'(W,A,X) \mid Z,A,X\right]\right] \\
&= \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)\mathbb{E}\left[Y \mid Z,A,X\right]\right].
\end{aligned}$$

It follows that

$$\mathbb{E}\left[\tilde{\phi}_{\text{REG}}\left(O;h\right)\right] - J = \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)\mathbb{E}\left[h(W,A,X) - Y \mid Z,A,X\right]\right].$$

■

I.2 Proofs for Section 2

Proof of Lemma 1. We first prove the identification via ϕ_{REG}:

$$\begin{aligned}
J &= \mathbb{E}\left\{\int Y(a)\pi(a|X)\mathrm{d}\mu(a)\right\} \\
&= \mathbb{E}\left\{\int \mathbb{E}[Y(a)|U,X]\pi(a|X)\mathrm{d}\mu(a)\right\} && \text{(Tower property)} \\
&= \mathbb{E}\left\{\int \mathbb{E}[Y(a)|U,X,A=a]\pi(a|X)\mathrm{d}\mu(a)\right\} && (Y(a)\perp A|U,X) \\
&= \mathbb{E}\left\{\int \mathbb{E}[Y|U,X,A=a]\pi(a|X)\mathrm{d}\mu(a)\right\} = \mathbb{E}\left[\phi_{\mathrm{REG}}\left(Y,A,U,X;k_0\right)\right]. && \text{(Consistency)}
\end{aligned}$$

We next prove the identification via ϕ_{IPW}:

$$\begin{aligned}
J &= \mathbb{E}\left\{\int \mathbb{E}[Y|U,X,A=a]\pi(a|X)\mathrm{d}\mu(a)\right\} \\
&= \mathbb{E}\left\{\int \mathbb{E}[Y|U,X,A=a]\{\pi(a|X)/f(a|X,U)\}\times f(a|X,U)\mathrm{d}\mu(a)\right\} \\
&= \mathbb{E}\left\{\mathbb{E}[Y|U,X,A]\pi(A|X)/f(A|X,U)\right\} \\
&= \mathbb{E}\left\{Y\pi(A|X)/f(A|X,U)\right\} = \mathbb{E}\left[\phi_{\mathrm{IPW}}\left(Y,A,U,X;f\right)\right].
\end{aligned}$$

Finally, the identification via ϕ_{DR} follows from the fact that

$$\mathbb{E}\left[\phi_{\mathrm{DR}}\left(Y,A,U,X;k_0,f\right)\right] = \mathbb{E}\left[\phi_{\mathrm{REG}}\left(Y,A,U,X;k_0\right)\right] = J.$$

■

Proof for Lemma 2. We derive the equations in Lemma 2 one by one.

First,

$$\begin{aligned}
\mathbb{E}\left[\tilde{\phi}_{\mathrm{IPW}}(O;q_0)\right] &= \mathbb{E}\left[\mathbb{E}\left[\pi(A\mid X)q_0(Z,A,X)Y\mid A,U,X\right]\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\pi(A\mid X)q_0(Z,A,X)\mid A,U,X\right]\mathbb{E}\left[Y\mid A,U,X\right]\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\pi(A\mid X)/f(A\mid U,X)\mid A,U,X\right]\mathbb{E}\left[Y\mid A,U,X\right]\right] \\
&= \mathbb{E}\left[\pi(A\mid X)Y/f(A\mid U,X)\right] \\
&= \mathbb{E}\left[\int \pi\left(a\mid X\right)\mathbb{E}\left[Y(a)\mid A=a,U,X\right]\mathrm{d}\mu(a)\right] \\
&= \mathbb{E}\left[\int \pi\left(a\mid X\right)\mathbb{E}\left[Y(a)\mid U,X\right]\mathrm{d}\mu(a)\right] = J.
\end{aligned}$$

Here the second equality follows from $Y\perp Z\mid A,U,X$ and the third equality follows from the definition of q_0 according to Equation (3).

Second,

$$\mathbb{E}\left[\tilde{\phi}_{\mathrm{REG}}(O;h_0)\right] = \mathbb{E}\left[\int \pi(a\mid X)h_0(W,a,X)\mathrm{d}\mu(a)\right]$$

$$
\begin{aligned}
&= \mathbb{E}\left[\int \pi(a \mid X)\mathbb{E}\left[h_0(W,a,X) \mid A=a,U,X\right]\mathrm{d}\mu(a)\right] \\
&= \mathbb{E}\left[\int \pi(a \mid X)\mathbb{E}\left[Y \mid A=a,U,X\right]\mathrm{d}\mu(a)\right] \\
&= J.
\end{aligned}
$$

Here the second equality follows from $W \perp A \mid U, X$ and the third equality follows from the definition of h_0 according to Equation (2).

Finally,

$$
\begin{aligned}
\mathbb{E}\left[\tilde{\phi}_{\mathrm{DR}}(O;h_0,q_0)\right] &= \mathbb{E}\left[\tilde{\phi}_{\mathrm{REG}}(O;h_0,q_0)\right] + \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)\left(Y-h_0(W,A,X)\right)\right] \\
&= J + \mathbb{E}\left[\mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X)\left(Y-h_0(W,A,X)\right) \mid A,U,X\right]\right] \\
&= J + \mathbb{E}\left[\mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X) \mid A,U,X\right]\mathbb{E}\left[Y-h_0(W,A,X) \mid A,U,X\right]\right] \\
&= J.
\end{aligned}
$$

Here the third equality follows from that $Z \perp Y \mid A, U, X$ and the last equality follows from the definition of h_0 according to Equation (2). ■

Proof for Lemma 3. For any $h_0 \in \mathbb{H}_0$, we have that

$$
\mathbb{E}\left[Y-h_0(W,A,X) \mid Z,A,U,X\right] = \mathbb{E}\left[Y-h_0(W,A,X) \mid A,U,X\right] = 0,
$$

where the first equality holds because $(W,Y) \perp Z \mid A,U,X$. Therefore,

$$
\mathbb{E}[Y-h_0(W,A,X) \mid Z,A,X] = \mathbb{E}\left[\mathbb{E}\left[Y-h_0(W,A,X) \mid Z,A,U,X\right] \mid Z,A,X\right] = 0.
$$

For any $q_0 \in \mathbb{Q}_0$, we have that

$$
\mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X) \mid W,A,U,X\right] = \mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X) \mid A,U,X\right] = \frac{\pi(A \mid X)}{f(A \mid U,X)},
$$

where the first equality follows from $Z \perp W \mid A,U,X$. Therefore,

$$
\begin{aligned}
\mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X) \mid W,A,X\right] &= \mathbb{E}\left[\mathbb{E}\left[\pi(A \mid X)q_0(Z,A,X) \mid W,A,U,X\right] \mid W,A,X\right] \\
&= \mathbb{E}\left[\frac{\pi(A \mid X)}{f(A \mid U,X)} \mid W,A,X\right].
\end{aligned}
$$

Then the conclusion follows from the fact that

$$
\begin{aligned}
\mathbb{E}\left[\frac{1}{f(A \mid U,X)} \mid W,A,X\right] &= \int \frac{1}{f(A \mid u,X)} f(u \mid W,A,X)\mathrm{d}(u) \\
&= \int \frac{f(W,A \mid u,X)f(u,X)}{f(A \mid u,X)f(W,A,X)}\mathrm{d}\mu(u) \\
&= \int \frac{f(W \mid u,X)f(u,X)}{f(W,A,X)}\mathrm{d}\mu(u) \\
&= \frac{1}{f(A \mid W,X)}.
\end{aligned}
$$

Here the third equality follows from the fact that $W \perp A \mid U, X$.

In summary, the condition that $(W,Y) \perp Z \mid A,U,X$ ensures Equation (8) holds, while the condition that $(Z,A) \perp W \mid U,X$ ensures Equation (9) holds. ■

Proof for Theorem 4. Obviously, the conclusions in statements 1 and 2 directly follow from Lemma 49. We only need to prove the statement 3.

When the condition in statement (1) holds, for any $h_0 \in \mathbb{H}_0^{\text{obs}}$ and $q_0 \in \mathbb{Q}_0^{\text{obs}}$,

$$\begin{aligned}\mathbb{E}\left[\tilde{\phi}_{\text{DR}}(O; h_0, q_0)\right] &= \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X)\left(Y - h_0(W, A, X)\right)\right] + \mathbb{E}\left[(\mathcal{T} h_0)(X, W)\right] \\ &= \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X) \mathbb{E}\left[Y - h_0(W, A, X) \mid Z, A, X\right]\right] + \mathbb{E}\left[\tilde{\phi}_{\text{REG}}(O; h_0)\right] \\ &= \mathbb{E}\left[\tilde{\phi}_{\text{REG}}(O; h_0)\right] = J.\end{aligned}$$

When the condition in statement (2) holds, for any $h_0 \in \mathbb{H}_0^{\text{obs}}$ and $q_0 \in \mathbb{Q}_0^{\text{obs}}$,

$$\begin{aligned}\mathbb{E}\left[\tilde{\phi}_{\text{DR}}(O; h_0, q_0)\right] &= \mathbb{E}\left[\tilde{\phi}_{\text{IPW}}(O; q_0)\right] + \mathbb{E}\left[(\mathcal{T} h_0)(X, W)\right] - \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X) h_0(W, A, X)\right] \\ &= \mathbb{E}\left[\tilde{\phi}_{\text{IPW}}(O; q_0)\right] = J.\end{aligned}$$

Here the second equality follows from Equation (49). ■

I.3 Proofs for Section 3

Lemma 50 (Bridge Functions in Linear Models)**.** *Consider the linear models in Example 4, where $\alpha_Z \in \mathbb{R}^{p_z \times p_u}, \alpha_W \in \mathbb{R}^{p_w \times p_u}$ both have full column rank. Then bridge functions h_0 and q_0 have the following forms:*

$$\begin{aligned}h_0(W, A, X) &= \left(\theta_W + \omega_Y\right)^\top W + \left(\beta_Y - \theta_W^\top \beta_W\right) X + \gamma_Y A, \quad \forall \theta_W \text{ s.t. } \theta_W^\top \alpha_W = \alpha_Y^\top, \\ q_0(Z, A, X) &= \theta_Z^\top Z + \left(\overline{\beta}_A - \underline{\beta}_A - \theta_Z^\top \beta_Z\right) X - \theta_Z^\top \gamma_Z A, \quad \forall \theta_Z \text{ s.t. } \theta_Z^\top \alpha_Z = \overline{\alpha}_A^\top - \underline{\alpha}_A^\top.\end{aligned}$$

Proof for Lemma 50. Here we show the details of deriving bridge functions in Example 4.

We first derive the function $h_0(W, A, X)$. If $\theta_w^\top \alpha_W = \alpha_Y^\top$,then

$$\theta_W^\top W = \alpha_Y^\top U + \theta_W^\top \beta_W X + \theta_W^\top \epsilon_W \implies \alpha_Y^\top U = \theta_W^\top W - \theta_W^\top \beta_W X - \theta_W^\top \epsilon_W.$$

Therefore,

$$\begin{aligned}Y &= \left(\theta_W + \omega_Y\right)^\top W + \left(\beta_Y^\top - \theta_W^\top \beta_W\right) X + \gamma_Y A + \epsilon_Y - \theta_W^\top \epsilon_W \\ &= h_w(W, A, X) + \epsilon_Y - \theta_W^\top \epsilon_W.\end{aligned}$$

It follows from the independence of ϵ_Y, ϵ_W with A, U, X that

$$\mathbb{E}\left[Y \mid A, U, X\right] = \mathbb{E}\left[h_0(W, A, X) \mid A, U, X\right] + \mathbb{E}\left[\epsilon_Y - \theta_W^\top \epsilon_W\right] = \mathbb{E}\left[h_0(W, A, X) \mid A, U, X\right].$$

Now we derive $q_0(Z, A, X)$. First note that

$$1/f(A \mid U, X) = \left(\left(\overline{\alpha}_A - \underline{\alpha}_A\right) U + \left(\overline{\beta}_A - \underline{\beta}_A\right) X\right).$$

Because $\theta_Z^\top \alpha_Z = \overline{\alpha}_A^\top - \underline{\alpha}_A^\top$,

$$\theta_Z^\top Z = \left(\overline{\alpha}_A - \underline{\alpha}_A\right)^\top U + \theta_Z^\top \beta_Z X + \theta_Z^\top \gamma_Z A + \theta_Z^\top \epsilon_Z,$$

which means that

$$\begin{aligned} q_0(Z,A,X) &= \theta_Z^\top Z + \left(\overline{\beta}_A - \underline{\beta}_A - \theta_Z^\top \beta_Z\right) X - \theta_Z^\top \gamma_Z A \\ &= \left(\overline{\alpha}_A - \underline{\alpha}_A\right)^\top U + \left(\overline{\beta}_A - \underline{\beta}_A\right)^\top X + \theta_Z^\top \epsilon_Z. \end{aligned}$$

Therefore,

$$\mathbb{E}\left[q_0(Z,A,X) \mid A,U,X\right] = \left(\overline{\alpha}_A - \underline{\alpha}_A\right)^\top U + \left(\overline{\beta}_A - \underline{\beta}_A\right)^\top X + \mathbb{E}\left[\theta_Z^\top \epsilon_Z\right] = 1/f(A \mid U,X).$$

■

I.4 Proofs for Section 5

We first prove the error bounds and then derive rates for common function classes.

I.4.1 Proofs for Section 5.1

Proof of Theorem 7. We start by proving Equation (27). We fix functions $h_0 \in \mathbb{H}_0^{\text{obs}}, q_0 \in \mathbb{Q}_0^{\text{obs}}$. Recall that

$$\hat{q} \in \arg\min_{q\in\mathbb{Q}} \max_{h\in\mathbb{H}'} \left(\mathbb{E}_n[\pi(A|X)q(Z,A,X)h(W,A,X) - (\mathcal{T}h)(W,X)]\right)^2 = \arg\min_{q\in\mathbb{Q}} \max_{h\in\mathbb{H}'} \left|\mathbb{E}_n f_q\left(q,h\right)\right|,$$

where

$$f_q(q,h) \coloneqq -\pi q h + \mathcal{T}h.$$

First Step. We first define $J_{\text{IPW}} \coloneqq \mathbb{E}[\hat{q}\pi Y]$. Obviously,

$$|J_{\text{IPW}} - \hat{J}_{\text{IPW}}| \le \sup_{q\in\mathbb{Q}} |(\mathbb{E} - \mathbb{E}_n)[q\pi Y]|.$$

In addition, we have

$$\begin{aligned} |J_{\text{IPW}} - J| &\le |\mathbb{E}[(\hat{q}-q_0)\pi Y]| = |\mathbb{E}[\mathbb{E}[(\hat{q}-q_0)\pi Y|A,X,U]]| \\ &= |\mathbb{E}[(\hat{q}-q_0)\pi h_0]| && (Z \perp Y \mid A,X,U) \\ &\le \sup_{h\in\mathbb{H}'} |\mathbb{E}[(\hat{q}-q_0)\pi h]| + \inf_{h\in\mathbb{H}'} |\mathbb{E}[(\hat{q}-q_0)\pi(h_0-h)]| \\ &\le |\mathbb{E}\left[f_q(\hat{q},\overline{h})\right]| + \sup_{q\in\mathbb{Q}} \inf_{h\in\mathbb{H}'} |\mathbb{E}[(q-q_0)\pi(h_0-h)]|, && \overline{h} \in \arg\max_{h\in\mathbb{H}'} |\mathbb{E}[(\hat{q}-q_0)\pi h]| \\ &\le |\mathbb{E}\left[f_q(\hat{q},\overline{h})\right]| + \inf_{h\in\mathbb{H}'} \sup_{q\in\mathbb{Q}} |\mathbb{E}[(q-q_0)\pi(h_0-h)]|. \end{aligned}$$

Second Step. Define $q' = \arg\min_{q\in\mathbb{Q}} \sup_{h\in\mathbb{H}'} |\mathbb{E}[f_q(q,h)]|$. Then

$$\begin{aligned} |\mathbb{E}\left[f_q(\hat{q},\overline{h})\right]| &\le |\mathbb{E}_n\left[-\pi\hat{q}\overline{h} + \mathcal{T}\overline{h}\right]| + |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi\hat{q}\overline{h} + \mathcal{T}\overline{h}\right]| \\ &\le \sup_{h\in\mathbb{H}'} |\mathbb{E}_n\left[-\pi\hat{q}h + \mathcal{T}h\right]| + |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi\hat{q}\overline{h} + \mathcal{T}\overline{h}\right]| \\ &\le \sup_{h\in\mathbb{H}'} |\mathbb{E}_n\left[-\pi q'h + \mathcal{T}h\right]| + |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi\hat{q}\overline{h} + \mathcal{T}\overline{h}\right]| && (\text{Definition of } \hat{q}) \\ &\le \sup_{h\in\mathbb{H}'} |\mathbb{E}\left[-\pi q'h + \mathcal{T}h\right]| + \sup_{h\in\mathbb{H}'} |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi q'h + \mathcal{T}h\right]| + |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi\hat{q}\overline{h} + \mathcal{T}\overline{h}\right]| \end{aligned}$$

$$
\begin{aligned}
&\leq \sup_{h\in\mathbb{H}'} |\mathbb{E}\left[-\pi q' h + \mathcal{T}h\right]| + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi q h + \mathcal{T}h\right]| \\
&= \inf_{q\in\mathbb{Q}} \sup_{h\in\mathbb{H}'} |\mathbb{E}\left[-\pi q h + \mathcal{T}h\right]| + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi q h + \mathcal{T}h\right]| \qquad \text{(Definition of } q'\text{)} \\
&= \inf_{q\in\mathbb{Q}} \sup_{h\in\mathbb{H}'} |\mathbb{E}\left[\pi\left(q_0 - q\right)h\right]| + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi q h + \mathcal{T}h\right]|
\end{aligned}
$$

