DOUBLY ROBUST PROXY CAUSAL LEARNING WITH NEURAL MEAN EMBEDDINGS

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose a novel doubly robust neural mean-embedding framework for proxy causal learning (PCL) for continuous and structured treatments. PCL uses proxy variables and bridge functions to identify causal effects under unobserved confounding: either an outcome bridge or a treatment bridge can identify the target, while doubly robust methods combine both. Our framework learns both bridges with data-adaptive neural features and combines them into doubly robust estimators of population, treatment-conditional, and heterogeneous dose-response functions. Two features of our result are particularly notable: the treatment bridge can be learned without density ratios, and the same treatment bridge yields both population and heterogeneous responses, eliminating the need for a separate bridge fit. Our analysis shows that the doubly robust error is controlled by final-stage error and the smaller of the two projected bridge errors, with no separate density-ratio estimation error. Across challenging continuous- and structured-treatment benchmarks, our method outperforms neural single-bridge estimators and fixed-feature doubly robust estimators.

1 INTRODUCTION AND RELATED WORK

Causal inference from observational data is difficult when treatment assignment is confounded. Standard adjustment assumes that all relevant confounders are observed, yielding $Y^{(a)} \perp A \mid X$ under consistency, no interference, and positivity. This condition underlies covariate adjustment and propensity-score methods (Rubin, 1980; Rosenbaum & Rubin, 1983; Imbens, 2004), outcome-regression methods (Hill, 2011), and representation-learning estimators for treatment effects (Johansson et al., 2016; Yao et al., 2018). Yet important determinants of treatment and outcome may be unrecorded or only indirectly measured. Instrumental-variable methods offer one route beyond observed-confounder adjustment (Reiersøl, 1945; Stock & Trebbi, 2003; Newey & Powell, 2003), but require their own exclusion and relevance conditions that are often restrictive.

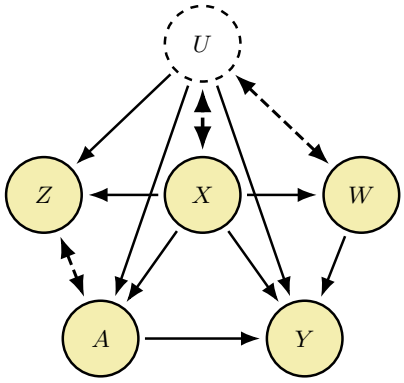

Figure 1: Causal graph for PCL setting.

Proxy causal learning (PCL) uses a different source of information: two observed proxy variables Z and W that are informative about the unobserved confounding (Kuroki & Pearl, 2014; Miao et al., 2018; Tchetgen Tchetgen et al., 2024). Figure 1 illustrates the standard PCL setting. Here, X denotes the pre-treatment covariates, while Z and W provide complementary treatment-side and outcome-side information about the latent confounder U. The treatment A and outcome Y are confounded by observed covariates X and latent variables U. Conditional independence and completeness conditions permit identification without recovering U itself. Identification is expressed through bridge functions whose conditional expectations reproduce regression or weighting equations. Learning these bridges replaces latent-confounder inference by the solution of Fredholm integral equations of the first kind.

Existing PCL estimators mainly follow outcome-bridge or treatment-bridge routes. Outcome-bridge methods learn a function $h(A, X, W)$ satisfying $\mathbb{E}\{Y - h(A, X, W) \mid A, X, Z\} = 0$, and then recover response curves by averaging or conditioning the learned bridge. Linear and semiparametric

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

bridge estimators were developed in proximal causal inference (Tchetgen Tchetgen et al., 2024; Cui et al., 2024); sieve and proxy-control estimators were studied by Deaner (2023); and RKHS-methods provide flexible nonparametric bridge estimators (Mastouri et al., 2021; Singh, 2023). Neural outcome-bridge estimators replace fixed features with learned representations, improving scalability to nonlinear and high-dimensional settings (Xu et al., 2021b; Kompa et al., 2022). These methods learn the regression-type bridge $h(A, X, W)$, but leave open the neural estimation of the complementary weighting-type bridge based on Z , i.e. treatment bridge functions.

Treatment bridges give the complementary weighting route and are analogous to inverse-probability weighting (Rosenbaum & Rubin, 1983; Robins et al., 2000). For a given causal estimand, a treatment bridge $\varphi(A, X, Z)$ is characterized by a moment equation of the form $\mathbb{E}\{r_0(A, X, W) - \varphi(A, X, Z) \mid A, X, W\} = 0$, where r_0 denotes the estimand-specific Radon–Nikodym derivative; e.g., for the population dose-response curve, $r_{0,\mathrm{ATE}} := d(\mathbb{P}_A \otimes \mathbb{P}_{X,W})/d\mathbb{P}_{A,X,W}$. Semiparametric (Cui et al., 2024), minimax (Kallus et al., 2021), and kernel-based treatment-bridge estimators (Bozkurt et al., 2025a) have been developed for proxy causal inference, including extensions to continuous treatments and treatment-conditional dose responses (Bozkurt et al., 2025a). Recent neural treatment-bridge work moves toward adaptive treatment-side estimation (Zhang et al., 2025), but is developed for population binary treatment effects rather than continuous or structured treatments. Restricting treatment to two options cannot capture responses across many modern settings of interest, including in medicine (drug combinations and doses) and advertising (ad designs). Learned neural features let our estimators use these richer treatment descriptions and adapt their representation to the data.

Doubly robust PCL combines the outcome- and treatment-bridge routes, mirroring classical doubly robust causal inference, where validity can be retained when one of two complementary components is correctly specified (Bang & Robins, 2005; Chernozhukov et al., 2018; Kennedy, 2024). This is especially useful in PCL because the two bridges use the proxies differently, and it is rarely known in advance whether the treatment or outcome bridge will be more reliable. Semiparametric and kernel doubly robust proxy estimators have been developed for population responses (Cui et al., 2024; Wu et al., 2024; Bozkurt et al., 2025b), but continuous-treatment methods rely on fixed kernel representations. Beyond population responses, Sverdrup & Cui (2023) develop a proximal P-learner for binary-treatment CATE contrasts. More generally, applications may require full heterogeneous response surfaces $(a, v) \mapsto f_{\mathrm{CATE}}(a, v)$, describing responses in the subpopulation $V = v$, or treatment-conditional surfaces $(a, a') \mapsto f_{\mathrm{ATT}}(a, a')$, describing how units observed at $A = a'$ would respond under intervention a. Existing neural doubly robust PCL methods, however, do not cover these targets for continuous or structured treatments.

We address this gap by proposing a neural doubly robust PCL framework that learns both bridge routes for continuous and structured treatments. Its treatment-side component, TREATMENTNET, is a two-stage neural mean embedding estimator that learns the projected treatment bridge with data-adaptive features by minimizing a density-ratio-free risk. Combining it with a neural outcome bridge and a final residual regression, we propose DR-PCL-NET for population, heterogeneous, and treatment-conditional response surfaces. For heterogeneous responses, we derive a novel conditional self-normalization identity showing that one population treatment bridge can be reused over all supported (a, v). For treatment-conditional responses, a joint objective learns the full (a, a') treatment-bridge surface in one fit. We adapt the proximal closed-form updates of Galashov et al. (2025) to train both bridges. By retaining information from earlier batches, these updates address minibatch instability in the outcome-bridge method of Xu et al. (2021b).