Third Step. Combining all results,

$$
\begin{aligned}
|J - \hat{J}_{\mathrm{IPW}}| &\leq |J_{\mathrm{IPW}} - \hat{J}_{\mathrm{IPW}}| + |J_{\mathrm{IPW}} - J| \\
&\leq \sup_{q\in\mathbb{Q}} |(\mathbb{E} - \mathbb{E}_n)[q\pi Y]| + |\mathbb{E}[f_q(\hat{q}, \bar{h})]| + \inf_{h\in\mathbb{H}'} \sup_{q\in\mathbb{Q}} |\mathbb{E}[(q - q_0)\pi(h_0 - h)]| \\
&\leq \sup_{q\in\mathbb{Q}} |(\mathbb{E} - \mathbb{E}_n)[q\pi Y]| + \inf_{q\in\mathbb{Q}} \sup_{h\in\mathbb{H}'} |\mathbb{E}\left[\pi\left(q_0 - q\right)h\right]| \\
&\quad + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |\left(\mathbb{E}_n - \mathbb{E}\right)\left[-\pi q h + \mathcal{T}h\right]| + \inf_{h\in\mathbb{H}'} \sup_{q\in\mathbb{Q}} |\mathbb{E}[(q - q_0)\pi(h_0 - h)]|.
\end{aligned}
$$

Next, we turn to proving Equation (28). We fix functions $h_0 \in \mathbb{H}_0^{\mathrm{obs}}, q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$. Recall that

$$
\hat{h} \in \arg\min_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} \left(\mathbb{E}_n[q(Z, A, X)\left(h(W, A, X) - Y\right)]\right)^2 = \arg\min_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} \left|\mathbb{E}_n f_h\left(q, h\right)\right|,
$$

where

$$
f_h(q, h) \coloneqq q\left(Y - h\right).
$$

First step We first define $J_{\mathrm{REG}} \coloneqq \mathbb{E}[\mathcal{T}h]$. Then,

$$
|J_{\mathrm{REG}} - \hat{J}_{\mathrm{REG}}| \leq \sup_{h\in\mathbb{H}} (\mathbb{E} - \mathbb{E}_n)|[\mathcal{T}h]|.
$$

In addition,

$$
\begin{aligned}
|J_{\mathrm{REG}} - J| &\leq |\mathbb{E}[\mathcal{T}(\hat{h} - h_0)| = |\mathbb{E}[q_0\pi(\hat{h} - h_0)| \\
&\leq \sup_{q\in\mathbb{Q}'} |\mathbb{E}[q(\hat{h} - h_0)| + \inf_{q\in\mathbb{Q}'} |\mathbb{E}[(q_0\pi - q)(\hat{h} - h_0)| \\
&\leq |\mathbb{E}[\bar{q}(\hat{h} - h_0)| + \sup_{h\in\mathbb{H}} \inf_{q\in\mathbb{Q}'} |\mathbb{E}[(q_0\pi - q)(h - h_0)|, \qquad \bar{q} \in \arg\max_{q\in\mathbb{Q}'} |\mathbb{E}[q(\hat{h} - h_0)| \\
&\leq |\mathbb{E}[\bar{q}(\hat{h} - h_0)| + \inf_{q\in\mathbb{Q}'} \sup_{h\in\mathbb{H}} |\mathbb{E}[(q_0\pi - q)(h - h_0)|.
\end{aligned}
$$

Second step. Define $h' = \arg\min_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} |\mathbb{E}[f_h(q, h)]|$. Then,

$$
\begin{aligned}
&|\mathbb{E}[\bar{q}(\hat{h} - h_0)| = |\mathbb{E}[\bar{q}(\hat{h} - Y)| \\
&\leq |(\mathbb{E} - \mathbb{E}_n)[\bar{q}(\hat{h} - Y)| + |\mathbb{E}_n[\bar{q}(\hat{h} - Y)]| \\
&\leq |(\mathbb{E} - \mathbb{E}_n)[\bar{q}(\hat{h} - Y)| + \sup_{q\in\mathbb{Q}'} |\mathbb{E}_n[q(\hat{h} - Y)]| \\
&\leq |(\mathbb{E} - \mathbb{E}_n)[\bar{q}(\hat{h} - Y)| + \sup_{q\in\mathbb{Q}'} |\mathbb{E}_n[q(h' - Y)]| \qquad \text{(Definition of } \hat{h}\text{)} \\
&\leq |(\mathbb{E} - \mathbb{E}_n)[\bar{q}(\hat{h} - Y)| + \sup_{q\in\mathbb{Q}'} |(\mathbb{E}_n - \mathbb{E})[q(h' - Y)]| + \sup_{q\in\mathbb{Q}'} |\mathbb{E}[q(h' - Y)]|
\end{aligned}
$$

$$\leq 2 \sup_{q\in\mathbb{Q}',h\in\mathbb{H}} |(\mathbb{E}_n - \mathbb{E})[q(Y-h)]| + \inf_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} |\mathbb{E}[q(h_0-h)]|. \qquad \text{(Definition of } h')$$

Third step.
Combining all results,

$$\begin{aligned}
|J - \hat{J}_{\mathrm{REG}}| &\leq |J_{\mathrm{REG}} - \hat{J}_{\mathrm{REG}}| + |J_{\mathrm{REG}} - J| \\
&\leq \sup_{h\in\mathbb{H}} |(\mathbb{E}-\mathbb{E}_n)[\mathcal{T}h]| + |\mathbb{E}[\bar{q}(\hat{h}-h_0)| + \inf_{q\in\mathbb{Q}'} \sup_{h\in\mathbb{H}} |\mathbb{E}[(q_0\pi - q)(h-h_0)| \\
&\leq \sup_{h\in\mathbb{H}} |(\mathbb{E}-\mathbb{E}_n)[\mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q}',h\in\mathbb{H}} |(\mathbb{E}-\mathbb{E}_n)[q(Y-h)]| \\
&+ \inf_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} |\mathbb{E}[\{h_0-h\}q]| + \inf_{q\in\mathbb{Q}'} \sup_{h\in\mathbb{H}} |\mathbb{E}[(q_0\pi - q)(h-h_0)|.
\end{aligned}$$

■

Proof of Theorem 9. We fix functions $h_0 \in \mathbb{H}_0^{\mathrm{obs}}, q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$.

First Statement. First, we prove

$$\begin{aligned}
|\hat{J}_{\mathrm{DR}} - J| \leq \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E}-\mathbb{E}_n)[q\pi Y - q\pi h + \mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}'} |(\mathbb{E}-\mathbb{E}_n)[-q\pi h + \mathcal{T}h]| + \\
+ \inf_{q\in\mathbb{Q}} \sup_{h\in\mathbb{H}'} |\mathbb{E}[(q_0-q)\pi h]| + \inf_{h\in\mathbb{H}'} \sup_{h'\in\mathbb{H}} \sup_{q\in\mathbb{Q}} |\mathbb{E}[(q-q_0)\pi(h_0-h'-h)]|.
\end{aligned}$$

We define

$$J_{\mathrm{DR}} \coloneqq \mathbb{E}[\hat{q}\pi\{Y-\hat{h}\} + \mathcal{T}\hat{h}].$$

Then,

$$|J_{\mathrm{DR}} - \hat{J}_{\mathrm{DR}}| \leq \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E}-\mathbb{E}_n)[q\pi\{Y-h\} + \mathcal{T}h]|.$$

In addition,

$$\begin{aligned}
|J_{\mathrm{DR}} - J| &\leq |\mathbb{E}[(\hat{q}-q_0)\pi(h_0-\hat{h})]| \\
&\leq \sup_{h\in\mathbb{H}'} |\mathbb{E}[(\hat{q}-q_0)\pi h]| + \inf_{h\in\mathbb{H}'} |\mathbb{E}[(\hat{q}-q_0)\pi(h_0-\hat{h}-h)]| \\
&= |\mathbb{E}[(\hat{q}-q_0)\pi\bar{h}]| + \sup_{h'\in\mathbb{H}} \sup_{q\in\mathbb{Q}} \inf_{h\in\mathbb{H}'} |\mathbb{E}[(q-q_0)\pi(h_0-h'-h)]|, \qquad \bar{h} \coloneqq \arg\max_{h\in\mathbb{H}'} |\mathbb{E}[(\hat{q}-q_0)\pi h]| \\
&= |\mathbb{E}[(\hat{q}-q_0)\pi\bar{h}]| + \inf_{h\in\mathbb{H}'} \sup_{h'\in\mathbb{H},q\in\mathbb{Q}} |\mathbb{E}[(q-q_0)\pi(h_0-h'-h)]| \\
&= |\mathbb{E}[f_q(\hat{q},\bar{h})]| + \inf_{h\in\mathbb{H}'} \sup_{h'\in\mathbb{H},q\in\mathbb{Q}} |\mathbb{E}[(q-q_0)\pi(h_0-h'-h)]|
\end{aligned}$$

The rest of the proof is the same as that of Equation (27) in Theorem 7.

Second Statement. Second, we prove

$$\begin{aligned}
|\hat{J}_{\mathrm{DR}} - J| \leq \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E}-\mathbb{E}_n)[q\pi Y - q\pi h + \mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q}',h\in\mathbb{H}} |(\mathbb{E}-\mathbb{E}_n)[q(Y-h)]| \\
+ \inf_{h\in\mathbb{H}} \sup_{q\in\mathbb{Q}'} |\mathbb{E}[\{h_0-h\}q]| + \inf_{q\in\mathbb{Q}'} \sup_{h\in\mathbb{H},q'\in\mathbb{Q}} |\mathbb{E}[(q_0\pi - q' - q)(h-h_0)|.
\end{aligned}$$

We define

$$J_{\mathrm{DR}} \coloneqq \mathbb{E}[\hat{q}\pi\{Y - \hat{h}\} + \mathcal{T}\hat{h}].$$

Then,

$$|J_{\mathrm{DR}} - \hat{J}_{\mathrm{DR}}| \leq \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E} - \mathbb{E}_n)[q\pi\{Y - h\} + \mathcal{T}h]|.$$

In addition,

$$\begin{aligned}
|J_{\mathrm{DR}} - J| &\leq |\mathbb{E}[(\hat{q} - q_0)\pi(h_0 - \hat{h})]| \\
&\leq \sup_{q\in\mathbb{Q}'} |\mathbb{E}[q(\hat{h} - h_0)| + \inf_{q\in\mathbb{Q}'} |\mathbb{E}[(q_0\pi - \hat{q}\pi - q)\hat{h} - h_0)| \\
&\leq |\mathbb{E}[\bar{q}(\hat{h} - h_0)]| + \sup_{h\in\mathbb{H},q'\in\mathbb{Q}} \inf_{q\in\mathbb{Q}'} |\mathbb{E}[(q_0\pi - q'\pi - q)(h - h_0)|, \qquad \bar{q} \coloneqq \arg\max_{q\in\mathbb{Q}'} |\mathbb{E}[q(\hat{h} - h_0)| \\
&\leq |\mathbb{E}[\bar{q}(\hat{h} - h_0)]| + \inf_{q\in\mathbb{Q}'} \sup_{h\in\mathbb{H},q'\in\mathbb{Q}} |\mathbb{E}[(q_0\pi - q'\pi - q)(h - h_0)| \\
&= |\mathbb{E}[f_h(\bar{q},\hat{h})]| + \inf_{q\in\mathbb{Q}'} \sup_{h\in\mathbb{H},q'\in\mathbb{Q}} |\mathbb{E}[(q_0\pi - q'\pi - q)(h - h_0)|
\end{aligned}$$

The rest of the proof is the same as that of Equation (28) in Theorem 7. ∎

I.4.2 Proofs for Section 5.2

Proof of Corollary 11. Given that $\mathbb{H}' = \mathbb{H}$ and $\mathbb{Q}' = \pi\mathbb{Q}$, Corollary 8 states that

$$\begin{aligned}
|\hat{J}_{\mathrm{IPW}} - J| &\leq \sup_{q\in\mathbb{Q}} |(\mathbb{E} - \mathbb{E}_n)[q\pi Y]| + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E} - \mathbb{E}_n)[-q\pi h + \mathcal{T}h]|, \\
|\hat{J}_{\mathrm{REG}} - J| &\leq \sup_{h\in\mathbb{H}} |(\mathbb{E} - \mathbb{E}_n)[\mathcal{T}h]| + 2 \sup_{q\in\mathbb{Q},h\in\mathbb{H}} |(\mathbb{E} - \mathbb{E}_n)[q\pi(Y - h)]|.
\end{aligned}$$

We define three function classes for the analysis:

$$\begin{aligned}
\mathbb{A}_1 &= \{(y,z,a,x) \mapsto q(z,a,x)\pi\left(a \mid x\right) y : q \in \mathbb{Q}\}, \\
\mathbb{A}_2 &= \{(y,w,z,a,x) \mapsto q(z,a,x)\pi\left(a \mid x\right) h(w,a,x) : q \in \mathbb{Q}, h \in \mathbb{H}\}, \\
\mathbb{A}_3 &= \{(w,a,x) \mapsto (\mathcal{T}h)(w,x) \in \mathbb{H}\}.
\end{aligned}$$

Then, from Theorem 4.10 in Wainwright (2019) and Corollary 8,there exists a universal positive constant c such that

$$\max\left\{|\hat{J}_{\mathrm{REG}} - J|, |\hat{J}_{\mathrm{IPW}} - J|\right\} \leq c\{\mathcal{R}(\infty;\mathbb{A}_1) + \mathcal{R}(\infty;\mathbb{A}_2) + \mathcal{R}(\infty;\mathbb{A}_3) + \sqrt{\log(1/\delta)/n}\}.$$

We first note that

$$\mathcal{R}(\infty;\mathbb{A}_2) \leq \|\pi(A \mid X)\|_\infty \mathcal{R}(\infty;\mathbb{H}\mathbb{Q})$$

and

$$\begin{aligned}
\mathcal{R}(\infty;\mathbb{H}\mathbb{Q}) &\leq \mathcal{R}(\infty;0.25\left\{(\mathbb{H} + \mathbb{Q})^2 - (\mathbb{H} - \mathbb{Q})^2\right\}) \\
&\leq 0.25\left\{\mathcal{R}(\infty;(\mathbb{H} - \mathbb{Q})^2) + \mathcal{R}(\infty;(\mathbb{H} + \mathbb{Q})^2)\right\}
\end{aligned}$$

$$\leq 0.5(\|\mathbb{H}\|_\infty + \|\mathbb{Q}\|_\infty)\{\mathcal{R}(\infty;(\mathbb{H}-\mathbb{Q})) + \mathcal{R}(\infty;(\mathbb{H}+\mathbb{Q}))\}$$
(Contraction property (Mendelson, 2002))
$$\leq (\|\mathbb{H}\|_\infty + \|\mathbb{Q}\|_\infty)\,(\mathcal{R}(\infty;\mathbb{H}) + \mathcal{R}(\infty;\mathbb{Q}))$$

Thus

$$\mathcal{R}(\infty;\mathbb{A}_2) \leq \|\pi(A \mid X)\|_\infty\,(\|\mathbb{H}\|_\infty + \|\mathbb{Q}\|_\infty)\,(\mathcal{R}(\infty;\mathbb{H}) + \mathcal{R}(\infty;\mathbb{Q}))$$

Moreover, we have

$$\mathcal{R}(\infty;\mathbb{A}_1) \leq \|\pi(A \mid X)\|_\infty \|Y\|_\infty \mathcal{R}(\infty;\mathbb{Q}).$$

From the Dudley integral (Lemma 42), the covering number of the VC-subgraph class (Lemma 43), and the boundedness of $\|\pi(A \mid X)\|_\infty$ and $\|Y\|_\infty$, we have

$$\mathcal{R}(\infty;\mathbb{H}) = O\left(\sqrt{V(\mathbb{H})/n}\right), \mathcal{R}(\infty;\mathbb{Q}) = O\left(\sqrt{V(\mathbb{Q})/n}\right).$$

In addition,

$$\mathcal{R}(\infty;\mathbb{A}_3) = O\left(\sqrt{V(\mathbb{A}_3)/n}\right) = O\left(\sqrt{V(\mathbb{H})/n}\right).$$

In the second equality, we use the proof of (Uehara et al., 2021, Corollary 9) Therefore, we have

$$\max\left\{|\hat{J}_{\mathrm{REG}} - J|, |\hat{J}_{\mathrm{IPW}} - J|\right\} = O\left(\sqrt{V(\mathbb{Q})/n} + \sqrt{V(\mathbb{H})/n} + \sqrt{\log(1/\delta)/n}\right).$$