In summary, our contributions are fourfold. (i) We introduce TREATMENTNET that utilizes density-ratio-free treatment-bridge learning with adaptive neural features and full response functions for continuous or structured-treatments. (ii) We combine neural treatment and outcome bridges in DR-PCL-NET, giving doubly robust population, heterogeneous, and treatment-conditional responses. A single population bridge serves all supported subgroups; a joint fit covers the treatment-conditional surface. (iii) We derive nonasymptotic projected error bounds for approximate empirical minimizers that separate approximation, empirical-process, final-stage, and optimization errors. Under bounded complementary errors, the doubly robust remainder is controlled by the smaller projected bridge error, without strong bridge recovery or a separate density-ratio estimation error. (iv) In our experiements, we first construct synthetic examples that isolate challenging regimes in each individual bridge route, which are addressed when both are combined in a doubly robust approach. We then evaluate population, heterogeneous and treatment-conditional responses on challenging benchmarks, including image-valued treatments. Our method outperforms competing approaches in the majority of scenarios.

2 PROBLEM SETUP AND TARGET CAUSAL FUNCTIONS

We observe i.i.d. copies $O = (Y, A, X, Z, W) \sim \mathbb{P}_O$, where $A \in \mathcal{A}$ is a possibly continuous or structured treatment, $Y \in \mathcal{Y} \subseteq \mathbb{R}^{d_Y}$, $X \in \mathcal{X}$ contains measured covariates, $Z \in \mathcal{Z}$ is a treatment-inducing proxy, and $W \in \mathcal{W}$ is an outcome-inducing proxy. The latent confounder is denoted by $U \in \mathcal{U}$. For heterogeneous effects write $X = (S, V)$, where V indexes the subgroup conditioned on and S contains the remaining covariates. We use the notation $\mathbb{P}$ for probability laws; lowercase p denotes a density or mass function when it exists. We study the following causal response functions

$$f_{\mathrm{ATE}}(a) \coloneqq \mathbb{E}\{Y^{(a)}\}, \quad f_{\mathrm{CATE}}(a, v) \coloneqq \mathbb{E}\{Y^{(a)} \mid V = v\}, \quad f_{\mathrm{ATT}}(a, a') \coloneqq \mathbb{E}\{Y^{(a)} \mid A = a'\}$$

These are, respectively, the population, heterogeneous, and treatment-conditional dose-response functions, with domains $\mathcal{A}$, $\mathcal{A} \times \mathcal{V}$, and $\mathcal{A}^2$. The subscripts follow standard discrete-treatment terminology: ATE, CATE, and ATT refer to average, conditional average, and average-on-the-treated treatment effects. Here, they index the corresponding response functions rather than discrete-treatment contrasts. Assume consistency, $Y = Y^{(A)}$, and latent exchangeability, $Y^{(a)} \perp A \mid (U, X)$. For the population, heterogeneous, and treatment-conditional targets, respectively, impose $\mathbb{P}_{X,U} \ll \mathbb{P}_{X,U|A=a}$, $\mathbb{P}_{S,U|V=v} \ll \mathbb{P}_{S,U|A=a,V=v}$, $\mathbb{P}_{X,U|A=a'} \ll \mathbb{P}_{X,U|A=a}$. Under these conditions, the targets admit the latent-variable representations: (i) $f_{\mathrm{ATE}}(a) = \mathbb{E}[\mathbb{E}[Y \mid A = a, X, U]]$, (ii) $f_{\mathrm{CATE}}(a, v) = \mathbb{E}[\mathbb{E}[Y \mid A = a, S, V = v, U] \mid V = v]$, and (iii) $f_{\mathrm{ATT}}(a, a') = \mathbb{E}[\mathbb{E}[Y \mid A = a, X, U] \mid A = a']$. These representations show that the causal functions would be identified by adjustment if the latent confounder U were observed. Proxy causal learning instead identifies the same functions without observing or estimating U, by using the proxies Z and W to learn bridge functions.

Intuitively, Z is informative about the latent confounding through treatment assignment, while W is informative about the same latent confounding through the outcome. We formalize these proxy roles through the following conditional independence assumptions, and then impose completeness to ensure that the proxies carry enough information for bridge identification.

Assumption 2.1 (Proxy conditional independencies)**.** *The structural causal model is compatible with Figure 1 and obeys* $Y \perp Z \mid (U, A, X)$, $W \perp (A, Z) \mid (U, X)$.

Assumption 2.1 formalizes the roles of Z and W as treatment and outcome proxies, respectively. In addition, identification requires that these proxies be sufficiently informative about the latent confounder U conditional on (A, X) (Miao et al., 2018; Tchetgen Tchetgen et al., 2024; Cui et al., 2024). This is captured by the following completeness assumption.

Assumption 2.2 (Completeness)**.** *For* $\mathbb{P}_{A,X}$*-almost all* $(a, x) \in \mathcal{A} \times \mathcal{X}$*, and for every square-integrable function* $\ell \in L^2(\mathbb{P}_{U|A=a,X=x})$*: (i)* $\mathbb{E}[\ell(U) \mid A = a, X = x, Z] = 0$ $\mathbb{P}_{Z|A=a,X=x}$*-a.e.* $\implies$ $\ell(U) = 0$ $\mathbb{P}_{U|A=a,X=x}$*-a.e.*, *(ii)* $\mathbb{E}[\ell(U) \mid A = a, X = x, W] = 0$ $\mathbb{P}_{W|A=a,X=x}$*-a.e.* $\implies$ $\ell(U) = 0$ $\mathbb{P}_{U|A=a,X=x}$*-a.e.*.

Assumption 2.2 is the standard sufficient condition used in PCL and nonparametric instrumental-variable regression (Miao et al., 2018; Newey & Powell, 2003; Cui et al., 2024; Tchetgen Tchetgen et al., 2024). For fixed a, x, it requires that no nonzero square-integrable function of U have zero conditional mean given the relevant proxy. In finite discrete models it reduces to a full-rank condition, while sufficiently rich exponential-family conditional laws provide familiar continuous sufficient conditions. Appendix A reviews these examples from (Newey & Powell, 2003; Tchetgen Tchetgen et al., 2024) and separately discusses the additional range conditions required for bridge existence.

3 IDENTIFICATION OF CAUSAL FUNCTIONS

This section first reviews the established outcome-bridge and treatment-bridge representations, then develops our general-treatment extensions. We derive doubly robust identification of the full heterogeneous response surface $f_{\mathrm{CATE}}(a, v)$ and the treatment-conditional response surface $f_{\mathrm{ATT}}(a, a')$ for continuous or structured treatments. We also give an alternative CATE identification result: after conditional normalization, any valid population treatment bridge induces a valid CATE bridge. Our proposed method uses this identity to construct a density-ratio-free heterogeneous-response estimator that reuses the population bridge rather than fitting a separate target-specific bridge.

Outcome bridge identification. The standard route in PCL uses outcome bridge function (Miao et al., 2018; Mastouri et al., 2021; Xu et al., 2021b; Kompa et al., 2022): it seeks a function $h(A, X, W)$

such that $\mathbb{E}[h(A, X, W) \mid A, X, Z] = \mathbb{E}[Y \mid A, X, Z]$. Once such a bridge exists, response functions are obtained by averaging or conditioning h over the relevant observed distribution.

Theorem 3.1 (Identification via the outcome bridge)**.** *Suppose Assumption 2.1 and Assumption 2.2 (i) hold. Furthermore, suppose that for $\mathbb{P}_{A,X}$-almost all $(a, x) \in \mathcal{A} \times \mathcal{X}$, there exists an outcome bridge function $h_0(a, x, \cdot) \in L^2\big(\mathbb{P}_{W|A=a,X=x}\big)$ such that*

$$\mathbb{E}[Y \mid A = a, X = x, Z] = \mathbb{E}[h_0(a, x, W) \mid A = a, X = x, Z], \qquad \mathbb{P}_{Z|A=a,X=x}\text{-a.e.} \tag{1}$$

Then, the causal functions are identified by (i) $f_{\mathrm{ATE}}(a) = \mathbb{E}[h_0(a, X, W)]$, (ii) $f_{\mathrm{CATE}}(a, v) = \mathbb{E}[h_0(a, X, W) \mid V = v]$, (iii) $f_{\mathrm{ATT}}(a, a') = \mathbb{E}[h_0(a, X, W) \mid A = a']$.

The population dose-response formula is the standard outcome-bridge result of Miao et al. (2018); the conditional formulas follow by conditioning the same representation (Singh, 2023, Theorem 1).

Treatment bridge identification. We next turn to treatment bridge identification, which complements the outcome-bridge view and is closer in spirit to inverse propensity weighting (Rosenbaum & Rubin, 1983; Robins et al., 2000). This perspective has been developed in semiparametric, minimax, and kernel forms, including recent extensions to continuous treatments and density-ratio-free formulations (Cui et al., 2024; Kallus et al., 2021; Wu et al., 2024; Bozkurt et al., 2025a;b). It is also the identification route that will motivate our treatment-side estimator.

Define the target specific Radon–Nikodym derivatives for population, heterogeneous, and treatment-conditional response curves: (i) $r_{0,\mathrm{ATE}} := d(\mathbb{P}_A \otimes \mathbb{P}_{X,W})/d\mathbb{P}_{A,X,W}$, (ii) $r_{0,\mathrm{CATE}}(\cdot, v, \cdot, \cdot) := d(\mathbb{P}_{A|V=v} \otimes \mathbb{P}_{S,W|V=v})/d\mathbb{P}_{A,S,W|V=v}$, (iii) $r_{0,\mathrm{ATT}}(a, a', \cdot, \cdot) := d\mathbb{P}_{X,W|A=a'}/d\mathbb{P}_{X,W|A=a}$. When common dominating measures exist, these are, respectively, $p_A(a)/p_{A|X,W}(a \mid x, w)$, $p_{A|V}(a \mid v)/p_{A|S,V,W}(a \mid s, v, w)$, and $p_{X,W|A}(x, w \mid a')/p_{X,W|A}(x, w \mid a)$, on their denominator supports.

Theorem 3.2 (Identification via treatment bridge)**.** *Suppose Assumptions 2.1 and 2.2 (ii) hold. Assume that for $\mathbb{P}_{A,X}$-almost all $(a, x) \in \mathcal{A} \times \mathcal{X}$, there exist treatment bridge functions $\varphi_0^{\mathrm{ATE}}(a, x, \cdot) \in L^2\big(\mathbb{P}_{Z|A=a,X=x}\big)$ and, for any $a' \in \mathcal{A}$, $\varphi_0^{\mathrm{ATT}}(a, a', x, \cdot) \in L^2\big(\mathbb{P}_{Z|A=a,X=x}\big)$. In addition, assume that for $\mathbb{P}_{A,S,V}$-almost all $(a, s, v) \in \mathcal{A} \times \mathcal{S} \times \mathcal{V}$, there exists $\varphi_0^{\mathrm{CATE}}(a, v, s, \cdot) \in L^2\big(\mathbb{P}_{Z|A=a,S=s,V=v}\big)$. Suppose these bridge functions satisfy, almost everywhere with respect to the corresponding conditional law of W,*

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a, X, W] = r_{0,\mathrm{ATE}}(a, X, W), \tag{2}$$

$$\mathbb{E}[\varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, S, V = v, W] = r_{0,\mathrm{CATE}}(a, v, S, W), \tag{3}$$

$$\mathbb{E}[\varphi_0^{\mathrm{ATT}}(a, a', X, Z) \mid A = a, X, W] = r_{0,\mathrm{ATT}}(a, a', X, W). \tag{4}$$

Then, the causal functions are identified by (i) $f_{\mathrm{ATE}}(a) = \mathbb{E}[Y\varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a]$, (ii) $f_{\mathrm{CATE}}(a, v) = \mathbb{E}[Y\varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, V = v]$, (iii) $f_{\mathrm{ATT}}(a, a') = \mathbb{E}[Y\varphi_0^{\mathrm{ATT}}(a, a', X, Z) \mid A = a]$.

The ATE and ATT statements follow from the treatment-bridge identification results of Bozkurt et al. (2025a). We prove the new CATE extension in Appendix B.

Doubly robust identification. Combining outcome and treatment bridges yields doubly robust formulas that remain valid when either bridge is correctly specified and underlie semiparametric and kernel DR proxy methods (Cui et al., 2024; Wu et al., 2024; Bozkurt et al., 2025b). The population dose-response formula recovers existing results; the heterogeneous and treatment-conditional formulas are new extensions developed here.

The doubly robust identification theorem below extends the bridge structure that appears in the influence functions derived in Appendix D for discrete treatments. That appendix treats A as discrete, and V as discrete for the heterogeneous target. Here, we state the corresponding bridge identities directly as identification formulas for general treatments and covariates, including continuous A and V.

Theorem 3.3 (Doubly robust identification)**.** *Suppose the causal and target-specific latent overlap conditions stated in Section 2, together with Assumption 2.1, hold. Let the candidate bridges h, φ^{ATE}, φ^{CATE}, and φ^{ATT} be measurable and square-integrable, and define*

$$\Psi_{\mathrm{ATE}}(a; h, \varphi^{\mathrm{ATE}}) = \mathbb{E}[h(a, X, W)] + \mathbb{E}[\varphi^{\mathrm{ATE}}(a, X, Z)\{Y - h(a, X, W)\} \mid A = a], \tag{5}$$

$$\Psi_{\mathrm{CATE}}(a, v; h, \varphi^{\mathrm{CATE}}) = \mathbb{E}[h(a, X, W) \mid V = v]$$

$$+ \mathbb{E}[\varphi^{\mathrm{CATE}}(a, v, S, Z)\{Y - h(a, X, W)\} \mid A = a, V = v], \tag{6}$$

$$\Psi_{\mathrm{ATT}}(a, a'; h, \varphi^{\mathrm{ATT}}) = \mathbb{E}[h(a, X, W) \mid A = a'] + \mathbb{E}[\varphi^{\mathrm{ATT}}(a, a', X, Z)\{Y - h(a, X, W)\} \mid A = a]. \tag{7}$$

For each target, the corresponding functional equals the causal response if either (i) h satisfies Eq. 1 and Assumption 2.2 (i) holds, or (ii) its treatment-bridge candidate satisfies the corresponding equation in Eqs. 2–4 and Assumption 2.2 (ii) holds.

The ATE representation in Eq. 5 is the covariate-adjusted analogue of the doubly robust dose-response formula in Bozkurt et al. (2025b). Appendix C proves our novel heterogeneous and treatment-conditional extensions.

One population treatment bridge for heterogeneous responses. We show that conditional normalization of any valid population treatment bridge produces a valid CATE bridge. This gives a new identifying formula and, subsequently, a density-ratio-free estimator that reuses a single population bridge over all supported (a, v). For any valid ATE bridge, define $D_0(a, v) := \mathbb{E}\{\varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a, V = v\} = d(\mathbb{P}_A \otimes \mathbb{P}_V)/d\mathbb{P}_{A,V}(a, v)$, where the second equality follows from the conditional-expectation property of Radon–Nikodym derivatives.