■

Proof of Corollary 12. From the proof of Corollary 11, we have

$$|\hat{J} - J| \leq c(\mathcal{R}(\infty;\mathbb{Q}) + \mathcal{R}(\infty;\mathbb{H}) + \mathcal{R}(\infty;\mathcal{T}\mathbb{H}) + \sqrt{\log(1/\delta)/n}).$$

for some universal constant c. According to Lemma 9 in Uehara et al. (2021), we have

$$\log \mathcal{N}(\epsilon, \mathcal{T}\mathbb{H}, \|\cdot\|_\infty) \leq \log \mathcal{N}(\epsilon, \mathbb{H}, \|\cdot\|_\infty).$$

Then, from Lemma 42, we have

$$|\hat{J} - J| \leq c \inf_{\tau \geq 0}\left\{\tau + \int_\tau^{\sup_{f\in\mathbb{Q}}\sqrt{\mathbb{E}_n[f^2]}} \sqrt{\frac{\log \mathcal{N}(\beta, \mathbb{Q}, \|\cdot\|_n)}{n}} d\beta\right\}$$
$$+ c \inf_{\tau \geq 0}\left\{\tau + \int_\tau^{\sup_{f\in\mathbb{H}}\sqrt{\mathbb{E}_n[f^2]}} \sqrt{\frac{\log \mathcal{N}(\beta, \mathbb{H}, \|\cdot\|_n)}{n}} d\beta\right\} + c\sqrt{\log(1/\delta)/n}.$$

By calculating Dudley integral for each term, the statement is concluded following the proof of Corollary 2 in Uehara et al. (2021). ■

I.5 Proofs for Section 6

Proof of Theorem 13. We use notation in the proof of Theorem 7. Then,

$$\begin{aligned}\mathbb{E}[\{P_z(\hat{h}-h_0)\}^2] &\leq \sup_{q\in\mathbb{Q}'}\left|\mathbb{E}[P_z(\hat{h}-h_0)q]\right| && (P_z(\mathbb{H}-h_0)\subseteq\mathbb{Q}')\\ &\leq \sup_{q\in\mathbb{Q}'}\left|\mathbb{E}[(\hat{h}-h_0)q]\right|\\ &\leq 2\sup_{q\in\mathbb{Q}',h\in\mathbb{H}}|(\mathbb{E}-\mathbb{E}_n)[\{y-h\}q]|,\end{aligned}$$

where the last inequality follows from the second step in the proof for $\hat{J}_{\mathrm{REG}}$ in Theorem 7.

$$\begin{aligned}\mathbb{E}[\{\pi P_w(\hat{q}-q_0)\}^2] &\leq \sup_{h\in\mathbb{H}'}|\mathbb{E}[\pi P_w(\hat{q}-q_0)h]| && (\pi P_w(\mathbb{Q}-q_0)\subseteq\mathbb{H}')\\ &\leq \sup_{h\in\mathbb{H}'}|\mathbb{E}[\pi(\hat{q}-q_0)h]|\\ &\leq 2\sup_{q\in\mathbb{Q},h\in\mathbb{H}'}|(\mathbb{E}-\mathbb{E}_n)[-q\pi h+\mathcal{T}h]|,\end{aligned}$$

where the last inequality follows from the second step in the proof for $\hat{J}_{\mathrm{IPW}}$ in Theorem 7. ■

Proof of Theorem 14. *IPW estimators* The error is decomposed into the three terms:

$$\begin{aligned}&|\hat{J}_{\mathrm{IPW}}-J|\\ =&|\mathbb{E}_n[\hat{q}\pi y]-J|\\ \leq&|\{\mathbb{E}_n[(\hat{q}-q_0)\pi y]-\mathbb{E}[(\hat{q}-q_0)\pi y]\}|+|\{\mathbb{E}[\hat{q}\pi y]-\mathbb{E}[q_0\pi y]\}|+|\mathbb{E}_n[q_0\pi y]-J|\\ \leq&|\{\mathbb{E}_n[\hat{q}\pi y]-\mathbb{E}[\hat{q}\pi y]\}|+|\{\mathbb{E}[\hat{q}\pi y]-\mathbb{E}[q_0\pi y]\}|+|\{\mathbb{E}_n[q_0\pi y]-\mathbb{E}[q_0\pi y]\}|+|\mathbb{E}_n[q_0\pi y]-J|\end{aligned} \tag{50}$$

The first term in (50) is upper bounded by

$$\sup_{q\in\mathbb{Q}}|(\mathbb{E}_n-\mathbb{E})[q\pi y]|.$$

The second term in (50) is

$$\begin{aligned}\mathbb{E}[\hat{q}\pi Y]-\mathbb{E}[q_0\pi Y] &= \mathbb{E}[\{(\hat{q}-q_0)\pi\}Y]=\mathbb{E}[\{(\hat{q}-q_0)\pi\}\mathbb{E}[Y\mid Z,A,X]]\\ &= \mathbb{E}[\{(\hat{q}-q_0)\pi\}\mathbb{E}[h_0\mid Z,A,X]] \qquad \text{(Use the assumption } \mathbb{H}_0\neq\emptyset)\\ &= \mathbb{E}[\{(\hat{q}-q_0)\pi\}h_0]\\ &= \mathbb{E}[P_w\{(\hat{q}-q_0)\pi\}h_0].\end{aligned}$$

Thus, from CS inequality, this term is upper-bounded by

$$\|P_w\{(\hat{q}-q_0)\pi\}\|_2\|h_0\|_2.$$

The third and fourth terms in (50) are upper-bounded by Bernstein inequality. This concludes

$$|\hat{J}_{\mathrm{IPW}}-J|\leq c_1\mathcal{R}(\infty;\pi\mathbb{Q})+\|P_w\{(\hat{q}-q_0)\pi\}\|_2\|h_0\|_2+c_1\sqrt{\log(c_2/\delta)/n}. \tag{51}$$

REG estimators

The error is decomposed into the three terms:

$$\begin{aligned}
&|\hat{J}_{\mathrm{REG}} - J| \\
\leq &|\{\mathbb{E}_n[\mathcal{T}(\hat{h} - h_0)] - \mathbb{E}[\mathcal{T}(\hat{h} - h_0)]\}| + |\{\mathbb{E}[\mathcal{T}\hat{h}] - \mathbb{E}[\mathcal{T}h_0]\}| + |\mathbb{E}_n[\mathcal{T}h_0] - J| \\
\leq &|\{\mathbb{E}_n[\mathcal{T}\hat{h}] - \mathbb{E}[\mathcal{T}\hat{h}]\}| + |\{\mathbb{E}[\mathcal{T}\hat{h}] - \mathbb{E}[\mathcal{T}h_0]\}| + |\{\mathbb{E}_n[\mathcal{T}h_0] - \mathbb{E}[\mathcal{T}h_0]\}| + |\mathbb{E}_n[\mathcal{T}h_0] - J| \quad (52)
\end{aligned}$$

The first term in Equation (52) is upper bounded by

$$|\sup_{h \in \mathbb{H}} (\mathbb{E}_n - \mathbb{E})[\mathcal{T}h]|.$$

The second term in Equation (52) is upper-bounded as follows:

$$\begin{aligned}
|\mathbb{E}[\mathcal{T}\hat{h}] - \mathbb{E}[\mathcal{T}h_0]| &= |\mathbb{E}[\pi(A \mid X)/f(A \mid W, X)\{\hat{h} - h_0\}]| \\
&= |\mathbb{E}[\mathbb{E}[\pi(A \mid X) q_0(Z, A, X) \mid W, A, X]\{\hat{h} - h_0\}]| \\
&\qquad \text{(Use the assumption } \mathbb{Q}_0 \neq \emptyset) \\
&= |\mathbb{E}[\pi(A \mid X) q_0(Z, A, X)\{\hat{h} - h_0\}]| \\
&= |\mathbb{E}[\pi(A \mid X) q_0(Z, A, X) P_z\{\hat{h} - h_0\}]| \leq \|q_0 \pi\|_2 \|P_z\{\hat{h} - h_0\}\|_2.
\end{aligned}$$

The third and fourth terms in Equation (52) are upper-bounded by Bernstein inequality. This concludes

$$|\hat{J}_{\mathrm{REG}} - J| \leq c_1 \mathcal{R}(\infty; \mathcal{T}\mathbb{H}) + \|q_0 \pi\|_2 \|P_z\{\hat{h} - h_0\}\|_2 + c_1 \sqrt{\log(c_2/\delta)/n}). \quad (53)$$

■

Proof of Theorem 15. The proof is similar to Theorem 14. The error is decomposed into the three terms:

$$\begin{aligned}
|\hat{J}_{\mathrm{DR}} - J| \leq &\left|(\mathbb{E}_n - \mathbb{E})\left[\{\pi\hat{q}\{Y - \hat{h}\} + \mathcal{T}\hat{h}\} - \{\pi q_0\{Y - h_0\} + \mathcal{T}h_0\}\right]\right| \\
&+ |\{\mathbb{E}[\{\pi\hat{q}\{Y - \hat{h}\} + \mathcal{T}\hat{h}\}] - \mathbb{E}[\{\pi q_0\{Y - h_0\} + \mathcal{T}h_0\}]\}| \\
&+ |\mathbb{E}_n[\{\pi q_0\{Y - h_0\} + \mathcal{T}h_0\}] - J|
\end{aligned}$$

It follows that

$$\begin{aligned}
|\hat{J}_{\mathrm{DR}} - J| \leq &\left|(\mathbb{E}_n - \mathbb{E})\left[\{\pi\hat{q}\{Y - \hat{h}\} + \mathcal{T}\hat{h}\}\right]\right| \\
&+ |\{\mathbb{E}[\{\pi\hat{q}\{Y - \hat{h}\} + \mathcal{T}\hat{h}\}] - \mathbb{E}[\{\pi q_0\{Y - h_0\} + \mathcal{T}h_0\}]\}| \\
&+ |(\mathbb{E}_n - \mathbb{E})\left[\{\pi q_0\{Y - h_0\} + \mathcal{T}h_0\}\right]| \\
&+ |\mathbb{E}_n[\{\pi q_0\{Y - h_0\} + \mathcal{T}h_0\}] - J| \quad (54)
\end{aligned}$$

The first term above is upper-bounded by

$$|\sup_{h \in \mathbb{H}, q \in \mathbb{Q}} (\mathbb{E}_n - \mathbb{E})[\pi q\{Y - h\} + \mathcal{T}h]|.$$

The second term above is equal to

$$\left|\mathbb{E}\left[\pi\left(\hat{q}-q_0\right)\left(\hat{h}-h_0\right)\right]\right|.$$

This term can be upper bounded by both

$$\|P_w\{(\hat{q}-q_0)\pi\}\|_2 \sup_{h\in\mathbb{H}}\|h_0-h\|_2$$

and

$$\|P_z\{\hat{h}-h_0\}\|_2 \sup_{q\in\mathbb{Q}}\|\{q_0-q\}\pi\|_2.$$

The third and fourth terms are upper-bounded by Bernstein's inequality. Then, $|\hat{J}_{\mathrm{DR}}-J|$ is upper-bounded as the statements. ■

I.6 Proofs for Section 7

Here we prove Theorems 16 and 18. Conclusions in Corollaries 17 and 19 then follow immediately.

Proof of Theorem 16. Define

$$\begin{aligned}
\Phi(h,q) &= \mathbb{E}[\{Y-h(W,A,X)\}q(Z,A,X)]\\
\Phi_n(h,q) &= \mathbb{E}_n[\{Y-h(W,A,X)\}q(Z,A,X)]\\
\Phi^\lambda(h,q) &= \mathbb{E}[\{Y-h(W,A,X)\}q(Z,A,X)]-\lambda\|q\|_n^2\\
\Phi_n^\lambda(h,q) &= \mathbb{E}_n[\{Y-h(W,A,X)\}q(Z,A,X)]-\lambda\|q\|_2^2.
\end{aligned}$$

where $\|q\|_n=\{\mathbb{E}_n[q^2]\}^{1/2}$ and $\|q\|_2=\{\mathbb{E}[q^2]\}^{1/2}$. From Lemma 46, we have

$$\forall q\in\mathbb{Q}', |\|q\|_n^2-\|q\|_2^2|\le 0.5\|q\|_2^2+0.5(\eta_h')^2 \tag{55}$$

for our choice of $\eta_h':=\eta_h+\sqrt{c_0\log(c_1/\delta)/n}$, where η_h upper bounds the critical radius of $\mathbb{Q}'$.

Step I.

By definition of the estimator $\hat{h}$ and the assumption $h'\in\mathbb{H}$, we have

$$\sup_{q\in\mathbb{Q}'}\Phi_n^\lambda(\hat{h},q)\le\sup_{q\in\mathbb{Q}'}\Phi_n^\lambda(h',q). \tag{56}$$

From Lemma 47, we have that there exists a positive constant c such that

$$\forall q\in\mathbb{Q}': |\Phi_n(h',q)-\Phi(h',q)|\le cC_1\{\eta_h'\|q\|_2+(\eta_h')^2\}. \tag{57}$$

To prove this, we apply Lemma 47 to $l(a_1,a_2):=a_1a_2, a_1=q(Z,A,X), a_2=Y-h'(W,A,X)$ that is C_1-Lipschitz with respect to a_1 by noting $Y-h'(W,A,X)$ is in $[-C_1,C_1]$ with some constants $C_1=\|Y\|_\infty+\|\mathbb{H}\|_\infty$:

$$|l(a_1,a_2)-l(a_1',a_2)|\le C_1|a_1-a_1'|.$$

Thus,

$$\begin{aligned}
\sup_{q\in\mathbb{Q}'} \Phi_n^\lambda(h',q) &\le \sup_{q\in\mathbb{Q}'}\{\Phi_n(h',q) - \lambda\|q\|_n^2\} \\
&\le \sup_{q\in\mathbb{Q}'}\{\Phi(h',q) + cC_1\{\eta_h'\|q\|_2 + (\eta_h')^2\} - \lambda\|q\|_n^2\} \\
&\quad \text{From Equation (57)} \\
&\le \sup_{q\in\mathbb{Q}'}\{\Phi(h',q) + cC_1\{\eta_h'\|q\|_2 + (\eta_h')^2\} - 0.5\lambda\|q\|_2^2 + 0.5\lambda(\eta_h')^2\} \\
&\quad \text{From Equation (55)} \\
&\le \sup_{q\in\mathbb{Q}'}\{\Phi(h',q) - 0.25\lambda\|q\|_2^2 + cC_1\{\eta_h'\|q\|_2 + (\eta_h')^2\} - 0.25\lambda\|q\|_2^2 + \lambda(\eta_h')^2\} \\
&\le \sup_{q\in\mathbb{Q}'}\{\Phi(h',q) - 0.25\lambda\|q\|_2^2 + (\lambda + c^2C_1^2/\lambda + cC_1)(\eta_h')^2\}.
\end{aligned}$$

In the last line, we use a general inequality that for any $a, b > 0$:

$$\sup_{q\in\mathbb{Q}'}(a\|q\|_2 - b\|q\|_2^2) \le a^2/4b. \tag{58}$$

Moreover,

$$\begin{aligned}
\sup_{q\in\mathbb{Q}'} \Phi_n^\lambda(\hat{h},q) &= \sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) + \Phi_n(h',q) - \lambda\|q\|_n^2\} \\
&\ge \sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) - 2\lambda\|q\|_n^2\} + \inf_{q\in\mathbb{Q}'}\{\Phi_n(h',q) + \lambda\|q\|_n^2\} \\
&= \sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) - 2\lambda\|q\|_n^2\} + \inf_{-q\in\mathbb{Q}'}\{\Phi_n(h',-q) + \lambda\|q\|_n^2\} \\
&= \sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) - 2\lambda\|q\|_n^2\} + \inf_{-q\in\mathbb{Q}'}\{-\Phi_n(h',q) + \lambda\|q\|_n^2\} \\
&= \sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) - 2\lambda\|q\|_n^2\} - \sup_{-q\in\mathbb{Q}'}\{\Phi_n(h',q) - \lambda\|q\|_n^2\} \\
&= \sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) - 2\lambda\|q\|_n^2\} - \sup_{-q\in\mathbb{Q}'}\{\Phi_n^\lambda(h',q)\} \\
&= \sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) - 2\lambda\|q\|_n^2\} - \sup_{q\in\mathbb{Q}'}\{\Phi_n^\lambda(h',q)\},
\end{aligned}$$

where in the last equation we use the symmetry of $\mathbb{Q}'$.

Therefore,

$$\begin{aligned}
\sup_{q\in\mathbb{Q}'}\{\Phi_n(\hat{h},q) - \Phi_n(h',q) - 2\lambda\|q\|_n^2\} &\le \sup_{q\in\mathbb{Q}'}\{\Phi_n^\lambda(\hat{h},q)\} + \sup_{q\in\mathbb{Q}'}\{\Phi_n^\lambda(h',q)\} \\
&\le 2\sup_{q\in\mathbb{Q}'}\{\Phi_n^\lambda(h',q)\} \qquad \text{(From (56))} \\
&\le 2\sup_{q\in\mathbb{Q}'}\{\Phi(h',q) - 0.25\lambda\|q\|_2^2 + (\lambda + c^2C_1^2/\lambda + cC_1)(\eta_h')^2\} \\
&\le 2\sup_{q\in\mathbb{Q}'}\{\|P_z(h_0 - h')\|_2\|q\|_2 - 0.25\lambda\|q\|_2^2 + (\lambda + c^2C_1^2/\lambda + cC_1)(\eta_h')^2\} \\
&\le 2\{\|P_z(h_0 - h')\|_2^2/\lambda + (\lambda + c^2C_1^2/\lambda + cC_1)(\eta_h')^2\},
\end{aligned}$$

where the last inequality again uses Equation (58).

Step II. For the fixed $h' \in \mathbb{H}$ and for any $h \in \mathbb{H}$ define

$$q_h := \arg\min_{q\in\mathbb{Q}'} \|q - P_z(h' - h)\|.$$

Further define

$$\epsilon_n = \sup_{h \in \mathbb{H}} \inf_{q \in \mathbb{Q}'} \|q - P_z(h - h')\|_2.$$

According to the asserted assumptions, we have $\|q_h - P_z(h' - h)\| \leq \epsilon_n$.

Suppose $\|q_{\hat{h}}\|_2 \geq \eta'_h$, and let $r = \eta'_h / \{2\|q_{\hat{h}}\|_2\} \in [0, 0.5]$. Then, noting $\mathbb{Q}'$ is star-shaped and symmetric, we have $rq_{\hat{h}} \in \mathbb{Q}$ and

$$\begin{aligned} r^2\|q_{\hat{h}}\|_n^2 &\leq r^2\{1.5\|q_{\hat{h}}\|_2^2 + 0.5(\eta'_h)^2\} && \text{From Equation (55)} \\ &\leq (\eta'_h)^2. && \text{From definition of } r. \end{aligned}$$

It follows that

$$\begin{aligned} \sup_{q \in \mathbb{Q}'} \{\Phi_n(\hat{h}, q) - \Phi_n(h', q) - 2\lambda\|q\|_n^2\} &\geq r\{\Phi_n(\hat{h}, q_{\hat{h}}) - \Phi_n(h', q_{\hat{h}})\} - 2\lambda r^2\|q_{\hat{h}}\|_n^2 \\ &\geq r\{\Phi_n(\hat{h}, q_{\hat{h}}) - \Phi_n(h', q_{\hat{h}})\} - 2\lambda(\eta'_h)^2. \end{aligned}$$

Observe that

$$\Phi_n(h, q_{\hat{h}}) - \Phi_n(h', q_{\hat{h}}) = \mathbb{E}_n[\{h' - h\}q_{\hat{h}}(Z, A, X)].$$

We now invoke Lemma 47 with $l(a_1, a_2), a_1 = (h - h')q$ and $a_2 = 1$, by noting that η'_h upper bounds the critical radius of $\text{star}(\mathcal{G}_h)$:

$$\begin{aligned} &|\Phi_n(h, q_h) - \Phi_n(h', q_h) - \{\Phi(h, q_h) - \Phi(h', q_h)\}| \\ &\leq c(\eta'_h\|\{h - h'\}q_{\hat{h}}\|_2 + (\eta'_h)^2) \leq (cC_2\eta'_h\|q_{\hat{h}}\|_2 + c(\eta'_h)^2), \end{aligned}$$

where the last line uses $\|h - h'\|_\infty \leq C_2$ for $C_2 = 2\|\mathbb{H}\|_\infty$.

Thus,

$$\begin{aligned} r\{\Phi_n(\hat{h}, q_{\hat{h}}) - \Phi_n(h', q_{\hat{h}})\} &\geq r\{\Phi(\hat{h}, q_{\hat{h}}) - \Phi(h', q_{\hat{h}})\} - rc(C_2\eta'_h\|q_{\hat{h}}\|_2 + (\eta'_h)^2) \\ &\geq r\{\Phi(\hat{h}, q_{\hat{h}}) - \Phi(h', q_{\hat{h}})\} - rcC_2\eta'_h\|q_{\hat{h}}\|_2 - 0.5c(\eta'_h)^2 \\ &\qquad (r \in [0, 0.5]) \\ &\overset{(a)}{=} r\mathbb{E}[P_z(-\hat{h} + h')q_{\hat{h}}] - rcC_2\eta'_h\|q_{\hat{h}}\|_2 - 0.5c(\eta'_h)^2 \\ &= \frac{\eta'_h}{2\|q_{\hat{h}}\|_2}\{\mathbb{E}[P_z(-\hat{h} + h')q_{\hat{h}}] - cC_2\|q_{\hat{h}}\|_2\eta'_h\} - 0.5c(\eta'_h)^2 \\ &\qquad \text{(Definition of } r) \\ &\overset{(b)}{\geq} 0.5\eta'_h\{\|P_z(-\hat{h} + h')\|_2 - 2\epsilon_n\} - 0.5c(1 + C_2)(\eta'_h)^2, \end{aligned}$$

where (a) follows from

$$\Phi(\hat{h}, q_{\hat{h}}) - \Phi(h', q_{\hat{h}}) = \mathbb{E}[\{-\hat{h} + h'\}q_{\hat{h}}] = \mathbb{E}[P_z(-\hat{h} + h')q_{\hat{h}}],$$

and (b) follows from

$$\frac{\mathbb{E}[P_z(-\hat{h} + h')q_{\hat{h}}]}{\|q_{\hat{h}}\|_2} = \frac{\mathbb{E}[\{-q_{\hat{h}} + q_{\hat{h}} + P_z(-\hat{h} + h')\}q_{\hat{h}}]}{\|q_{\hat{h}}\|_2}$$

$$
\begin{aligned}
&\geq \frac{\|q_{\hat{h}}\|_2^2 - \|\{-q_{\hat{h}} + P_z(-\hat{h} + h')\}\|_2 \|q_{\hat{h}}\|_2}{\|q_{\hat{h}}\|_2} \\
&= \|q_{\hat{h}}\|_2 - \|\{-q_{\hat{h}} + P_z(-\hat{h} + h')\}\|_2 \\
&\geq \|q_{\hat{h}}\|_2 - \epsilon_n \geq \|P_z(-\hat{h} + h')\|_2 - 2\epsilon_n.
\end{aligned}
$$

Step III: Combining All Results.

Thus, we have either $\|q_{\hat{h}}\|_2 \leq \eta_h'$ or with high probability $1 - \delta$,

$$
0.5\eta_h'\{\|P_z(\hat{h} - h')\|_2 - 2\epsilon_n\} - (0.5c + 0.5cC_2 + 2\lambda)(\eta_h')^2 \leq \frac{2\|P_z(h_0 - h')\|_2^2}{\lambda} + 2(\lambda + c^2 C_1^2/\lambda + cC_1)(\eta_h')^2.
$$

Therefore, we have either

$$
\|P_z(\hat{h} - h')\|_2 \leq \|P_z(h_0 - h')\|_2 + \|q_{\hat{h}}\|_2 + \|P_z(\hat{h} - h') - q_{\hat{h}}\|_2 \leq \|P_z(h_0 - h')\|_2 + \eta_h' + \epsilon_n
$$

or

$$
\|P_z(\hat{h} - h')\|_2 \lesssim (1 + \lambda + 1/\lambda)\eta_h' + \frac{\|P_z(h_0 - h')\|_2^2}{\eta_h'\lambda} + \epsilon_n.
$$

Thus, from triangle inequality,

$$
\|P_z(h_0 - \hat{h})\|_2 \lesssim (1 + \lambda + 1/\lambda)\eta_h' + \frac{\|P_z(h_0 - h')\|_2^2}{\eta_h'\lambda} + \epsilon_n + \|P_z(h_0 - h')\|_2.
$$

■

Proof of Theorem 18. Define

$$
\begin{aligned}
\Phi(q, h) &= \mathbb{E}[-q(Z, A, X)\pi(A \mid X)h(W, A, X) + \mathbb{E}_{a \sim \pi(A|X)}[h(W, A, X)]] \\
\Phi_n(q, h) &= \mathbb{E}_n[-q(Z, A, X)\pi(A \mid X)h(W, A, X) + \mathbb{E}_{a \sim \pi(A|X)}[h(W, A, X)]] \\
\Phi^\lambda(q, h) &= \mathbb{E}[-q(Z, A, X)\pi(A \mid X)h(W, A, X) + \mathbb{E}_{a \sim \pi(A|X)}[h(W, A, X)]] - \lambda\mathbb{E}[h^2] \\
\Phi_n^\lambda(q, h) &= \mathbb{E}_n[-q(Z, A, X)\pi(A \mid X)h(W, A, X) + \mathbb{E}_{a \sim \pi(A|X)}[h(W, A, X)]] - \lambda\mathbb{E}_n[h^2].
\end{aligned}
$$

From Lemma 46, we have

$$
\forall h \in \mathbb{H}',\ |\|h\|_n^2 - \|h\|_2^2| \leq 0.5\|h\|_2^2 + 0.5\left(\eta_q'\right)^2 \tag{59}
$$

for our choice of $\eta_q' := \eta_q + \sqrt{c_0 \log(c_1/\delta)/n}$, where η_q upper bounds the critical radius of $\mathbb{H}'$.

First Part.

By definition of $\hat{q}$ and $q' \in \mathbb{Q}$, we have

$$
\sup_{h \in \mathbb{H}'} \Phi_n^\lambda(\hat{q}; h) \leq \sup_{h \in \mathbb{H}'} \Phi_n^\lambda(q'; h). \tag{60}
$$

From Lemma 47, there exists a universal constant $c > 0$ and $C_3 > 0$ such that with probability $1 - \delta$,

$$
\forall h \in \mathbb{H}', |\Phi_n(q', h) - \Phi(q', h)| \leq cC_3\{\eta_q'\|h\|_2 + \left(\eta_q'\right)^2\}. \tag{61}
$$

To prove this, we first apply Lemma 47 to $l(a_1, a_2) = a_1 a_2$ with $a_1 = -q'(Z, A, X)\pi(A \mid X), a_2 = h(W, A, X)$. This function is $C_{3,1}$-Lipschitz in a_2 where $C_{3,1} = \|\pi\mathbb{Q}\|_\infty$. This means that

$$|\mathbb{E}_n[-q'(Z, A, X)\pi(A \mid X)h(W, A, X)] - \mathbb{E}[-q'(Z, A, X)\pi(A \mid X)h(W, A, X)]| \leq cC_{3,1}\{\eta'_q\|h\|_2 + \left(\eta'_q\right)^2\}.$$

Similarly, we have

$$|\mathbb{E}_n[\mathcal{T}h] - \mathbb{E}[\mathcal{T}h]| \leq c\{\eta'_q\|\mathcal{T}h\|_2 + \left(\eta'_q\right)^2\} \leq cC_{3,2}\{\eta'_q\|h\|_2 + \left(\eta'_q\right)^2\}, \quad C_{3,2} = \|\pi(A \mid X)/f(A \mid X, W)\|_2.$$

Combining the above two equations gives Equation (61) with $C_3 = C_{3,1} + C_{3,2}$.

Thus,

$$\begin{aligned}
\sup_{h\in\mathbb{H}'} \Phi_n^\lambda(q', h) &= \sup_{h\in\mathbb{H}'}\{\Phi_n(q', h) - \lambda\|h\|_n^2\} \\
&\leq \sup_{h\in\mathbb{H}'}\{\Phi(q', h) + cC_3(\eta'_q\|h\|_2 + \left(\eta'_q\right)^2) - \lambda\|h\|_n^2\} && \text{(Use Equation (61))} \\
&\leq \sup_{h\in\mathbb{H}'}\{\Phi(q', h) + cC_3(\eta'_q\|h\|_2 + \left(\eta'_q\right)^2) - 0.5\lambda\|h\|_2^2 + 0.5\lambda\left(\eta'_q\right)^2\} \\
& && \text{(Use Equation (59))} \\
&\leq \sup_{h\in\mathbb{H}'}\{\Phi(q', h) - 0.25\lambda\|h\|_2^2 + cC_3(\eta'_q\|h\|_2 + \left(\eta'_q\right)^2) - 0.25\lambda\|h\|_2^2 + \lambda\left(\eta'_q\right)^2\} \\
&\leq \sup_{h\in\mathbb{H}'}\{\Phi(q', h) - 0.25\lambda\|h\|_2^2 + \{\lambda + c^2C_3^2/\lambda + cC_3\}\left(\eta'_q\right)^2\}.
\end{aligned}$$

In the last line, we use a general inequality that for any $a, b > 0$,

$$a\|h\|_2 - b\|h\|_2^2 \leq a^2/4b.$$

Moreover,

$$\begin{aligned}
&\sup_{h\in\mathbb{H}'} \Phi_n^\lambda(\hat{q}, h) \\
&= \sup_{h\in\mathbb{H}'}\{\Phi_n(\hat{q}, h) - \Phi_n(q', h) + \Phi_n(q', h) - \lambda\|h\|_n^2\} \\
&\geq \sup_{h\in\mathbb{H}'}\{\Phi_n(\hat{q}, h) - \Phi_n(q', h) - 2\lambda\|h\|_n^2\} + \inf_{h\in\mathbb{H}'}\{\Phi_n(q', h) - \lambda\|h\|_n^2\} \\
&= \sup_{h\in\mathbb{H}'}\{\Phi_n(\hat{q}, h) - \Phi_n(q', h) - 2\lambda\|h\|_n^2\} + \inf_{-h\in\mathbb{H}'}\{\Phi(q', -h) - \lambda\|h\|_n^2\} \\
&= \sup_{h\in\mathbb{H}'}\{\Phi_n(\hat{q}, h) - \Phi_n(q', h) - 2\lambda\|h\|_n^2\} - \sup_{-h\in\mathbb{H}'}\{\Phi_n(q', h) + \lambda\|h\|_n^2\} \\
&= \sup_{h\in\mathbb{H}'}\{\Phi_n(\hat{q}, h) - \Phi_n(q', h) - 2\lambda\|h\|_n^2\} - \sup_{h\in\mathbb{H}'}\{\Phi_n^\lambda(q', h)\}.
\end{aligned}$$

where in the last equation we use the symmetry of $\mathbb{H}'$.

Therefore,

$$\begin{aligned}
\sup_{h\in\mathbb{H}'}\{\Phi_n(\hat{q}, h) - \Phi_n(q', h) - 2\|h\|_n^2\} &\leq \sup_{h\in\mathbb{H}'}\{\Phi_n^\lambda(\hat{q}, h)\} + \sup_{h\in\mathbb{H}'}\{\Phi_n^\lambda(q', h)\} \\
&\leq 2\sup_{h\in\mathbb{H}'}\{\Phi_n^\lambda(q', h)\} && \text{(From Equation (60))} \\
&\leq 2\sup_{h\in\mathbb{H}'}\{\Phi(q', h) - 0.25\lambda\|h\|_2^2 + \{\lambda + c^2C_3^2/\lambda + cC_3\}\left(\eta'_q\right)^2\}
\end{aligned}$$

$$\leq \sup_{h\in\mathbb{H}'} 2\{\|\pi P_w(q'-q_0)\|_2\|h\|_2 - 0.25\lambda\|h\|_2^2 + c\{C_2^2/\lambda+\lambda+C_2\}\left(\eta_q'\right)^2\}$$
$$\leq 2\{\|\pi P_w(q'-q_0)\|_2^2/\lambda + c\{C_2^2/\lambda+\lambda+C_2\}\left(\eta_q'\right)^2\}$$

Second Part. For the fixed q' and for any $q\in\mathbb{Q}$ we define

$$h_q \coloneqq \arg\min_{h\in\mathbb{H}'} \|h - P_w\{(q-q')\pi\}\|.$$

Further define

$$\tilde{\epsilon}_n = \sup_{q\in\mathbb{Q}} \inf_{h\in\mathbb{H}'} \|h - \pi P_w\big(q-q'\big)\|_2.$$

According to the asserted assumptions, we have $\|h_q - P_w\{(q-q')\pi\}\| \leq \tilde{\epsilon}_n$.

Suppose $\|h_{\hat{q}}\| \geq \eta_q'$, and let $r = \eta_q'/\{2\|h_{\hat{q}}\|_2\} \in (0, 0.5]$. Then, noting $\mathbb{H}'$ is star-shaped and symmetric, we have $rh_{\hat{q}} \in \mathbb{H}'$ and

$$r^2\|h_{\hat{q}}\|_n^2 \leq r^2\{1.5\|h_{\hat{q}}\|_2^2 + 0.5\left(\eta_q'\right)^2\}$$
$$\leq \left(\eta_q'\right)^2. \qquad \text{(Definition of } r\text{)}$$

It follows that

$$\sup_{h\in\mathbb{H}'}\{\Phi_n(\hat{q},h) - \Phi_n(q',h) - 2\lambda\|h\|_n^2\} \geq r\{\Phi_n(\hat{q},h_{\hat{q}}) - \Phi_n(q',h_{\hat{q}})\} - 2\lambda r^2\|h_{\hat{q}}\|_n^2$$
$$\geq r\{\Phi_n(\hat{q},h_{\hat{q}}) - \Phi_n(q',h_{\hat{q}})\} - 2\lambda\eta_n^2.$$

Observe that

$$\Phi_n(q,h_q) - \Phi_n(q',h_q) = \mathbb{E}_n[(-q(Z,A,X) + q'(Z,A,X))\pi(a|x)h_q(W,A,X)].$$

We now invoke Lemma 47 with $l(a_1,a_2) = a_1a_2$ with $a_1 = (q-q')\pi h_q, a_2 = 1$, by noting that η_q' upper bounds the critical radius of $\mathcal{G}_q$:

$$|\Phi_n(q,h_q) - \Phi_n(q',h_q) - \{\Phi(q,h_q) - \Phi(q',h_q)\}|$$
$$\leq c(\eta_q'\mathbb{E}[\{(q'(Z,A,X) - q(Z,A,X))\pi(A\mid X)h_q(W,A,X)\}^2]^{1/2} + \left(\eta_q'\right)^2)$$
$$\leq c(C_4\eta_q'\|h_q\|_2 + \left(\eta_q'\right)^2).$$

where $\|(q'(Z,A,X) - q(Z,A,X))\pi(A\mid X)\|_\infty \leq C_4$ for $C_4 = 2\|\pi\mathbb{Q}\|_\infty$.