Proposition 3.4 (ATE bridge induces CATE bridge)**.** *At points where $0 < D_0(a, v) < \infty$, the Radon–Nikodym chain rule gives $r_{0,\mathrm{ATE}}(a, (s, v), w) = D_0(a, v) r_{0,\mathrm{CATE}}(a, v, s, w)$ a.e. Hence $\varphi_0^{\mathrm{CATE}}(a, v, s, z) := \varphi_0^{\mathrm{ATE}}(a, (s, v), z)/D_0(a, v)$ solves Eq. 3, and*

$$f_{\mathrm{CATE}}(a, v) = \frac{\mathbb{E}\{Y \varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a, V = v\}}{D_0(a, v)}.$$

For an arbitrary ATE-scale candidate $\varphi(a, X, Z)$, let $D_\varphi(a, v) = \mathbb{E}\{\varphi(a, X, Z) \mid A = a, V = v\}$. Whenever all terms are finite and $D_\varphi(a, v) \neq 0$, define

$$\Psi^{\mathrm{SN}}_{\mathrm{CATE}}(a, v; h, \varphi) = \mathbb{E}\{h(a, X, W) \mid V = v\} + \frac{\mathbb{E}[\varphi(a, X, Z)\{Y - h(a, X, W)\} \mid A = a, V = v]}{D_\varphi(a, v)}.$$

This functional identifies $f_{\mathrm{CATE}}(a, v)$ if either h is a valid outcome bridge or φ is a valid ATE treatment bridge.

Thus, heterogeneous responses reuse one population bridge and retain double robustness, without fitting another treatment bridge or sampling treatments conditional on V (Appendix F.2).

4 Neural mean embedding approach for dose-response curve

We now construct our novel TreatmentNet and the doubly robust estimator DR-PCL-Net in combination with OutcomeNet for the population dose-response.

Notation and sample splitting. Let $b \in \{h, \varphi\}$ index the outcome- and treatment-bridge routes, respectively. We use uppercase letters for random variables and lowercase letters for their realizations. Given $\mathcal{D} = \{o_i\}_{i=1}^n$ with $o_i = (a_i, x_i, w_i, z_i, y_i)$, partition the observation indices into three disjoint base folds: $\mathcal{I}_s^{(b)}$ denotes the base fold assigned to Stage s of route b. Stage 1 uses $\mathcal{I}_1^{(b)}$ to estimate the conditional mean embedding of the relevant proxy features. Stage 2 uses $\mathcal{I}_2^{(b)}$ to learn the bridge after substituting the Stage 1 estimate. Stage 3 uses the held-out observations to estimate the causal response, using $\mathcal{I}_3^{(h)}$ for OutcomeNet, $\mathcal{I}_3^{(\varphi)}$ for TreatmentNet, and $\mathcal{I}_3^{(\mathrm{DR})}$ for DR-PCL-Net.

TreatmentNet and the ratio-free risk. Define the conditional operator $(T_\varphi \varphi)(a, x, w) := \mathbb{E}\{\varphi(a, x, Z) \mid A = a, X = x, W = w\}$ and denote $g_\varphi := T_\varphi \varphi$. For $Q_\varphi := (A, X, W) \sim \mathbb{P}_{Q_\varphi}$ and independent copies $Q_{\varphi,j} = (A_j, X_j, W_j)$, $j \in \{1, 2\}$, set

$$\mathcal{R}^{\mathrm{RF}}_{\mathrm{ATE}}(u) := \mathbb{E}\{u(A_1, X_1, W_1)^2\} - 2\mathbb{E}\{u(A_1, X_2, W_2)\}. \tag{8}$$

The linear term integrates under $\mathbb{P}_A \otimes \mathbb{P}_{X,W}$, rather than the observed joint law.

Lemma 4.1 (Ratio-free risk characterization)**.** *If $\mathbb{P}_A \otimes \mathbb{P}_{X,W} \ll \mathbb{P}_{A,X,W}$, with derivative $r_{0,\mathrm{ATE}} \in L^2(\mathbb{P}_{Q_\varphi})$, then $\mathcal{R}^{\mathrm{RF}}_{\mathrm{ATE}}(u) - \mathcal{R}^{\mathrm{RF}}_{\mathrm{ATE}}(r_{0,\mathrm{ATE}}) = \|u - r_{0,\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_{Q_\varphi})}$.*

Under treatment-bridge existence, the lemma shows that minimizing $\mathcal{R}^{\mathrm{RF}}_{\mathrm{ATE}}(T_\varphi\varphi)$ targets $T_\varphi\varphi = r_{0,\mathrm{ATE}}$ without estimating or evaluating the Radon–Nikodym derivative. However, evaluating this loss requires $T_\varphi\varphi$ which is not readily accessible. For a generic nonlinear bridge, evaluating the projected loss would itself require solving an inner conditional-expectation problem that changes with the bridge parameters. A DeepIV-style alternative would model $\mathbb{P}_{Z|A,X,W}$ and approximate this projection by Monte Carlo integration during training (Hartford et al., 2016). Our construction follows a neural mean-embedding approach (Xu & Gretton, 2023; Xu et al., 2021b). In particular, to evaluate the ratio-free criterion at $u = T_\varphi\varphi$, we parameterize the treatment bridge by a linear head over separable learnable features, $\varphi(a,x,z) = \boldsymbol{\varphi}^\top\left[\phi^{(\varphi)}_{AX,2}(a,x)\otimes\phi^{(\varphi)}_{Z,2}(z)\right]$. For this class of treatment bridge functions, define the conditional feature mean $\mu^{(\varphi)}_Z(a,x,w) := \mathbb{E}\Big\{\phi^{(\varphi)}_{Z,2}(Z)\mid A=a, X=x, W=w\Big\}$. Linearity of conditional expectation then gives

$$(T_\varphi\varphi)(a,x,w) = \boldsymbol{\varphi}^\top\left[\phi^{(\varphi)}_{AX,2}(a,x)\otimes\mu^{(\varphi)}_Z(a,x,w)\right]. \tag{9}$$

Stage 1 estimates this vector-valued conditional mean through $\widehat{\mu}^{(\varphi)}_Z(a,x,w) := \{\widehat{\boldsymbol{V}}^{(\varphi)}\}^\top\phi^{(\varphi)}_{AXW,1}(a,x,w;\widehat{\theta}^{(\varphi)}_1)$, where $(\widehat{\theta}^{(\varphi)}_1, \widehat{\boldsymbol{V}}^{(\varphi)})$ are obtained by regressing $\phi^{(\varphi)}_{Z,2}(Z)$ on $\phi^{(\varphi)}_{AXW,1}(A,X,W)$. For fixed neural features, this single Stage 1 regression evaluates the projection for every linear bridge head: Stage 2 simply replaces $\mu^{(\varphi)}_Z$ by $\widehat{\mu}^{(\varphi)}_Z$ in Equation 9. It therefore avoids full conditional-law modeling and repeated proxy sampling. The same factorization applies to OUTCOMENET, with $W \mid A, X, Z$ (Xu et al., 2021a).

Proximal bridge objectives. At iteration t, let $\mathcal{B}^{(\varphi)}_{s,t} \subset \mathcal{I}^{(\varphi)}_s$ be the Stage-s mini-batch and $m^{(\varphi)}_{s,t} := |\mathcal{B}^{(\varphi)}_{s,t}|$. The first-stage feature-regression objective is

$$\widehat{\mathcal{L}}^{\mathrm{prox}}_{\varphi,1,t}(\theta^{(\varphi)}_1, \boldsymbol{V}^{(\varphi)}) = \frac{1}{m^{(\varphi)}_{1,t}}\sum_{i\in\mathcal{B}^{(\varphi)}_{1,t}}\left\|\phi^{(\varphi)}_{Z,2}(z_i) - \{\boldsymbol{V}^{(\varphi)}\}^\top\phi^{(\varphi)}_{AXW,1}(a_i,x_i,w_i)\right\|_2^2 + \lambda^{(\varphi)}_1\|\boldsymbol{V}^{(\varphi)} - \widehat{\boldsymbol{V}}^{(\varphi)}_t\|_F^2,$$

Here, $\theta^{(\varphi)}_1$ denotes the parameters of the first-stage neural network $\phi^{(\varphi)}_{AXW,1}(a_i,x_i,w_i)$. Let $\hat{\mu}^{(\varphi)}_{Z,t}$ denote the fitted conditional feature mean. For $i,j \in \mathcal{B}^{(\varphi)}_{2,t}$, define $\boldsymbol{u}_{ij,t} := \phi^{(\varphi)}_{AX,2}(a_i,x_j)\otimes\hat{\mu}^{(\varphi)}_{Z,t}(a_i,x_j,w_j)$. The second-stage proximal training objective is

$$\widehat{\mathcal{L}}^{\mathrm{RF,U}}_{\varphi,2,t}(\theta^{(\varphi)}_2, \boldsymbol{\varphi}) = \frac{1}{m^{(\varphi)}_{2,t}}\sum_{i\in\mathcal{B}^{(\varphi)}_{2,t}}(\boldsymbol{\varphi}^\top\boldsymbol{u}_{ii,t})^2 - \frac{2}{m^{(\varphi)}_{2,t}\left(m^{(\varphi)}_{2,t}-1\right)}\sum_{\substack{i,j\in\mathcal{B}^{(\varphi)}_{2,t}\\ i\neq j}}\boldsymbol{\varphi}^\top\boldsymbol{u}_{ij,t} + \lambda^{(\varphi)}_2\|\boldsymbol{\varphi} - \widehat{\boldsymbol{\varphi}}_t\|_2^2,$$

where $\theta^{(\varphi)}_2$ collects the parameters of the second stage networks $\phi^{(\varphi)}_{AX,2}$ and $\phi^{(\varphi)}_{Z,2}$. Minimizing this objective yields the treatment bridge estimate $\hat{\varphi}$.

OUTCOMENET. OUTCOMENET follows the two-stage neural conditional-embedding construction of Xu et al. (2021b;a). Stage 1 estimates $\mu^{(h)}_W(a,x,z) = \mathbb{E}\{\phi^{(h)}_{W,2}(W) \mid A=a, X=x, Z=z\}$. Stage 2 regresses the outcome Y on $\phi^{(h)}_{AX,2}(A,X)\otimes\hat{\mu}^{(h)}_W(A,X,Z)$, using $h(a,x,w) = \{\boldsymbol{H}^{(h)}\}^\top[\phi^{(h)}_{AX,2}(a,x)\otimes\phi^{(h)}_{W,2}(w)]$. We employ a new training procedure, centering the penalties on both linear layers at their previous values, following Galashov et al. (2025). The full objectives and closed-form updates are given in Appendix E.1. On the evaluation fold, $\hat{f}^{(h)}_{\mathrm{ATE}}(a) = \frac{1}{|\mathcal{I}^{(h)}_3|}\sum_{i\in\mathcal{I}^{(h)}_3}\hat{h}(a,x_i,w_i)$.

Training schedule. For each paired $\mathcal{B}^{(b)}_{1,t}$ and $\mathcal{B}^{(b)}_{2,t}$, $b\in\{h,\varphi\}$, we alternate $T_{b,1}$ first-stage updates and $T_{b,2}$ second-stage updates. Stage 1 holds the proxy-side target feature map fixed, updates the Stage-1 input featurizer by gradient-based steps, and then updates $\boldsymbol{V}^{(b)}$ by its history-centered closed-form solve. To account for the dependence of the Stage-1 optimum on the proxy feature representation, Stage 2 differentiates through an auxiliary closed-form Stage-1 solve; the linear bridge heads are then updated analytically Galashov et al. (2025); Xu et al. (2021b;a). Appendix E gives the complete objectives and algorithm.

Final response regressions. On held-out folds $\mathcal{I}_3^{(\varphi)}$ and $\mathcal{I}_3^{(\mathrm{DR})}$, define the TREATMENTNET and DR-PCL-NET pseudo-outcomes $\gamma_i^{(\varphi)} := y_i \widehat{\varphi}(a_i, x_i, z_i)$, $\gamma_i^{(\mathrm{DR})} := \widehat{\varphi}(a_i, x_i, z_i)\{y_i - \widehat{h}(a_i, x_i, w_i)\}$. Regressing these pseudo-outcomes on a_i gives functions

$$\widehat{k}_\varphi(a) \approx \mathbb{E}[Y\widehat{\varphi}(a, X, Z) \mid A = a], \;\; \widehat{k}_{\mathrm{DR}}(a) \approx \mathbb{E}\Big[\widehat{\varphi}(a, X, Z)\{Y - \widehat{h}(a, X, W)\} \mid A = a\Big]. \quad (10)$$

The resulting estimators for TREATMENTNET and DR-PCL-NET are given by $\widehat{f}_{\mathrm{ATE}}^{(\varphi)}(a) = \widehat{k}_\varphi(a)$, $\widehat{f}_{\mathrm{ATE}}^{\mathrm{DR}}(a) = \widehat{f}_{\mathrm{ATE}}^{(h)}(a) + \widehat{k}_{\mathrm{DR}}(a)$, respectively, where the OUTCOMENET average in the second expression can be evaluated on the same held-out fold. These are the sample analogues of the treatment-bridge and doubly robust identification formulas in Section 3.

Extensions beyond population response. For heterogeneous responses, Proposition 3.4 leads to a self-normalized estimator: a joint regression on (A, V) learns the bridge-weighted numerator and bridge normalizer, whose stabilized ratio estimates $f_{\mathrm{CATE}}(a, v)$; see Appendix F.2. The doubly robust version replaces the numerator by the corresponding bridge-weighted residual correction as outlined in Appendix F.3. For treatment-conditional responses, our pairwise objective learns $\varphi^{\mathrm{ATT}}(a, a', x, z)$ jointly over both treatment arguments. A single final regression then estimates $f_{\mathrm{ATT}}(a, a')$, avoiding a separate fit for each reference level a' (Appendices G.2–G.3).

5 PROJECTED-ERROR ANALYSIS

Appendix H derives nonasymptotic bounds for the sample-split estimators under stagewise approximate empirical minimization (van der Vaart & Wellner, 1996, Section 3.2). The DR-PCL-NET error decomposes into final averaging and regression errors and a doubly robust remainder. Under bounded complementary bridge errors, this remainder is controlled by the smaller of the outcome- and treatment-side projected bridge errors. Consequently, once the final-stage errors vanish, projected convergence of either bridge is sufficient for consistency. Here we summarize the resulting population dose-response rate consequences; the corresponding CATE and ATT results are given in Appendices H.7 and H.8–H.10. The bridge error bounds control projected conditional-moment errors (Bennett et al., 2025) rather than strong L^2-recovery of either solution to the underlying ill-posed inverse problems, and no separate density-ratio estimation error appears.

Let $Q_h := (A, X, Z), Q_\varphi := (A, X, W)$, and $m_0(q_h) := \mathbb{E}[Y \mid Q_h = q_h] \in \mathbb{R}^{d_Y}$. For vector-valued u and scalar v, respectively, define the conditional expectation operators $(T_h u)(a, x, z) := \mathbb{E}[u(a, x, W) \mid A = a, X = x, Z = z]$, and $(T_\varphi v)(a, x, w) := \mathbb{E}[v(a, x, Z) \mid A = a, X = x, W = w]$. We measure the bridge learning through projected conditional-moment residuals: $\mathcal{R}_h^{\mathrm{proj}} := \|T_h \hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h}; \mathbb{R}^{d_Y})}$, and $\mathcal{R}_\varphi^{\mathrm{proj}} := \|T_\varphi \hat{\varphi} - r_{0,\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_{Q_\varphi})}$. When valid bridges exist, these equal $\|T_h(\widehat{h} - h_0)\|_2^2$ and $\|T_\varphi(\widehat{\varphi} - \varphi_0^{\mathrm{ATE}})\|_2^2$, respectively. At empirical-risk stage j, the optimization gap $\xi_j \geq 0$ is defined by $\widehat{R}_j(\widehat{\vartheta}_j) \leq \inf_{\vartheta \in \mathcal{V}_j} \widehat{R}_j(\vartheta) + \xi_j$.