Thus,

$$r\{\Phi_n(\hat{q},h_{\hat{q}}) - \Phi_n(q',h_{\hat{q}})\}$$
$$\geq r\{\Phi(\hat{q},h_{\hat{q}}) - \Phi(q',h_{\hat{q}})\} - rc(C_4\eta_q'\|h_{\hat{q}}\|_2 + \left(\eta_q'\right)^2)$$
$$\geq r\{\Phi(\hat{q},h_{\hat{q}}) - \Phi(q',h_{\hat{q}})\} - rc(C_4\eta_q'\|h_{\hat{q}}\|_2) - 0.5c\left(\eta_q'\right)^2 \qquad (r\in[0,0.5])$$
$$\overset{(a)}{=} r\mathbb{E}[P_w\left[(-q'+\hat{q})\pi h_{\hat{q}}\right]] - rc(C_4\eta_q'\|h_{\hat{q}}\|_2) - 0.5c\left(\eta_q'\right)^2$$
$$= \frac{\eta_q'}{2\{\|h_{\hat{q}}\|_2\}}(\mathbb{E}[P_w\left[(-q'+\hat{q})\pi h_{\hat{q}}\right]] - cC_4\eta_q'\|h_{\hat{q}}\|_2) - 0.5c\left(\eta_q'\right)^2 \qquad \text{(Definition of } r\text{)}$$

$$\overset{(b)}{\geq} 0.5\eta_q'(\|P_w\{(-q'+\hat{q})\pi\}\|_2 - 2\tilde{\epsilon}_n) - 0.5c\{1+C_4\}\left(\eta_q'\right)^2,$$

where (a) follows from

$$\begin{aligned}\Phi(q,h_q)-\Phi(q',h_q) &= \mathbb{E}[\{-q(Z,A,X)+q'(Z,A,X)\}\pi(A\mid X)h_q(W,A,X)]\\ &= \mathbb{E}[\{\mathbb{E}[-q(Z,A,X)+q'(Z,A,X)\mid W,A,X]\}\pi(A\mid X)h_q(W,A,X)]\\ &= \mathbb{E}[P_w\{q'-q\}(W,A,X)\pi(A\mid X)h_q(W,A,X)],\end{aligned}$$

and (b) follows from

$$\begin{aligned}\frac{\mathbb{E}[\pi P_w\{-q'+\hat{q}\}h_q]}{\|h_{\hat{q}}\|_2} &= \frac{\mathbb{E}[[-h_{\hat{q}}+h_{\hat{q}}+\pi P_w\{-q'+\hat{q}\}]h_q]}{\|h_{\hat{q}}\|_2}\\ &= \frac{\|h_{\hat{q}}\|_2^2-\|h_{\hat{q}}\|_2\|-h_{\hat{q}}+\pi P_w\{-q'+\hat{q}\}\|_2}{\|h_{\hat{q}}\|_2}\\ &\geq \|h_{\hat{q}}\|_2-\tilde{\epsilon}_n \geq \|\pi P_w\{-q'+\hat{q}\}\|_2-2\tilde{\epsilon}_n.\end{aligned}$$

Combining All Results. Thus, $\|h_{\hat{q}}\| \leq \eta_q'$ or

$$\begin{aligned}0.5\eta_q'(\|P_w\{(-q'+\hat{q})\pi\}\|_2-2\tilde{\epsilon}_n) - \{0.5c+0.5cC_4+2\lambda\}\left(\eta_q'\right)^2&\\ \leq 2\{\|\pi P_w(q'-q_0)\|_2^2/\lambda + c\{C_2^2/\lambda+\lambda+C_2\}\left(\eta_q'\right)^2\}.&\end{aligned}$$

Therefore, we have either

$$\|P_w(\pi\hat{q}-\pi q')\|_2 \leq \|P_w(\pi q_0-\pi q')\|_2 + \|h_{\hat{q}}-\pi P_w(\pi\hat{q}-\pi q')\|_2 + \|h_{\hat{q}}\|_2 \leq \|P_w(\pi q_0-\pi q')\|_2 + \eta_q'+\tilde{\epsilon}_n.$$

or

$$\|P_w(\pi\hat{q}-\pi q')\|_2 \lesssim \|\pi P_w(q'-q_0)\|_2^2/\lambda\eta_q' + \{1+1/\lambda+\lambda\}\eta_q' + \tilde{\epsilon}_n.$$

Finally, from triangle inequality,

$$\|P_w(\pi\hat{q}-\pi q_0)\|_2 \lesssim \|\pi P_w(q'-q_0)\|_2^2/\lambda\eta_q' + \{1+1/\lambda+\lambda\}\eta_q' + \tilde{\epsilon}_n + \|P_w(\pi q_0-\pi q')\|_2.$$

■

Proof of Corollary 20. For a function $f(O)$, we denote its empirical L^2-norm as $\|f\|_{n,2} = \{1/n\sum_{i=1}^n f(O_i)^2\}^{1/2}$ and L^∞-norm as $\|f(O)\|_{n,\infty} = \max_{1\leq i\leq n}|f(O_i)|$.

Note that

$$\begin{aligned}&\log\mathcal{N}(t,(\mathbb{H}-h_0)\mathbb{Q}',\|\cdot\|_{n,\infty})\\ &\leq \log\mathcal{N}(t,\mathbb{H}\mathbb{Q}',\|\cdot\|_{n,\infty}) + \log\mathcal{N}(t,h_0\mathbb{Q}',\|\cdot\|_{n,\infty})\\ &\leq \log\mathcal{N}(t/\{2\|\mathbb{H}\|_\infty\},\mathbb{Q}',\|\cdot\|_{n,\infty}) + \log\mathcal{N}(t/\{2\|\mathbb{Q}'\|_\infty\},\mathbb{H},\|\cdot\|_{n,\infty}) + \log\mathcal{N}(t/\|\mathbb{H}\|_\infty,\mathbb{Q}',\|\cdot\|_{n,\infty}).\end{aligned}$$

It follows that

$$\int_0^\eta \sqrt{\frac{\log\mathcal{N}(t,(\mathbb{H}-h_0)\mathbb{Q}',\|\cdot\|_{n,2})}{n}}\mathrm{d}(t)$$

$$
\begin{aligned}
&\leq \int_0^\eta \sqrt{\frac{\log \mathcal{N}(t, (\mathbb{H}-h_0)\mathbb{Q}', \|\cdot\|_{n,\infty})}{n}}\mathrm{d}(t) && (\|f\|_{n,2} \leq \|f\|_{n,\infty} \text{ for any } f)\\
&\leq \int_0^\eta \sqrt{\frac{2\log \mathcal{N}(t/2\|\mathbb{H}\|_\infty, \mathbb{Q}', \|\cdot\|_{n,\infty})}{n}} + \sqrt{\frac{\log \mathcal{N}(t/2\|\mathbb{Q}'\|_\infty, \mathbb{H}, \|\cdot\|_{n,\infty})}{n}}\mathrm{d}(t)\\
&\leq \int_0^\eta \sqrt{\frac{2\log \mathcal{N}(t/2\|\mathbb{H}\|_\infty, \mathbb{Q}', \sqrt{n}\|\cdot\|_{n,2})}{n}} + \sqrt{\frac{\log \mathcal{N}(t/2\|\mathbb{Q}'\|_\infty, \mathbb{H}, \sqrt{n}\|\cdot\|_{n,2})}{n}}\mathrm{d}(t).\\
& && (\|f\|_{n,\infty} \leq \|f\|_{n,2}\sqrt{n} \text{ for any } f)\\
&= O\left(\frac{\log n}{\sqrt{n}} \int_0^\eta \sqrt{\max(V(\mathbb{H}), V(\mathbb{Q}')\log(1/t)}\,\mathrm{d}(t)\right) && \text{(Lemma 43)}\\
&= O\left(\sqrt{\max(V(\mathbb{H}), V(\mathbb{Q}')}\eta\log(1/\eta)\frac{\log n}{\sqrt{n}}\right).
\end{aligned}
$$

Then, the critical inequality in Lemma 45 becomes

$$
O\left(\sqrt{\max(V(\mathbb{H}), V(\mathbb{Q})}\eta\log(1/\eta)\frac{\log n}{\sqrt{n}}\right) \leq \eta^2.
$$

This is satisfied with

$$
\eta_h = O\left(\sqrt{\max(V(\mathbb{H}), V(\mathbb{Q})}\frac{\log n}{\sqrt{n}}\log\left(\sqrt{\max(V(\mathbb{H}), V(\mathbb{Q})}\frac{\log n}{\sqrt{n}}\right)\right).
$$

According to Lemma 45, the η_h above upper bounds the critical radius of $\mathcal{G}_h$. By a similar calculation, it can be shown that this η_h can also upper bound the critical radius of $\mathbb{Q}'$. Plugging it into Corollary 17 then proves the conclusion of Corollary 20. ■

Proof for Corollary 21. Again, we have

$$
\begin{aligned}
&\log \mathcal{N}(t, (\mathbb{H}-h_0)\mathbb{Q}', \|\cdot\|_{n,\infty})\\
&\leq \log \mathcal{N}(t, \mathbb{H}\mathbb{Q}', \|\cdot\|_{n,\infty}) + \log \mathcal{N}(t, h_0\mathbb{Q}', \|\cdot\|_{n,\infty})\\
&\leq 2\log \mathcal{N}(t/\{2\|\mathbb{H}\|_\infty\}, \mathbb{Q}', \|\cdot\|_{n,\infty}) + \log \mathcal{N}(t/\{2\|\mathbb{Q}'\|_\infty\}, \mathbb{H}, \|\cdot\|_{n,\infty}).
\end{aligned}
$$

and

$$
\begin{aligned}
&\int_0^\eta \sqrt{\frac{\log \mathcal{N}(t, (\mathbb{H}-h_0)\mathbb{Q}', \|\cdot\|_{n,2})}{n}}\,\mathrm{d}(t)\\
&\leq \int_0^\eta \sqrt{\frac{\log \mathcal{N}(t, (\mathbb{H}-h_0)\mathbb{Q}', \|\cdot\|_{\infty})}{n}}\,\mathrm{d}(t)\\
&\leq \int_0^\eta \sqrt{\frac{2\log \mathcal{N}(t/2\|\mathbb{H}\|_\infty, \mathbb{Q}', \|\cdot\|_{\infty})}{n}} + \sqrt{\frac{\log \mathcal{N}(t/2\|\mathbb{Q}'\|_\infty, \mathbb{H}, \|\cdot\|_{\infty})}{n}}\,\mathrm{d}(t).\\
&= \begin{cases} O(n^{-1/2}\eta^{1-\frac{1}{2}\beta}), (\beta \geq 2)\\ O(n^{-1/2}\log(\eta)), (\beta = 2)\\ O(n^{-1/2}\eta^{-\beta+2}), (\beta \leq 2). \end{cases}
\end{aligned}
$$

Then we solve the following inequalities:

$$\begin{cases} n^{-1/2}\eta^{1-\frac{1}{2}\beta} & \leq \eta^2, \ (\beta \geq 2) \\ n^{-1/2}\log(\eta) & \leq \eta^2, (\beta = 2) \\ n^{-1/2}\eta^{-\beta+2} & \leq \eta^2, (\beta \leq 2). \end{cases}$$

which gives

$$\eta_h = \begin{cases} O(n^{-1/(2+\beta)}) & \beta < 2 \\ O(n^{-1/4}\log n) & \beta = 2 \\ O(n^{-1/(2\beta)}) & \beta > 2 \end{cases}$$

According to Lemma 45, the η_h above upper bounds the critical radius of $\mathcal{G}_h$. By a similar calculation, it can be shown that this η_h can also upper bound the critical radius of $\mathbb{Q}'$. Plugging it into Corollary 17 then proves the conclusion of Corollary 20. ■

I.7 Proofs for Section 7.2

Proof of Theorem 22. We denote

$$\eta_h' = \eta_h + \sqrt{1 + \log(1/\delta)/n}, \ \eta_q' = \eta_q + \sqrt{1 + \log(1/\delta)/n}.$$

Recall that we use sample splitting with two data subsamples $\mathcal{D}_1$ and $\mathcal{D}_0$, and the estimator is

$$\hat{J}_{\mathrm{DR}} = \frac{1}{2}\mathbb{E}_{n_1}[\tilde{\phi}_{\mathrm{DR}}(O; \hat{q}^{(0)}, \hat{h}^{(0)})] + \frac{1}{2}\mathbb{E}_{n_0}[\tilde{\phi}_{\mathrm{DR}}(O; \hat{q}^{(1)}, \hat{h}^{(1)})].$$

where $\mathbb{E}_{n_1}[\cdot]$ is the empirical average over $\mathcal{D}_1$, $\mathbb{E}_{n_0}[\cdot]$ is the empirical average over $\mathcal{D}_0$, and

$$\tilde{\phi}_{\mathrm{DR}}(O; h, q) = \pi(A \mid X)q(Z, A, X)\{Y - h(W, A, X)\} + (\mathcal{T}h)(W, X).$$

We take arbitrary elements h_0, q_0 s.t. $h_0 \in \mathbb{H}_0^{\mathrm{obs}}, q_0 \in \mathbb{Q}_0^{\mathrm{obs}}$ and $J = \mathbb{E}\left[\tilde{\phi}_{\mathrm{DR}}(O; q_0, h_0)\right]$ and denote $\hat{h} = \hat{h}^{(0)}$, $\hat{q} = \hat{q}^{(0)}$. Then, the bias is decomposed into the three terms:

$$\begin{aligned} \mathbb{E}_{n_1}[\tilde{\phi}_{\mathrm{DR}}(O; \hat{q}, \hat{h})] - J = (\mathbb{E}_{n_1} - \mathbb{E})[\tilde{\phi}_{\mathrm{DR}}(O; \hat{q}, \hat{h}) - \tilde{\phi}_{\mathrm{DR}}(O; q_0, h_0)] + \mathbb{E}[\tilde{\phi}_{\mathrm{DR}}(O; \hat{q}, \hat{h}) - \tilde{\phi}_{\mathrm{DR}}(O; q_0, h_0)] \\ + \mathbb{E}_{n_1}[\tilde{\phi}_{\mathrm{DR}}(O; q_0, h_0)] - J. \end{aligned} \tag{62}$$

According to Corollaries 17 and 19, we have that with probability at least $1 - \delta$,

$$\|P_z(\hat{h} - h_0)\|_2 \leq O(\eta_h'), \quad \|P_w(\hat{q} - q_0)\|_2 \leq O(\eta_q').$$

Given this event, Equation (34) implies that

$$\begin{aligned} &\|\mathbb{E}[\hat{h}(W, A, X) - Y \mid A, U, X]\|_2 = O(\tau_{1,n}^{\mathbb{H}}\eta_h'), \\ &\|\mathbb{E}[\pi(A \mid X)\{\hat{q}(Z, A, X) - 1/f(A \mid X, W)\} \mid A, U, X]\|_2 = O(\tau_{1,n}^{\mathbb{Q}}\eta_q'). \end{aligned}$$

In the rest of the proof, we always condition on this event. Now we bound each term in (62) respectively.