For $b \in \{h, \varphi\}$, let $\nu_b = (n_b, m_b)$ denote the two bridge-training sample sizes. Suppressing confidence arguments and fixed bounded multipliers given in the Appendix H, define

$$\rho_{h,\nu_h} := \kappa_{h,1,\nu_h} + \kappa_{h,2,\nu_h} + \Delta_{h,1,\nu_h} + \Delta_{h,2,\nu_h}, \;\; \rho_{\varphi,\nu_\varphi} := \kappa_{\varphi,1,\nu_\varphi} + \kappa_{\varphi,2,\nu_\varphi} + \Delta_{\varphi,1,\nu_\varphi} + \Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}.$$

Here the κ-terms are approximation errors and the Δ-terms are uniform empirical-process deviations for the corresponding loss classes; see also Bartlett et al. (2005) and Foster & Syrgkanis (2023). In particular, $\rho_{\varphi,\nu_\varphi}$ contains no separate density-ratio estimation error. Let t_h, t_φ, and t_{DR} be the held-out sample sizes for the OUTCOMENET average, TREATMENTNET response regression, and DR-PCL-NET final stage, respectively.

Suppressing confidence arguments and fixed bounded multipliers, also define $\rho_{\varphi,3,t_\varphi} := \kappa_{\varphi,3,t_\varphi} + \Delta_{\varphi,3,t_\varphi}$, $\rho_{\mathrm{DR},3,t_{\mathrm{DR}}} := \kappa_{\mathrm{DR},t_{\mathrm{DR}}} + \Delta_{\mathrm{DR},t_{\mathrm{DR}}}$. These are the final-stage regression errors for TREATMENTNET and the DR-PCL-NET residual correction in Eq. 10. Furthermore, writing $\xi_{h,1,\nu_h}, \xi_{h,2,\nu_h}$ and $\xi_{\varphi,1,\nu_\varphi}, \xi_{\varphi,2,\nu_\varphi}$ for the bridge-stage optimization gaps and $\xi_{\varphi,3,t_\varphi}, \xi_{\mathrm{DR},3,t_{\mathrm{DR}}}$ for those of the final regressions, define, again up to fixed bounded constants, $\tilde{\rho}_{h,\nu_h} := \rho_{h,\nu_h} + \xi_{h,1,\nu_h} + \xi_{h,2,\nu_h}$, $\tilde{\rho}_{\varphi,\nu_\varphi} := \rho_{\varphi,\nu_\varphi} + \xi_{\varphi,1,\nu_\varphi} + \xi_{\varphi,2,\nu_\varphi}$, $\tilde{\rho}_{\varphi,3,t_\varphi} := \rho_{\varphi,3,t_\varphi} + \xi_{\varphi,3,t_\varphi}$, $\tilde{\rho}_{\mathrm{DR},3,t_{\mathrm{DR}}} := \rho_{\mathrm{DR},3,t_{\mathrm{DR}}} + \xi_{\mathrm{DR},3,t_{\mathrm{DR}}}$. Corollaries H.9 and H.18 give, for the two bridge routes, $\mathcal{R}_b^{\mathrm{proj}} = O_p(\tilde{\rho}_{b,\nu_b}), \;\; b \in \{h, \varphi\}$.

Theorem 5.1 (Population dose-response error bounds)**.** *Suppose the conditions of Corollaries H.9 and H.18 and of Theorems H.13, H.14, H.24, and H.30 hold. For the* DR-PCL-NET *bound, also suppose Assumption H.31 holds. Then the corresponding squared-error rates are:*

(i) OUTCOMENET*:* $\|\hat{f}_{\mathrm{ATE}}^{(h)} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p\big(t_h^{-1} + \tilde{\rho}_{h,\nu_h}\big)$.

(ii) TREATMENTNET*:* $\|\hat{f}_{\mathrm{ATE}}^{(\varphi)} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p\big(\tilde{\rho}_{\varphi,3,t_\varphi} + \tilde{\rho}_{\varphi,\nu_\varphi}\big)$.

(iii) DR-PCL-NET*:* $\|\hat{f}_{\mathrm{ATE}}^{\mathrm{DR}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p\big(t_{\mathrm{DR}}^{-1} + \tilde{\rho}_{\mathrm{DR},3,t_{\mathrm{DR}}} + \min\{\tilde{\rho}_{h,\nu_h}, \tilde{\rho}_{\varphi,\nu_\varphi}\}\big)$.

Theorem 5.1 allows consistent causal-response estimation without L^2 convergence of either bridge. Under Assumption H.31, consistency follows if $t_{\mathrm{DR}} \to \infty$, $\widetilde{\rho}_{\mathrm{DR},3,t_{\mathrm{DR}}} = o_p(1)$, $\min\{\widetilde{\rho}_{h,\nu_h}, \widetilde{\rho}_{\varphi,\nu_\varphi}\} = o_p(1)$. Importantly, our result does not require controlling strong L^2-norm convergence of either $\widehat{h}$ to h_0 or $\widehat{\varphi}$ to φ_0^{ATE}. Since the bridge functions h_0 and φ_0^{ATE} are solutions to ill-posed inverse problems, controlling L^2-norms would require restrictive assumptions on measure of ill-posedness (Blundell et al., 2007; Chen & Christensen, 2018). The projected bridge-rate inputs are provided by Corollaries H.9 and H.18, while Theorem H.30 propagates them through the doubly robust remainder and the final-stage errors.

6 NUMERICAL EXPERIMENTS

We evaluate the proposed neural doubly robust proxy estimators against neural baselines TREATMENTNET, OUTCOMENET (Xu et al., 2021b, with our updated optimization), NMMR (Kompa et al., 2022), and fixed-kernel feature baselines DRKPV (Bozkurt et al., 2025b), PKDR (Wu et al., 2024), and KPV (Mastouri et al., 2021; Singh, 2023). Performance is measured by causal MSE between estimated and ground-truth response curves over the evaluation grid. The results use ten data-generating seeds and three split/initialization repetitions per seed; repetitions are averaged within each seed, and standard errors are computed across the ten seed-level means. Baselines are included only when compatible with the benchmark. Implementation details and neural architectures used in the experiments are given in Appendix I.

Bridge-route stress test. We use two analytic DGPs with closed-form bridges to isolate route-specific failure modes; Appendix I.4.1 gives the complete DGPs and bridge derivations. In Scenario A, increasing κ_2 strengthens treatment selection through the outcome-relevant latent factor U_2. The exact outcome bridge remains invariant, whereas worsening overlap increases the second moment of the treatment bridge. In Scenario B, decreasing λ_{W_2} weakens the outcome-proxy measurement of U_2 and amplifies the corresponding components of the outcome bridge. The treatment bridge remains invariant because treatment depends only on U_1. Oracle and learned errors exhibit the corresponding reversal, while DR-PCL-NET remains below the stressed route.

Table 1: Bridge-route stress tests at $n = 4000$. The oracle diagnostics replace the corresponding learned bridge by its closed-form value while retaining held-out final-stage for estimation.