First Term. Note that $\|\tilde{\phi}_{\mathrm{DR}}(O;h,q)\|_\infty \le \|\pi\mathbb{Q}\|_\infty(\|Y\|_\infty + \|\mathbb{H}\|_\infty) + \|\mathbb{H}\|_\infty$. So by Bernstein inequality, the first term in (62) is $O(\sqrt{\log(1/\delta)/n_1})$ with probability at least $1-\delta$.
Second Term. To bound the second term in (62), we note that

$$\begin{aligned}
&\mathbb{E}[\tilde{\phi}_{\mathrm{DR}}(O;\hat{q},\hat{h}) - \tilde{\phi}_{\mathrm{DR}}(O;q_0,h_0)|\mathcal{D}_0] \\
&= \mathbb{E}[\pi(\hat{q}-q_0)(Y-h_0)|\mathcal{D}_0] + \mathbb{E}[\pi q_0\{-\hat{h}+h_0\} + \mathcal{T}(\hat{h}-h_0)|\mathcal{D}_0] + \mathbb{E}[\pi\{\hat{q}-q_0\}\{h_0-\hat{h}\}|\mathcal{D}_0] \\
&= \mathbb{E}[\pi\{\hat{q}-q_0\}\{h_0-\hat{h}\}|\mathcal{D}_0] \\
&= \mathbb{E}[\mathbb{E}[\pi\{\hat{q}-q_0\} \mid A,U,X]\mathbb{E}[\{h_0-\hat{h}\} \mid A,U,X]|\mathcal{D}_0] && (Z \perp W \mid A,U,X) \\
&\le \|\mathbb{E}[\hat{h}(W,A,X) - Y \mid A,U,X]\|_2 \|\mathbb{E}[\pi(A \mid X)\{\hat{q}(Z,A,X) - 1/f(A \mid X,W)\} \mid A,U,X]\|_2 \\
& && \text{(CS inequality)} \\
&= O(\tau_{1,n}^{\mathbb{H}}\tau_{1,n}^{\mathbb{Q}}\eta_h'\eta_q').
\end{aligned}$$

Third Term. Again, we can use Bernstein inequality to bound the third term by $O(\sqrt{\log(1/\delta)/n_1})$ with probability at least $1-\delta$.
Combining All Terms. The above proves that with probability at least $1-\delta$, we have

$$|\mathbb{E}_{n_1}[\tilde{\phi}_{\mathrm{DR}}(O;\hat{q}^{(0)},\hat{h}^{(0)})] - J| = O\left(\tau_1^{\mathbb{H}}\tau_1^{\mathbb{Q}}\eta_h\eta_q + \sqrt{\{1+\log(1/\delta)\}/n}\right).$$

Similarly, we can prove the same bound for $|\mathbb{E}_{n_0}[\tilde{\phi}_{\mathrm{DR}}(O;\hat{q}^{(1)},\hat{h}^{(1)})] - J|$. Combining these two proves Equation (36). ∎

Proof of Theorem 23. Again, according to Corollaries 17 and 19, we have that with probability at least $1-2\delta$,

$$\|P_z(\hat{h}-h_0)\|_2 \le O(\eta_h'), \quad \|P_w(\hat{q}-q_0)\|_2 \le O(\eta_q').$$

Given this event, the assumption (35) implies that

$$\|\hat{h}-h_0\|_2 \le O(\tau_{2,n}^{\mathbb{H}}\eta_h'),\ \|\hat{q}-q_0\|_2 \le O(\tau_{2,n}^{\mathbb{Q}}\eta_q').$$

We again condition on this event in the rest of the proof, and analyze each term in Equation (62) respectively.
First Term. To bound the first term in Equation (62), we first note that

$$\begin{aligned}
&(\mathbb{E}_{n_1}-\mathbb{E})[\phi_{\mathrm{DR}}(\hat{q},\hat{h}) - \phi_{\mathrm{DR}}(q_0,h_0)|\mathcal{D}_0] && (63)\\
&= (\mathbb{E}_{n_1}-\mathbb{E})[\pi(\hat{q}-q_0)(Y-h_0)|\mathcal{D}_0] && (64)\\
&+ (\mathbb{E}_{n_1}-\mathbb{E})[\pi q_0\{-\hat{h}+h_0\} + \mathcal{T}(\hat{h}-h_0)|\mathcal{D}_0] && (65)\\
&+ (\mathbb{E}_{n_1}-\mathbb{E})[\pi\{\hat{q}-q_0\}\{h_0-\hat{h}\}|\mathcal{D}_0]. && (66)
\end{aligned}$$

From Bernstein's inequality, with probability $1-\delta$, Equation (64) is

$$\begin{aligned}
&|(\mathbb{E}_{n_1}-\mathbb{E})[\pi(\hat{q}-q_0)(Y-h_0)|\mathcal{D}_0]| \\
&\lesssim \sqrt{\frac{2\mathbb{E}[\{\pi(\hat{q}-q_0)\}^2\{Y-h_0\}^2|\mathcal{D}_0]\log(1/\delta)}{n_1}} + \frac{2\|\pi(\hat{q}-q_0)(Y-h_0)\|_\infty \log(1/\delta)}{3n_1}
\end{aligned}$$

$$= O\left(\sqrt{\frac{(\tau_{2,n}^{\mathbb{Q}}\eta_q')^2\log(1/\delta)}{n_1}} + \frac{\log(1/\delta)}{n_1}\right),$$

and similarly Equations (65) and (66) are

$$|(\mathbb{E}_{n_1} - \mathbb{E})[q_0\{-\hat{h} + h_0\} + \mathcal{T}(\hat{h} - h_0)|\mathcal{D}_0]| = O\left(\sqrt{\frac{(\tau_{2,n}^{\mathbb{H}}\eta_h')^2\log(1/\delta)}{n_1}} + \frac{\log(1/\delta)}{n_1}\right)$$

$$(\mathbb{E}_{n_1} - \mathbb{E})[\pi\pi\{\hat{q} - q_0\}\{h_0 - \hat{h}\}|\mathcal{D}_0] = O\left(\sqrt{\frac{(\tau_{2,n}^{\mathbb{H}}\eta_h')^2\log(1/\delta)}{n_1}} + \frac{\log(1/\delta)}{n_1}\right).$$

It follows that with probability at least $1 - 3\delta$,

$$|(\mathbb{E}_{n_1} - \mathbb{E})[\pi(\hat{q} - q_0)(Y - h_0)|\mathcal{D}_0]| = O\left(\sqrt{\frac{\log(1/\delta)}{n}}\left(\tau_{2,n}^{\mathbb{H}}\eta_h' + \tau_{2,n}^{\mathbb{Q}}\eta_q'\right) + \frac{\log(1/\delta)}{n}\right)$$

Second Term. To bound the second term in (62), we note that

$$\begin{aligned}
&\mathbb{E}[\phi_{\mathrm{DR}}(\hat{q}, \hat{h}) - \phi_{\mathrm{DR}}(q_0, h_0)|\mathcal{D}_0] \\
&= \mathbb{E}[\pi(\hat{q} - q_0)(y - h_0)|\mathcal{D}_0] + \mathbb{E}[\pi q_0\{-\hat{h} + h_0\} + \mathcal{T}(\hat{h} - h_0)|\mathcal{D}_0] + \mathbb{E}[\pi\{\hat{q} - q_0\}\{h_0 - \hat{h}\}|\mathcal{D}_0] \\
&= \mathbb{E}[\pi\{\hat{q} - q_0\}\{h_0 - \hat{h}\}|\mathcal{D}_0].
\end{aligned}$$

Thus we have

$$|\mathbb{E}[\phi_{\mathrm{DR}}(\hat{q}, \hat{h}) - \phi_{\mathrm{DR}}(q_0, h_0)|\mathcal{D}_0]| = |\mathbb{E}[\pi\{\hat{q} - q_0\}\mathbb{E}[\{h_0 - \hat{h}\} \mid Z, A, X]|\mathcal{D}_0]| \leq \|\{q_0 - \hat{q}\}\pi\|_2\|P_z(h_0 - \hat{h})\|_2,$$

and

$$|\mathbb{E}[\phi_{\mathrm{DR}}(\hat{q}, \hat{h}) - \phi_{\mathrm{DR}}(q_0, h_0)|\mathcal{D}_0]| = |\mathbb{E}[\mathbb{E}\left[\pi\{\hat{q} - q_0\} \mid W, A, X\right]\{h_0 - \hat{h}\}|\mathcal{D}_0]| \leq \|\pi P_w\{q_0 - \hat{q}\}\|_2\|h_0 - \hat{h}\|_2.$$

It follows that

$$\mathbb{E}[\phi_{\mathrm{DR}}(\hat{q}, \hat{h}) - \phi_{\mathrm{DR}}(q_0, h_0)|\mathcal{D}_0] = O\left(\min(\tau_{2,n}^{\mathbb{H}}, \tau_{2,n}^{\mathbb{Q}})\eta_h'\eta_q'\right).$$

Combining All Terms. The above proves that with probability at least $1 - \tilde{c}\delta$ for a universal positive constant $\tilde{c}$, we have

$$\begin{aligned}
&|\mathbb{E}_{n_1}[\tilde{\phi}_{\mathrm{DR}}(O; \hat{q}^{(0)}, \hat{h}^{(0)})] - \mathbb{E}_{n_1}[\tilde{\phi}_{\mathrm{DR}}(O; q_0, h_0)]| \\
&\qquad = O\left(\sqrt{\frac{\log(1/\delta)}{n}}\left(\tau_{2,n}^{\mathbb{H}}\eta_h' + \tau_{2,n}^{\mathbb{Q}}\eta_q'\right) + \frac{\log(1/\delta)}{n} + \min(\tau_{2,n}^{\mathbb{H}}, \tau_{2,n}^{\mathbb{Q}})\eta_h'\eta_q'\right).
\end{aligned}$$

Similarly, we can show

$$\begin{aligned}
&|\mathbb{E}_{n_0}[\tilde{\phi}_{\mathrm{DR}}(O; \hat{q}^{(1)}, \hat{h}^{(1)})] - \mathbb{E}_{n_0}[\tilde{\phi}_{\mathrm{DR}}(O; q_0, h_0)]| \\
&\qquad = O\left(\sqrt{\frac{\log(1/\delta)}{n}}\left(\tau_{2,n}^{\mathbb{H}}\eta_h' + \tau_{2,n}^{\mathbb{Q}}\eta_q'\right) + \frac{\log(1/\delta)}{n} + \min(\tau_{2,n}^{\mathbb{H}}, \tau_{2,n}^{\mathbb{Q}})\eta_h'\eta_q'\right).
\end{aligned}$$

Given that $\min(\tau_2^{\mathbb{H}}, \tau_2^{\mathbb{Q}})\eta_q\eta_h = o(n^{-1/2})$, $\tau_2^{\mathbb{H}}\eta_h = o(1)$ and $\tau_2^{\mathbb{Q}}\eta_q = o(1)$, we have

$$\hat{J} - J = \mathbb{E}_n[\tilde{\phi}_{\mathrm{DR}}(O; q_0, h_0)] - J + o_p\left(n^{-1/2}\right).$$

By the central limit theorem, we have

$$\sqrt{n}(\hat{J}_{\mathrm{DR}} - J) \xrightarrow{d} \mathcal{N}\left(0, \mathbb{E}\left[\left(\tilde{\phi}_{\mathrm{DR}}(O; h_0, q_0) - J\right)^2\right]\right).$$

■

I.8 Proofs for Appendix B

Proof for Lemma 24. According to Lemma 3, we have that $\mathbb{H}_0 \subseteq \mathbb{H}_0^{\mathrm{obs}}$. For any $h_0 \in \mathbb{H}_0^{\mathrm{obs}} \subseteq \mathbb{H}$,

$$\mathbb{E}\left[Y - h_0(W, A, X) \mid Z, A, X\right] = \mathbb{E}\left[\mathbb{E}\left[Y - h_0(W, A, X) \mid A, U, X\right] \mid Z, A, X\right] = 0.$$

By Assumption 2 condition (2), we have that $\mathbb{E}\left[Y - h_0(W, A, X) \mid A, U, X\right] = 0$, *i.e.*, $h_0 \in \mathbb{H}_0$. It follows that $\mathbb{H}_0^{\mathrm{obs}} = \mathbb{H}_0$. Similarly, we can prove that $\mathbb{Q}_0^{\mathrm{obs}} = \mathbb{Q}_0$ under Assumption 2 condition (1). ■

I.9 Proofs for Appendix C

Proof for Lemma 27. Consider two bridge functions $h_0, h_0' \in \mathbb{H}_0$:

$$\mathbb{E}\left[Y - h_0(W, A, X) \mid U, A, X\right] = \mathbb{E}\left[Y - h_0'(W, A, X) \mid U, A, X\right] = 0.$$

Thus we have

$$\mathbb{E}\left[h_0(W, A, X) - h_0'(W, A, X) \mid U, A, X\right] = 0.$$

Then Assumption 3 condition (1) implies that $h_0(W, A, X) = h_0'(W, A, X)$. Therefore, $\mathbb{H}_0$ is at most a singleton. We can similarly prove that $\mathbb{Q}_0$ is at most a singleton. ■

Proof for Lemma 28. Under Assumption 1, we have that

$$\mathbb{E}\left[g(W, A, X) \mid Z, A, X\right] = \mathbb{E}\left[\mathbb{E}\left[g(W, A, X) \mid U, A, X\right] \mid Z, A, X\right],$$
$$\mathbb{E}\left[g(Z, A, X) \mid W, A, X\right] = \mathbb{E}\left[\mathbb{E}\left[g(Z, A, X) \mid U, A, X\right] \mid W, A, X\right].$$

We first prove statement (1). According to Assumption 2 condition (2),

$$\mathbb{E}\left[g(W, A, X) \mid Z, A, X\right] = 0 \quad \text{if and only if} \quad \mathbb{E}\left[g(W, A, X) \mid U, A, X\right] = 0.$$

Assumption 3 condition (1) further ensures that this holds if and only if $g(W, A, X) = 0$. In other words, Assumption 2 conditions (2) and Assumption 3 condition (1) are sufficient for Assumption 4 condition (1). Similarly, we can show that Assumption 2 condition (1) and Assumption 3 condition (2) are sufficient for Assumption 4 condition (2).

Next, we prove statement (2). If

$$\mathbb{E}\left[g(W, A, X) \mid U, A, X\right] = 0,$$

then $\mathbb{E}[g(W,A,X) \mid Z,A,X] = 0$ as well. By Assumption 4 condition (1), this holds if and only if $g(W,A,X) = 0$. Therefore, Assumption 4 condition (1) is sufficient for Assumption 3 condition (1). Similarly, we can prove that Assumption 4 condition (2) is sufficient for Assumption 3 condition (2). ■

Proof for Lemma 31. We need to prove that the following equation of the first kind is solvable:

$$[K_{W|a,x}h](a,u,x) = k_0(a,u,x), \quad \text{a.e. } u,a,x \text{ w.r.t } \mathbb{P}.$$

Thus we only need to verify the assumptions in Picard's theorem in Lemma 30 with $K = K_{W|a,x}$ and $\phi = k_0$. Note that condition 2 in Lemma 30 is satisfied by our asserted assumptions. Thus we only need to show $k_0 \in \mathcal{N}\left(K^*_{W|a,x}\right)^\perp$.

Since $(K^*_{W|a,x}g)(w,a,x) = \mathbb{E}[g(U,a,x) \mid W = w, A = a, X = x]$. By Assumption 2 condition 1, $\mathcal{N}\left(K^*_{W|a,x}\right) = \{0\}$, which means that $\mathcal{N}\left(K^*_{W|a,x}\right)^\perp = \operatorname{dom} K^*_{W|a,x} = L_2(U \mid A = a, X = x)$. Therefore, $k_0 \in \mathcal{N}\left(K^*_{W|a,x}\right)^\perp$. ■

Proof for Lemma 32. The proof is completely analogous to the proof for Lemma 31. ■

I.10 Proofs for Appendix D

Before proving Theorem 33, we first introduce a generalized implicit function theorem below.

Lemma 51 (Implicit Function Theorem, Theorem 7.13-1 in Ciarlet (2013))**.** *Let $\mathcal{X}$ and $\mathcal{Y}$ be Banach spaces, $O \subset \mathcal{X} \times \mathcal{Y}$ be an open neighborhood containing $(\bar{x},\bar{y})$. Consider a mapping F satisfies the following conditions:*

1. *$F(\bar{x},\bar{y}) = 0$.*
2. *For any $(x,y) \in O$, F is continuous at (x,y), and $\frac{\partial F}{\partial y}(x,y)$ exists and is a continuous linear mapping over O.*
3. *The linear mapping $\frac{\partial F}{\partial y}(\bar{x},\bar{y})$ is a bijection.*
4. *The mapping F is differentiable at $(\bar{x},\bar{y})$.*

Then there exists an open neighborhood U of $\bar{x}$ in $\mathcal{X}$, an open neighborhood V of $\bar{y}$ in $\mathcal{Y}$ and a continuous implicit function $f : U \mapsto V$ such that $U \times V \subseteq O$,

$$F(x, f(x)) = 0, \quad \text{for any } x \in U,$$

and f is differentiable at $\bar{x}$ with derivative

$$f'(\bar{x}) = -\left(\frac{\partial F}{\partial y}(\bar{x},\bar{y})\right)^{-1} \frac{\partial F}{\partial x}(\bar{x},\bar{y}).$$

Proof for Theorem 33. **Step I: deriving the tangent space.** First, consider regular parametric submodel $\mathcal{P}_t = \{f_t(y,w,z,a,x) : t \in \mathbb{R}^s\}$ with $f_0(y,w,z,a,x)$ equals the true density $f(y,w,z,a,x)$ (with respect to an appropriate dominating measure). The associated

score function is denoted as $S(y,w,z,a,x) = \partial_t \log f_t(y,w,z,a,x)|_{t=0}$. The expectation w.r.t the distribution $f_t(y,w,z,a,x)$ is denoted by $\mathbb{E}_t$. We can similarly denote the score functions for any component of this parametric submodel. For example, the score function for $f_t(y,w \mid z,a,x)$ is denoted as $S(y,w \mid z,a,x) = \partial_t \log f_t(y,w \mid z,a,x)|_{t=0}$. It is easy to show that

$$S(Y,W,Z,A,X) = S(Z,A,X) + S(Y,W \mid Z,A,X), \mathbb{E}[S(Z,A,X)] = 0, \mathbb{E}[S(Y,W \mid Z,A,X) \mid Z,A,X] = 0.$$

Let h_t be a curve such that $h_t|_{t=0} = h_0$ and $\partial_t h_t|_{t=0}$ exists, and

$$\mathbb{E}_t[Y - h_t(W,A,X) \mid Z,A,X] = 0. \tag{67}$$

Equation (67) implies that

$$\partial_t \mathbb{E}_t[Y - h_t(W,A,X) \mid Z,A,X]|_{t=0} = 0,$$

which in turn implies that

$$\mathbb{E}[(Y - h_0(W,A,X)) S(W,Y \mid Z,A,X) \mid Z,A,X] = \mathbb{E}[\partial_t h_t(W,A,X)|_{t=0} \mid Z,A,X]. \tag{68}$$

This means that $S(W,Y \mid Z,A,X)$ must satisfy that

$$\mathbb{E}[(Y - h_0(W,A,X)) S(W,Y \mid Z,A,X) \mid Z,A,X] \in \text{Range}(P_z).$$

Therefore, all score vectors under $\mathcal{M}_{np}$ must lie in the set $\text{Closure}(\mathcal{S})$ where:

$$\begin{aligned}\mathcal{S} = \Big\{ & S(Y,W,Z,A,X) = S(Z,A,X) + S(Y,W \mid Z,A,X) : \\ & S(Z,A,X) \in L_2(Z,A,X), S(Y,W \mid Z,A,X) \in L_2(Y,W \mid Z,A,X), \\ & \mathbb{E}[S(Z,A,X)] = 0, \mathbb{E}[S(Y,W \mid Z,A,X) \mid Z,A,X] = 0, \\ & \mathbb{E}[(Y - h_0(W,A,X)) S(Y,W \mid Z,A,X) \mid Z,A,X] \in \text{Range}(P_z) \Big\}. \end{aligned} \tag{69}$$

Now we show that for any $S(Y,W,Z,A,X) \in \mathcal{S}$, we can find a parametric submodel f_t whose score function is $S(Y,W,Z,A,X) = S(Z,A,X) + S(Y,W \mid Z,A,X)$ such that Equation (67) holds for a certain h_t for t near 0 and $\partial_t h_t(W,A,X)|_t$ exists. We consider the parametric submodel of the following form:

$$f_t(z,a,x) = f(z,a,x)(1 + tS(z,a,x)), \quad f_t(y,w \mid z,a,x) = f(y,w \mid z,a,x)(1 + tS(y,w \mid z,a,x)),$$

where $t \in \mathbb{R}$ is small enough such that $1 + tS(z,a,x) \geq 0$ and $1 + tS(y,w \mid z,a,x) \geq 0$ for any y,w,z,a,x so that f_t defined above is a valid density function. It is easy to check that $\partial_t \log f_t(z,a,x)|_{t=0} = S(z,a,x)$ and $\partial_t \log f_t(y,w \mid z,a,x)|_{t=0} = S(y,w \mid z,a,x)$.

Consider the mapping $F(t,h) = P_{z,t}h - \mathbb{E}_t[Y \mid Z,A,X]$ for $t \in \mathbb{R}$ and $h \in L_2(W,A,X)$, where $P_{z,t} : L_2(W,A,X) \mapsto L_2(Z,A,X)$ is the conditional expectation operator defined as

$$[P_{z,t}h](Z,A,X) = \mathbb{E}_t[h(W,A,X) \mid Z,A,X].$$

Note that $F(0,h_0) = 0$, $\partial_h F(0,h_0) = P_z$ and

$$\partial_t F(0,h_0) = -\mathbb{E}[(Y - h_0(W,A,X)) S(Y,W,Z,A,X) \mid Z,A,X].$$

Since P_z is a bijective linear operator, Lemma 51 implies that for all t close to 0, there exists $h_t \in L_2(W, A, X)$ such that $F(t, h_t) = 0$ and $\partial_t h_t|_{t=0}$ exists with

$$\partial_t h_t(W, A, X)|_{t=0} = -\left(\partial_h F(0, h_0)\right)^{-1} \partial_t F(0, h_0).$$

This is equivalent to $\partial_t h_t|_{t=0}$ satisfying Equation (68).

Therefore, we have shown that the tangent space for the model $\mathcal{M}_{np}$ is Closure $(\mathcal{S})$.

Step II: deriving a preliminary influence function. Denote the target parameter under distribution $f_t(y, w, z, a, x)$ as $J_t = \mathbb{E}_t\left[\mathcal{T} h_t(W, X)\right]$. We need to derive

$$\partial_t J_t|_{t=0} = \mathbb{E}\left[\mathcal{T} h_0(W, X) S(W, X)\right] + \mathbb{E}\left[\partial_t \mathcal{T} h_t(W, X)|_{t=0}\right]$$

Note that

$$\begin{aligned}
\mathbb{E}\left[\partial_t \mathcal{T} h_t(W, X)|_{t=0}\right] &= \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X) \mathbb{E}\left[\partial_t h_t(W, A, X)|_{t=0} \mid Z, A, X\right]\right] \\
&= \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X) \left(Y - h_0(W, A, X)\right) S(W, Y \mid Z, A, X)\right] \\
&= \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X) \left(Y - h_0(W, A, X)\right) S(W, Y \mid Z, A, X)\right] \\
&+ \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X) \left(Y - h_0(W, A, X)\right) S(Z, A, X)\right] \\
&= \mathbb{E}\left[\pi(A \mid X) q_0(Z, A, X) \left(Y - h_0(W, A, X)\right) S(W, Y, Z, A, X)\right].
\end{aligned}$$

Here the second equality follows from Equation (68), and the third equality follows from the fact that

$$\mathbb{E}\left[Y - h_0(W, A, X) \mid Z, A, X\right] = 0.$$

Moreover,

$$\begin{aligned}
\mathbb{E}\left[\mathcal{T} h_0(W, X) S(W, X)\right] &= \mathbb{E}\left[\mathcal{T} h_0(W, X) S(W, X)\right] + \mathbb{E}\left[\mathcal{T} h_0(W, X) S(Y, Z, A \mid W, X)\right] \\
&= \mathbb{E}\left[\mathcal{T} h_0(W, X) S(Y, W, Z, A, X)\right] \\
&= \mathbb{E}\left[\left(\mathcal{T} h_0(W, X) - J\right) S(Y, W, Z, A, X)\right].
\end{aligned}$$

Here the first equality follows from the fact that $\mathbb{E}\left[S(Y, Z, A \mid W, X) \mid W, X\right] = 0$, and the last equality follows from $\mathbb{E}\left[S(Y, Z, A, W, X) \mid W, X\right] = 0$.

Therefore, we have that

$$\partial_t J_t|_{t=0} = \mathbb{E}\left[\left(\pi(A \mid X) q_0(Z, A, X)[Y - h_0(W, A, X)] + \mathcal{T} h_0(W, X) - J\right) S(Y, W, Z, A, X)\right].$$

This means that $\pi(A \mid X) q_0(Z, A, X)[Y - h_0(W, A, X)] + \mathcal{T} h_0(W, X) - J$ is a valid influence function for J under the model $\mathcal{M}_{np}$.

Step III: verifying efficient influence function. Now we verify that $\pi(A \mid X) q_0(Z, A, X)[Y - h_0(W, A, X)] + \mathcal{T} h_0(W, X) - J$ also belongs to Closure $(\mathcal{S})$ so that it is also the efficient influence function for J relative to the tangent space Closure $(\mathcal{S})$.

First, note that we can decompose this influence function in the following way:

$$\begin{aligned}
\tilde{S}(Y, W, Z, A, X) &\coloneqq \pi(A \mid X) q_0(Z, A, X)[Y - h_0(W, A, X)] + \mathcal{T} h_0(W, X) - J \\
&= \tilde{S}(Z, A, X) + \tilde{S}(Y, W \mid Z, A, X),
\end{aligned}$$

where

$$\tilde{S}(Z, A, X) = \mathbb{E}\left[\mathcal{T} h_0(W, X) - J \mid Z, A, X\right],$$

$$\tilde{S}(Y,W \mid Z,A,X) = \mathcal{T}h_0(W,X) - \mathbb{E}\left[\mathcal{T}h_0(W,X) \mid Z,A,X\right] + \pi(A \mid X)q_0(Z,A,X)[Y - h_0(W,A,X)].$$

It is easy to show that $\mathbb{E}\left[\tilde{S}(Z,A,X)\right] = 0$ and $\mathbb{E}\left[\tilde{S}(Y,W \mid Z,A,X) \mid Z,A,X\right] = 0$. Since P_z is bijective, we have that

$$\mathbb{E}\left[(Y - h_0(W,A,X))\,\tilde{S}(Y,W \mid Z,A,X) \mid Z,A,X\right] \in \text{Closure}\left(\text{Range}\left(P_z\right)\right).$$

Therefore,

$$\text{EIF}(J) = \pi(A \mid X)q_0(Z,A,X)[Y - h_0(W,A,X)] + \mathcal{T}h_0(W,X) - J \in \text{Closure}\left(\mathcal{S}\right),$$

which means that $\text{EIF}(J)$ is the efficient influence function of J and $\mathbb{E}\left[\text{EIF}^2(J)\right]$ is the corresponding semiparametric efficiency bound. ■

I.11 Proofs for Appendix E

Proof for Lemma 34. The conclusion follows easily by simple algebra. ■

Proof of Lemma 35. Define $S_z : \mathbb{R}^n \to \mathcal{L}_z$ and $S_w : \mathbb{R}^n \to \mathcal{L}_w$ by

$$S_z\alpha = \sum_{i=1}^{n} \alpha_i k_z((Z_i,A_i,X_i),\cdot),$$
$$S_w\alpha = \sum_{i=1}^{n} \alpha_i k_w((W_i,A_i,X_i),\cdot).$$

Their adjoints are the corresponding evaluation maps, so $S_z^*S_z = K_{z,n}$ and $S_w^*S_w = K_{w1,n}$. Write

$$B_z = n\gamma_1 I_{\mathcal{L}_z} + \lambda S_z S_z^*, \qquad R_z = n\gamma_1 I_n + \lambda K_{z,n},$$
$$B_w = n\gamma_2 I_{\mathcal{L}_w} + \lambda S_w S_w^*, \qquad R_w = n\gamma_2 I_n + \lambda K_{w1,n}.$$

These operators and matrices are invertible because $\gamma_1, \gamma_2 > 0$. The identities $B_zS_z = S_zR_z$ and $B_wS_w = S_wR_w$ imply

$$B_z^{-1}S_z = S_zR_z^{-1}, \qquad B_w^{-1}S_w = S_wR_w^{-1}.$$

For a fixed outcome bridge h, let $\psi = \psi_n(h)$. The inner objective in Equation (39) is

$$\frac{1}{n}\left\{-2\langle S_z\psi, q\rangle_{\mathcal{L}_z} - \langle q, B_zq\rangle_{\mathcal{L}_z}\right\}.$$

Completing the square gives the maximizer $q^* = -B_z^{-1}S_z\psi$ and optimal value

$$\begin{aligned}\frac{1}{n}\langle S_z\psi, B_z^{-1}S_z\psi\rangle_{\mathcal{L}_z} &= \frac{1}{n}\psi^\top K_{z,n}R_z^{-1}\psi \\ &= \frac{1}{n}\psi^\top K_{z,n}^{1/2}R_z^{-1}K_{z,n}^{1/2}\psi.\end{aligned}$$

Dropping the positive factor $1/n$ proves Equation (41).

For the action bridge, the Riesz representation theorem gives $\mu_w \in \mathcal{L}_w$ representing the bounded linear functional $h \mapsto \sum_{j=1}^n (\mathcal{T}h)(W_j, X_j)$, so that

$$\langle \mu_w, h \rangle_{\mathcal{L}_w} = \sum_{j=1}^n (\mathcal{T}h)(W_j, X_j).$$

The reproducing property and the definition of $K_{w2,n}$ give

$$S_w^* \mu_w = K_{w2,n} \mathbf{1}_n.$$

For a fixed q, write $\phi = \phi_n(q)$. The inner objective in Equation (40) is

$$\frac{1}{n}\big\{2\langle S_w\phi - \mu_w, h\rangle_{\mathcal{L}_w} - \langle h, B_w h\rangle_{\mathcal{L}_w}\big\}.$$

Its maximizer is $h^* = B_w^{-1}(S_w\phi - \mu_w)$, which need not lie in the range of S_w. The optimal value is

$$\begin{aligned}&\frac{1}{n}\langle S_w\phi - \mu_w, B_w^{-1}(S_w\phi - \mu_w)\rangle_{\mathcal{L}_w} \\ &\quad = \frac{1}{n}\big\{\phi^\top K_{w1,n} R_w^{-1}\phi - 2\phi^\top R_w^{-1} K_{w2,n}\mathbf{1}_n + C_w\big\},\end{aligned}$$

where $C_w = \langle \mu_w, B_w^{-1}\mu_w\rangle_{\mathcal{L}_w}$ does not depend on q. Since $K_{w1,n}$ commutes with R_w^{-1}, dropping C_w and the positive factor $1/n$ proves Equation (42). The argument does not require either Gram matrix to be nonsingular. ■

I.12 Proofs for Appendix F

We prove the second statement. The first statement is similarly proved.

We only need to prove

$$\inf_{h\in\mathcal{H}'} \sup_{q\in\mathcal{Q}} |\mathbb{E}[(q - q_0)\pi(h - h_0)]| = 0.$$

Here, letting $\langle \alpha_1, \phi(\cdot)\rangle = (q - q_0)$ and $h(\cdot) = \langle \beta_h, \psi\rangle$,

$$\mathbb{E}[(q - q_0)\pi(h - h_0)] = \alpha_1^\top \mathbb{E}[\tilde{\phi}(Z, A, X)\{\psi^\top(W, X, A)\beta_h - h_0\}].$$

Hence, there exists β_h such that

$$0 = \alpha_1^\top \mathbb{E}[\tilde{\phi}(Z, A, X)\{\psi^\top(W, X, A)\beta_h - h_0\}]$$

when $\mathbb{E}[\tilde{\phi}(Z, A, X)\psi^\top(W, X, A)]$ is full column rank. This shows

$$\inf_{h\in\mathcal{H}'} \sup_{q\in\mathcal{Q}} |\mathbb{E}[(q - q_0)\pi(h - h_0)]| = 0.$$

I.13 Proofs for Appendix G

Proof of Corollary 37. We prove the result for $\hat{J}_{\mathrm{IPW}}$. The result for $\hat{J}_{\mathrm{REG}}$ can be proved analogously so is omitted here.

Fix some element in $h_0 \in \mathbb{H}_0^{\mathrm{obs}}, \pi q_0 \in \pi\mathbb{Q}_0^{\mathrm{obs}}$. Then, from Theorem 7, with $1-\delta$, we have the following for a universal constant $c > 0$:

$$|\hat{J}_{\mathrm{IPW}} - J| \le c\Bigg\{\mathcal{R}(\infty;\mathbb{Q}) + \mathcal{R}(\infty;\mathbb{H}) + \mathcal{R}(\infty;\mathcal{T}\mathbb{H}) + \inf_{h\in\mathbb{H}'}\sup_{q\in\mathbb{Q}}|\mathbb{E}[(q-q_0)\pi(h_0-h)]| + \inf_{q\in\mathbb{Q}}\sup_{h\in\mathbb{H}'}|\mathbb{E}[(q_0-q)\pi h]| + \sqrt{\frac{\log(1/\delta)}{n}}\Bigg\}.$$

The first three terms are upper bounded by $O(\sqrt{k_n/n})$ since the VC dimension of $\mathbb{Q} = \mathcal{S}_{2,n}$ and $\mathbb{H}' = \mathcal{S}_{1,n}$ are both k_n. Besides, from the assumptions (44) and (45),

$$\inf_{h\in\mathbb{H}'}\sup_{q\in\mathbb{Q}}|\mathbb{E}[(q-q_0)\pi(h_0-h)]| \le 2\|\mathcal{S}_{2,n}\|_\infty\|\pi\left(A \mid X\right)\|_\infty \inf_{h\in\mathbb{H}'}\mathbb{E}[|h_0-h|] = O(k_n^{-\alpha/d}),$$

$$\inf_{q\in\mathbb{Q}}\sup_{h\in\mathbb{H}'}|\mathbb{E}[\pi(q_0-q)h]| \le 2\|\mathcal{S}_{1,n}\|_\infty\|\pi\left(A \mid X\right)\|_\infty \inf_{q\in\mathbb{Q}'}\mathbb{E}[|q-q_0|] = O(k_n^{-\alpha/d}).$$

In the end, the final error becomes

$$O(k_n^{-\alpha/d} + \sqrt{k_n/n} + \sqrt{\log(1/\delta)/n}),$$

where the second and third term are statistical error terms, which is derived in the proof of Corollary 11. By setting $k_n = O(n^{\frac{d}{d+2\alpha}})$, we obtain the optimal rate $O(n^{-\alpha/(2\alpha+d)} + \sqrt{\log(1/\delta)/n})$. ∎

Proof of Corollary 38. According to Theorem 16, for some $h' \in \mathbb{H}$ to be specified later and $h_0 \in \mathbb{H}_0^{\mathrm{obs}} \cap \Lambda^\alpha([0,1]^d)$, we have that with probability $1-\delta$,

$$\|P_z(\hat{h}-h_0)\|_2 = O\left(\eta_h' + \epsilon_n + \frac{\|P_z(h'-h_0)\|_2^2}{\eta_h'} + \|P_z(h'-h_0)\|_2\right),$$

where $\eta_h' = \eta_h + c_0\sqrt{\log(c_1/\delta)/n}$, and η_h is the maximum of critical radii of $\mathbb{Q}' = \mathcal{S}_{2,n}$ and $\mathcal{G}_h$, and $\epsilon_n = \sup_{h\in\mathbb{H}}\inf_{q\in\mathbb{Q}'}\|q - P_z(h-h')\|_2$.