		Learned estimators			Oracle diagnostics	
Scenario	Stress level	OUTCOMENET	TREATMENTNET	DR-PCL-NET	Outcome bridge	Treatment bridge
A	$\kappa_2 = 0.0$	3.799 ± 0.517	4.775 ± 0.414	$\mathbf{2.403 \pm 0.319}$	2.339 ± 0.351	2.551 ± 0.275
	$\kappa_2 = 0.2$	4.034 ± 0.610	4.998 ± 0.846	$\mathbf{2.714 \pm 0.311}$	2.339 ± 0.351	2.850 ± 0.429
	$\kappa_2 = 0.4$	4.877 ± 0.702	7.974 ± 1.230	$\mathbf{4.175 \pm 0.546}$	2.339 ± 0.351	4.375 ± 0.557
	$\kappa_2 = 0.6$	9.689 ± 1.077	27.388 ± 3.307	$\mathbf{8.907 \pm 1.274}$	2.339 ± 0.351	11.323 ± 1.166
B	$\lambda_{W_2} = 0.9$	1.259 ± 0.074	1.591 ± 0.157	$\mathbf{0.762 \pm 0.058}$	0.677 ± 0.125	1.130 ± 0.100
	$\lambda_{W_2} = 0.7$	1.979 ± 0.157	1.592 ± 0.186	$\mathbf{1.494 \pm 0.134}$	0.902 ± 0.124	1.130 ± 0.100
	$\lambda_{W_2} = 0.6$	2.605 ± 0.253	$\mathbf{1.944 \pm 0.216}$	2.408 ± 0.199	1.236 ± 0.163	1.130 ± 0.100
	$\lambda_{W_2} = 0.5$	4.759 ± 0.793	$\mathbf{2.400 \pm 0.255}$	3.205 ± 0.288	2.305 ± 0.316	1.130 ± 0.100
	$\lambda_{W_2} = 0.4$	7.604 ± 1.081	$\mathbf{2.936 \pm 0.359}$	5.057 ± 0.421	3.483 ± 0.636	1.130 ± 0.100

Table 1 demonstrates the results for both scenarios: the treatment route deteriorates while the oracle outcome bridge remains invariant in Scenario A, whereas the outcome route deteriorates while the oracle treatment bridge remains invariant in Scenario B. DR-PCL-NET stays below the deliberately stressed route throughout both sweeps, illustrating protection when the more stable bridge is unknown a priori rather than oracle route selection.

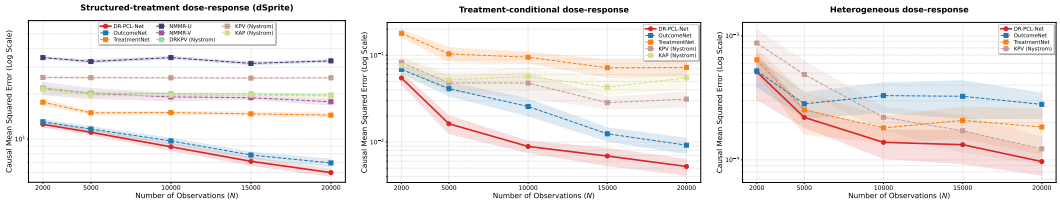

Figure 2: Causal-MSE size sweeps for dSprite (left), treatment-conditional dose-response (center), and heterogeneous dose-response (right). Each point averages repetitions within a data-generating seed; shaded regions are 95% confidence intervals across the seed-level means.

We next evaluate the performance across population, structured-treatment, heterogeneous, and treatment-conditional benchmarks.

Continuous treatment population response benchmarks. For population dose-response estimation, we use the Legalized Abortion and Crime benchmark (Donohue & Levitt, 2001; Mastouri et al., 2021), the calibrated semi-synthetic Deaner IM/IR grade-retention designs (Fruehwirth et al., 2016; Deaner, 2023), and the nonlinear Demand design of Xu et al. (2021b).

Structured-treatment population response benchmark. To assess scalability beyond low-dimensional treatments, we adopt the structured-treatment dSprite design (Matthey et al., 2017; Xu et al., 2021b), in which the treatment is a flattened 64×64 image.

Heterogeneous and treatment-conditional benchmarks. We also evaluate the continuous treatment-conditional surface $f_{\mathrm{ATT}}(a, a') = \mathbb{E}\{Y^{(a)} \mid A = a'\}$ and the heterogeneous curve $v \mapsto f_{\mathrm{CATE}}(1, v)$ in designs adapted from Wu et al. (2024) and Abrevaya et al. (2015), respectively.

Table 2: Population dose-response MSE (lower is better) for Legalized Abortion & Crime, Deaner, and demand design datasets. *Div.*: no screened configuration converged.

Benchmark (sample size; scale)	PKDR	DRKPV	NMMR-U	NMMR-V	OUTCOMENET	TREATMENTNET	DR-PCL-NET
Abortion & Crime (1.5k; $10^3\times$)	20.335 ± 0.904	19.646 ± 2.837	*Div.*	19.774 ± 3.601	9.366 ± 1.246	51.455 ± 4.126	$\mathbf{8.630 \pm 0.615}$
Deaner–IM (3k; $10^3\times$)	2.756 ± 0.366	2.073 ± 0.203	1.405 ± 0.089	0.958 ± 0.132	0.857 ± 0.045	5.729 ± 0.125	$\mathbf{0.805 \pm 0.050}$
Deaner–IR (3k; $10^3\times$)	3.947 ± 0.638	1.028 ± 0.064	1.635 ± 0.021	1.809 ± 0.149	0.790 ± 0.132	9.101 ± 0.139	$\mathbf{0.640 \pm 0.127}$
Demand (2k)	107.86 ± 23.96	89.18 ± 3.74	25.58 ± 0.77	$\mathbf{22.39 \pm 1.61}$	32.95 ± 3.08	199.67 ± 15.77	26.74 ± 2.66
Demand (5k)	129.54 ± 17.23	85.60 ± 2.65	22.30 ± 1.34	33.99 ± 3.72	23.62 ± 2.09	174.11 ± 16.50	$\mathbf{18.98 \pm 1.95}$
Demand (10k)	172.52 ± 17.28	82.53 ± 1.21	19.12 ± 1.81	53.18 ± 3.77	22.42 ± 1.81	153.84 ± 5.57	$\mathbf{17.37 \pm 1.07}$
Demand (15k)	134.48 ± 14.36	83.82 ± 2.26	$\mathbf{14.55 \pm 1.16}$	42.70 ± 2.37	21.15 ± 1.93	144.97 ± 4.24	16.38 ± 1.37
Demand (20k)	208.81 ± 21.97	80.30 ± 0.90	16.39 ± 2.50	48.64 ± 3.46	20.18 ± 1.61	145.76 ± 4.40	$\mathbf{15.77 \pm 1.14}$

Overall results. DR-PCL-NET attains the lowest reported mean MSE in six of the eight settings in Table 2, including Abortion & Crime, both Deaner benchmarks, and three of the five Demand sample sizes; it remains competitive in the other two. Figure 2 shows the same overall pattern on the structured-treatment, treatment-conditional, and heterogeneous benchmarks, where DR-PCL-NET is consistently among the strongest methods across sample sizes.

7 CONCLUSION

We introduced a density-ratio-free neural treatment bridge estimator and combined it with a neural outcome bridge to estimate population, heterogeneous, and treatment-conditional response curves under unobserved confounding. The resulting doubly robust estimator requires no auxiliary density-ratio estimate and is suitable for continuous and structured treatment settings. We also established projected-norm error bounds showing that the doubly robust error is controlled by the final averaging/regression errors and the smaller of the outcome- and treatment-side projected bridge errors. Empirically, DR-PCL-NET achieves the best mean performance in most reported settings.

Limitations and future work. The statistical analyzes approximate empirical minimizers, but does not establish convergence guarantees for the implemented nonconvex stochastic training procedure. Deriving such optimizer-specific guarantees, potentially using techniques from Chen et al. (2025), is left for future work.

AI USE STATEMENT

AI-assisted tools were used to improve the language and organization of the manuscript, assist with code editing, and help clarify and cross-check mathematical concepts and derivations. All resulting text, code, and mathematical arguments were independently reviewed and verified by the authors, who remain fully responsible for the content of the manuscript and the accompanying implementation.

REPRODUCIBILITY STATEMENT

The implementation used to produce the reported experimental results is provided in the supplementary material.