First, from Corollary 20, we have

$$\eta_h' = O\left(\sqrt{k_n}\frac{\log n}{\sqrt{n}}\log(\sqrt{k_n}\frac{\log n}{\sqrt{n}}) + \sqrt{(1+\log(1/\delta))/n}\right).$$

Next, according to Equation (44), we can find h' such that $\|h_0 - h'\| = O(k_n^{-\alpha/d})$. It follows that

$$\|P_z(h'-h_0)\| \le \|h'-h_0\| = O(k_n^{-\alpha/d}).$$

Besides, $P_z(\Lambda^\alpha([0,1]^d) - h_0) \subset \Lambda^\alpha([0,1]^d)$ implies that $P_z(h-h_0) \in \Lambda^\alpha([0,1]^d)$ for any $h \in \mathbb{H}$. According to Equation (45), for any $h \in \mathbb{H}$, we can find $q \in \mathbb{Q}'$ such that $\|q - P_z(h-h_0)\| = O(k_n^{-\alpha/d})$ and thus

$$\|q - P_z(h-h')\| \le \|q - P_z(h-h_0)\| + \|P_z(h'-h_0)\| = O(k_n^{-\alpha/d}).$$

This implies that $\epsilon_n = O(k_n^{-\alpha/d})$.

Combining all terns above, we have

$$\|P_w(h_0 - \hat{h})\|_2 = \tilde{O}\left(\sqrt{k_n/n} + k_n^{-\alpha/d} + \frac{k_n^{-2\alpha/d}}{\sqrt{k_n/n}} + \sqrt{(1+\log(1/\delta))/n}\right).$$

By setting $k_n = O(n^{\frac{d}{d+2\alpha}})$, we can obtain the optimal rate $\tilde{O}(n^{-\alpha/(2\alpha+d)})$. ■

Proof of Corollary 40. We prove the result for $\hat{J}_{\text{IPW}}$. The result for $\hat{J}_{\text{REG}}$ can be proved analogously so is omitted here.

Fix some element in $h_0 \in \mathbb{H}_0^{\text{obs}}, \pi q_0 \in \pi\mathbb{Q}_0^{\text{obs}}$. Then, from Theorem 7, with $1-\delta$, we have the following for a universal constant $c > 0$:

$$|\hat{J}_{\text{IPW}} - J| \leq c\Bigg\{\mathcal{R}(\infty;\mathbb{Q}) + \mathcal{R}(\infty;\mathbb{H}) + \mathcal{R}(\infty;\mathcal{T}\mathbb{H}) + \inf_{h\in\mathbb{H}'}\sup_{q\in\mathbb{Q}}|\mathbb{E}[(q-q_0)\pi(h_0-h)]| + \inf_{q\in\mathbb{Q}}\sup_{h\in\mathbb{H}'}|\mathbb{E}[(q_0-q)\pi h]| + \sqrt{\frac{\log(1/\delta)}{n}}\Bigg\}.$$

By Corollary 21 and lemma 48, we can compute the covering number of neural networks and upper bound the first three terms by $\tilde{O}(\sqrt{\Omega L/n})$.

According to Lemma 39,

$$\inf_{h\in\mathbb{H}'}\sup_{q\in\mathbb{Q}}|\mathbb{E}[(q-q_0)\pi(h_0-h)]| \leq 2\|\mathcal{F}_{2,n}\|_\infty\|\pi\left(A \mid X\right)\|_\infty \inf_{h\in\mathbb{H}'}\mathbb{E}[|h_0-h|] = \tilde{O}(\Omega^{-\alpha/d}),$$

$$\inf_{q\in\mathbb{Q}}\sup_{h\in\mathbb{H}'}|\mathbb{E}[\pi(q_0-q)h]| \leq 2\|\mathcal{F}_{1,n}\|_\infty\|\pi\left(A \mid X\right)\|_\infty \inf_{q\in\mathbb{Q}'}\mathbb{E}[|q-q_0|] = \tilde{O}(\Omega^{-\alpha/d}).$$

Combining the results above, the total error can be bounded by

$$\tilde{O}(\sqrt{\Omega L/n} + \Omega^{-\alpha/d} + \sqrt{\log(1/\delta)/n}).$$

By setting $L = \Theta(\log(n)), \Omega = \Theta(n^{d/(2\alpha+d)})$, we can obtain the optimal rate as $\tilde{O}(n^{-\alpha/(2\alpha+d)})$. ■

Proof of Corollary 41. According to Theorem 16, for some $h' \in \mathbb{H}$ to be specified later and $h_0 \in \mathbb{H}_0^{\text{obs}} \cap \Lambda^\alpha([0,1]^d)$, we have that with probability $1-\delta$,

$$\|P_z(\hat{h} - h_0)\|_2 = O\left(\eta_h' + \epsilon_n + \frac{\|P_z(h'-h_0)\|_2^2}{\eta_h'} + \|P_z(h'-h_0)\|_2\right),$$

where $\eta_h' = \eta_h + c_0\sqrt{\log(c_1/\delta)/n}$, and η_h is the maximum of critical radii of $\mathbb{Q}' = \mathcal{F}_{2,n}$ and $\mathcal{G}_h$, and $\epsilon_n = \sup_{h\in\mathbb{H}}\inf_{q\in\mathbb{Q}'}\|q - P_z(h-h')\|_2$.

First, from Corollary 20 and lemma 48, we have

$$\eta_h' = \tilde{O}(\sqrt{\Omega L/n} + c_0\sqrt{\log(c_1/\delta)/n}).$$

Next, according to Lemma 39, we can find h' such that $\|h_0 - h'\|_\infty = \tilde{O}(\Omega^{-\alpha/d})$. It follows that

$$\|P_z(h'-h_0)\| \leq \|h'-h_0\| = \tilde{O}(\Omega^{-\alpha/d}).$$

Besides, $P_z(\mathcal{S}^\alpha([0,1]^d) - h_0) \subset \mathcal{S}^\alpha([0,1]^d)$ implies that $P_z(h - h_0) \in \mathcal{S}^\alpha[0,1]^d)$ for any $h \in \mathbb{H}$. According to Lemma 39, for any $h \in \mathbb{H}$, we can find $q \in \mathbb{Q}'$ such that $\|q - P_z(h - h_0)\| = \tilde{O}(\Omega^{-\alpha/d})$ and thus

$$\|q - P_z(h - h')\| \leq \|q - P_z(h - h_0)\| + \|P_z(h' - h_0)\| = \tilde{O}(\Omega^{-\alpha/d}).$$

This implies that $\epsilon_n = \tilde{O}(\Omega^{-\alpha/d})$.

Combining all terns above, we have

$$\|P_w(h_0 - \hat{h})\|_2 = \tilde{O}\left(\sqrt{\Omega L/n} + c_0\sqrt{\log(c_1/\delta)/n} + \Omega^{-\alpha/d} + \frac{\Omega^{-2\alpha/d}}{\sqrt{\Omega L/n}}\right).$$

By setting $L = \Theta(\log(n)), \Omega = \Theta(n^{d/(2\alpha+d)})$, we can obtain the optimal rate as $\tilde{O}(n^{-\alpha/(2\alpha+d)})$. ■

References

Chunrong Ai and Xiaohong Chen. Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71:1795–1843, 2003.

Francis Bach. On the equivalence between kernel quadrature rules and random feature expansions. *Journal of Machine Learning Research*, 18(21):1–38, 2017.

H. Bang and J. M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61:962–973, 2005.

Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. *Ann. Statist.*, 33(4):1497–1537, 2005.

Andrew Bennett and Nathan Kallus. The variational method of moments. *arXiv preprint arXiv:2012.09422*, 2020.

Andrew Bennett, Nathan Kallus, and Tobias Schnabel. Deep generalized method of moments for instrumental variable analysis. In *Advances in Neural Information Processing Systems 32*, pages 3564–3574. 2019.

Richard Blundell, Xiaohong Chen, and Dennis Kristensen. Semi-nonparametric iv estimation of shape-invariant engel curves. *Econometrica*, 75(6):1613–1669, 2003.

Marine Carrasco, Jean-Pierre Florens, and Eric Renault. Chapter 77 linear inverse problems in structural econometrics estimation based on spectral decomposition and regularization. In *Handbook of Econometrics*, volume 6, pages 5633–5751. Elsevier B.V, 2007.

Xiaohong Chen. Chapter 76 large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics*, 6:5549–5632, 2007.

Xiaohong Chen and Demian Pouzo. Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. *Econometrica*, 80(1):277–321, 2012.

Xiaohong Chen and Demian Pouzo. Sieve wald and qlr inferences on semi/nonparametric conditional moment models. *Econometrica*, 83(3):1013–1079, 2015.

Xiaohong Chen and Yin Jia Jeff Qiu. Methods for nonparametric and semiparametric regressions with endogeneity: A gentle guide. *Annual Review of Economics*, (8):259–290, 2016.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21:C1–C68, 2018.

Victor Chernozhukov, Whitney Newey, Rahul Singh, and Vasilis Syrgkanis. Adversarial estimation of riesz representers. *arXiv preprint arXiv:2101.00009*, 2020.

Philippe G Ciarlet. *Linear and nonlinear functional analysis with applications*, volume 130. Siam, 2013.

Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Semiparametric proximal causal inference. *Journal of the American Statistical Association*, pages 1–12, 2023.

Serge Darolles, Yanqin Fan, Jean-Pierre Florens, and Eric Renault. Nonparametric instrumental regression. *Econometrica*, 79(5):1541–1565, 2010.

Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. *ICLR*, 2018.

Ben Deaner. Proxy controls and panel data. *arXiv preprint arXiv:1810.00283*, 2021.

Nishanth Dikkala, Greg Lewis, Lester Mackey, and Vasilis Syrgkanis. Minimax estimation of conditional moment models. In *Advances in Neural Information Processing Systems*, volume 33, pages 12248–12262, 2020.

Miroslav Dudik, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 29:485–511, 2014.

Xavier D'Haultfoeuille. On the completeness condition in nonparametric instrumental problems. *Econometric Theory*, 27(3):460–471, 2011.

Yihao Feng, Lihong Li, and Qiang Liu. A kernel loss for solving the bellman equation. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Dylan J. Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019.

AmirEmad Ghassami, Andrew Ying, Ilya Shpitser, and Eric Tchetgen Tchetgen. Minimax kernel machine learning for a class of doubly robust functionals. 2021.

Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In *Conference On Learning Theory*, pages 297–299. PMLR, 2018.

Rémi Gribonval, Gitta Kutyniok, Morten Nielsen, and Felix Voigtlaender. Approximation spaces of deep neural networks. *Constructive Approximation*, pages 1–109, 2021.

Peter Hall and Joel L. Horowitz. Nonparametric methods for inference in the presence of instrumental variables. *The Annals of Statistics*, 33(6):2904 – 2929, 2005.

Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50:1029–1054, 1982.

Jason Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Deep IV: A flexible approach for counterfactual prediction. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1414–1423, 2017.

K. Hirano, G. Imbens, and G. Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71:1161–1189, 2003.

Keisuke Hirano and Guido W. Imbens. The propensity score with continuous treatments. In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, chapter 7, pages 73–84. John Wiley & Sons, Ltd, 2004.

David Hirshberg and Stefan Wager. Augmented minimax linear estimation. *arXiv preprint arXiv:1712.00038*, 2019.

Yingyao Hu and Ji-Liang Shiu. Nonparametric identification using instrumental variables: Sufficient conditions for completeness. *Econometric Theory*, 34(3):659–693, 2011.

Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction.* 2015.

Prateek Jain and Purushottam Kar. Non-convex optimization for machine learning. *arXiv preprint arXiv:1712.07897*, 2017.

Nathan Kallus. Generalized optimal matching methods for causal inference. *J. Mach. Learn. Res.*, 21:62–1, 2020.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.

Keith Knight and Wenjiang Fu. Asymptotics for lasso-type estimators. *Annals of statistics*, pages 1356–1378, 2000.

Rainer Kress. *Linear Integral Equations*. Applied mathematical sciences (Springer-Verlag New York Inc.) ; v. 82. Springer New York : Imprint: Springer, New York, NY, 3rd ed. 2014. edition, 2014.

Thomas Kühn. Covering numbers of gaussian reproducing kernel hilbert spaces. *Journal of Complexity*, 27(5):489–499, 2011.

S. Lee and E. Bareinboim. Causal identification with matrix equations. *Columbia CausalAI Laboratory Technical Report (R-70)*, 2021.

Greg Lewis and Vasilis Syrgkanis. Adversarial generalized method of moments. *arXiv preprint arXiv:1803.07164*, 2018.

Luofeng Liao, You-Lin Chen, Zhuoran Yang, Bo Dai, Mladen Kolar, and Zhaoran Wang. Provably efficient neural estimation of structural equation models: An adversarial approach. In *Advances in Neural Information Processing Systems*, volume 33, pages 8947–8958, 2020.

Afsaneh Mastouri, Yuchen Zhu, Limor Gultchin, Anna Korba, Ricardo Silva, Matt Kusner, Arthur Gretton, and Krikamol Muandet. Proximal causal learning with kernels: Two-stage estimation and moment restriction. In *International Conference on Machine Learning*, pages 7512–7523. PMLR, 2021.

Shahar Mendelson. Improving the sample complexity using global data. *IEEE transactions on Information Theory*, 48(7):1977–1991, 2002.

Wang Miao, Zhi Geng, and Eric J Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993, 2018a.

Wang Miao, Xu Shi, and Eric Tchetgen Tchetgen. A confounding bridge approach for double negative control inference on causal effects. *arXiv preprint arXiv:1808.04945*, 2018b.

Krikamol Muandet, Arash Mehrjou, Si Kai Lee, and Anant Raj. Dual instrumental variable regression. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 2710–2721, 2020.

Ryumei Nakada and Masaaki Imaizumi. Adaptive approximation and generalization of deep neural network with intrinsic dimensionality. *Journal of Machine Learning Research*, 21 (174):1–38, 2020.

W. K. Newey and D. L Mcfadden. Large sample estimation and hypothesis testing. *Handbook of Econometrics*, IV:2113–2245, 1994.

Whitney K. Newey. Nonparametric instrumental variables estimation. *The American economic review*, 103(3):550–556, 2013.

Whitney K. Newey and James L. Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003.

Zhengling Qi, Rui Miao, and Xiaoke Zhang. Proximal learning for individualized treatment regimes under unmeasured confounding. 2021.

J. M. Robins, A. Rotnitzky, and L. P. Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89: 846–866, 1994.

Xu Shi, Wang Miao, Jennifer C Nelson, and Eric J Tchetgen Tchetgen. Multiply robust causal inference with double-negative control adjustment for categorical unmeasured confounding. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(2):521–540, 2020.

Rahul Singh. Kernel methods for unobserved confounding: Negative controls, proxies, and instruments. *arXiv preprint arXiv:2012.10315*, 2020.

Rahul Singh, Maneesh Sahani, and Arthur Gretton. Kernel instrumental variable regression. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Rahul Singh, Liyuan Xu, and Arthur Gretton. Kernel methods for policy evaluation: Treatment effects, mediation analysis, and off-policy planning. *arXiv preprint arXiv:2010.04855*, 2020.

Taiji Suzuki. Adaptivity of deep relu network for learning in besov and mixed smooth besov spaces: optimal rate and curse of dimensionality. *arXiv preprint arXiv:1810.08033*, 2018.

Zhiqiang Tan. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101:1619–1637, 2006.

Eric J Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal learning. *arXiv preprint arXiv:2009.10982*, 2020.

Guy Tennenholtz, Uri Shalit, and Shie Mannor. Off-policy evaluation in partially observable environments. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(06): 10276–10283, Apr. 2020.

Jin Tian. Identifying dynamic sequential plans. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*, UAI'08, page 554–561, 2008.

T. Tieleman and G. Hinton. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.

Masatoshi Uehara, Jiawei Huang, and Nan Jiang. Minimax weight and q-function learning for off-policy evaluation. *ICML 2020 (To appear)*, 2020.

Masatoshi Uehara, Masaaki Imaizumi, Nan Jiang, Nathan Kallus, Wen Sun, and Tengyang Xie. Finite sample analysis of minimax offline reinforcement learning: Completeness, fast rates and first-order efficiency. *arXiv preprint arXiv:2102.02981*, 2021.

A. W. van der Vaart. *Asymptotic statistics*. Cambridge University Press, Cambridge, UK, 1998.

Karel Vermeulen and Stijn Vansteelandt. Bias-reduced doubly robust estimation. *Journal of the American Statistical Association*, 110(511):1024–1036, 2015.

Martin J Wainwright. *High-Dimensional Statistics : A Non-Asymptotic Viewpoint*. Cambridge University Press, New York, 2019.

Liyuan Xu, Heishiro Kanagawa, and Arthur Gretton. Deep proxy causal learning and its application to confounded bandit policy evaluation. *Advances in Neural Information Processing Systems*, 34, 2021.

Mengjiao Yang, Ofir Nachum, Bo Dai, Lihong Li, and Dale Schuurmans. Off-policy evaluation via the regularized lagrangian. In *Advances in Neural Information Processing Systems*, volume 33, pages 6551–6561, 2020.

Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114, 2017.

Rui Zhang, Masaaki Imaizumi, Bernhard Schölkopf, and Krikamol Muandet. Maximum moment restriction for instrumental variable regression. *arXiv preprint arXiv:2010.07684*, 2020.

Wenjing Zheng and Mark J van der Laan. Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning: Causal Inference for Observational and Experimental Data*, Springer Series in Statistics, pages 459–474. Springer New York, New York, NY, 2011.