REFERENCES

Jason Abrevaya, Yu-Chin Hsu, and Robert P. Lieli. Estimating conditional average treatment effects. *Journal of Business & Economic Statistics*, 33(4):485–505, 2015. doi: 10.1080/07350015.2014.975555. URL https://doi.org/10.1080/07350015.2014.975555.

Heejung Bang and James M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, December 2005. doi: 10.1111/j.1541-0420.2005.00377.x. PMID: 16401269.

Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. *The Annals of Statistics*, 33(4):1497 – 1537, 2005. doi: 10.1214/009053605000000282. URL https://doi.org/10.1214/009053605000000282.

Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney K Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Inference on strongly identified functionals of weakly identified functions. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, pp. qkaf075, 2025.

Richard Blundell, Xiaohong Chen, and Dennis Kristensen. Semi-nonparametric iv estimation of shape-invariant engel curves. *Econometrica*, 75(6):1613–1669, 2007.

Bariscan Bozkurt, Ben Deaner, Dimitri Meunier, Liyuan Xu, and Arthur Gretton. Density ratio-based proxy causal learning without density ratios. In *The 28th International Conference on Artificial Intelligence and Statistics*, 2025a.

Bariscan Bozkurt, Houssam Zenati, Dimitri Meunier, Liyuan Xu, and Arthur Gretton. Density ratio-free doubly robust proxy causal learning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b. URL https://openreview.net/forum?id=a9HOg4f9Gh.

Xiaohong Chen and Timothy M Christensen. Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric iv regression. *Quantitative Economics*, 9(1):39–84, 2018.

Xiaohong Chen and Markus Reiss. On rate optimality for ill-posed inverse problems in econometrics. *Econometric Theory*, 27(3):497–521, 2011.

Zonghao Chen, Atsushi Nitanda, Arthur Gretton, and Taiji Suzuki. Towards a unified analysis of neural networks in nonparametric instrumental variable regression: Optimization and generalization. *arXiv preprint arXiv:2511.14710*, 2025.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.

Stephan Clémençon, Gábor Lugosi, and Nicolas Vayatis. Ranking and empirical minimization of u-statistics. *The Annals of Statistics*, 36, 04 2008. doi: 10.1214/009052607000000910.

Kyle Colangelo and Ying-Ying Lee. Double debiased machine learning nonparametric inference with continuous treatments. *Journal of Business & Economic Statistics*, pp. 1–26, 2025.

Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Semiparametric proximal causal inference. *Journal of the American Statistical Association*, 119(546):1348–1359, 2024. doi: 10.1080/01621459.2023.2191817. URL https://doi.org/10.1080/01621459.2023.2191817.

Ben Deaner. Proxy controls and panel data, 2023. URL https://arxiv.org/abs/1810.00283.

III Donohue, John J. and Steven D. Levitt. The Impact of Legalized Abortion on Crime*. *The Quarterly Journal of Economics*, 116(2):379–420, 05 2001. ISSN 0033-5533. doi: 10.1162/00335530151144050. URL https://doi.org/10.1162/00335530151144050.

Dylan J Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *The Annals of Statistics*, 51(3): 879–908, 2023.

Jane Cooley Fruehwirth, Salvador Navarro, and Yuya Takahashi. How the Timing of Grade Retention Affects Outcomes: Identification and Estimation of Time-Varying Treatment Effects. *Journal of Labor Economics*, 34(4):979–1021, 2016. doi: 10.1086/686262. URL https://ideas.repec.org/a/ucp/jlabec/doi10.1086-686262.html.

Alexandre Galashov, Nathaël Da Costa, Liyuan Xu, Philipp Hennig, and Arthur Gretton. Closed-form last layer optimization, 2025. URL https://arxiv.org/abs/2510.04606.

Jason S. Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Counterfactual prediction with deep instrumental variables networks. *ArXiv*, abs/1612.09596, 2016. URL https://api.semanticscholar.org/CorpusID:12435773.

Jennifer L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20:217–240, 03 2011. doi: 10.1198/jcgs.2010.08162.

Guido Imbens. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics*, 2004.

Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International Conference on Machine Learning*, 2016.

Nathan Kallus, Xiaojie Mao, and Masatoshi Uehara. Causal inference under unmeasured confounding with negative controls: A minimax learning approach, 2021.

Edward H. Kennedy. Semiparametric doubly robust targeted double machine learning: A review. In Eric Laber, Bibhas Chakraborty, Erica E. M. Moodie, Tianxi Cai, and Mark J. van der Laan (eds.), *Handbook of Statistical Methods for Precision Medicine*, pp. 207–236. Chapman and Hall/CRC, 2024. doi: 10.48550/arXiv.2203.06469.

Edward H. Kennedy, Zongming Ma, Matthew D. McHugh, and Dylan S. Small. Non-parametric Methods for Doubly Robust Estimation of Continuous Treatment Effects. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 79(4):1229–1245, 09 2017.

Juno Kim, Dimitri Meunier, Arthur Gretton, Taiji Suzuki, and Zhu Li. Optimality and adaptivity of deep neural features for instrumental variable regression. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 94163–94206, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/ead8e195df586b9e7c10fcd114a6b9d1-Paper-Conference.pdf.

Benjamin Kompa, David Bellamy, Tom Kolokotrones, Andrew Beam, et al. Deep learning methods for proximal inference via maximum moment restriction. *Advances in Neural Information Processing Systems*, 2022.

R. Kress. *Linear Integral Equations*. Applied Mathematical Sciences. Springer New York, 2013. ISBN 9781461495932. URL https://books.google.co.uk/books?id=-L-8BAAAQBAJ.

Manabu Kuroki and Judea Pearl. Measurement bias and effect restoration in causal inference. *Biometrika*, 101(2):423–437, 2014.

Afsaneh Mastouri, Yuchen Zhu, Limor Gultchin, Anna Korba, Ricardo Silva, Matt J. Kusner, Arthur Gretton, and Krikamol Muandet. Proximal causal learning with kernels: Two-stage estimation and moment restriction. In *International Conference on Machine Learning*, 2021.

Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. https://github.com/deepmind/dsprites-dataset/, 2017.

Dimitri Meunier, Zhu Li, Tim Christensen, and Arthur Gretton. Nonparametric instrumental regression via kernel methods is minimax optimal. *arXiv preprint arXiv:2411.19653*, 2024.

Wang Miao, Zhi Geng, and Eric Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987—993, 2018.

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. The MIT Press, 2nd edition, 2018. ISBN 0262039400.

Whitney K. Newey. Nonparametric instrumental variables estimation. *American Economic Review*, 103(3):550–56, May 2013. doi: 10.1257/aer.103.3.550. URL https://www.aeaweb.org/articles?id=10.1257/aer.103.3.550.

Whitney K. Newey and James L. Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1555512.

Olav Reiersøl. *Confluence analysis by means of instrumental sets of variables*. PhD thesis, Almqvist & Wiksell, 1945.

James M. Robins, Miguel A. Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, September 2000. doi: 10.1097/00001648-200009000-00011. PMID: 10955408.

P. Rosenbaum and D. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.

Donald B. Rubin. Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American Statistical Association*, 75(371):591–593, 1980. ISSN 01621459, 1537274X. URL http://www.jstor.org/stable/2287653.

Zikai Shen, Zonghao Chen, Dimitri Meunier, Ingo Steinwart, Arthur Gretton, and Zhu Li. Nonparametric instrumental variable regression with observed covariates, 2025. URL https://arxiv.org/abs/2511.19404.

Rahul Singh. Kernel methods for unobserved confounding: Negative controls, proxies, and instruments, 2023. URL https://arxiv.org/abs/2012.10315.

James Stock and Francesco Trebbi. Retrospectives: Who invented instrumental variable regression? *Journal of Economic Perspectives - J ECON PERSPECT*, 17:177–194, 09 2003. doi: 10.1257/089533003769204416.

Erik Sverdrup and Yifan Cui. Proximal causal learning of conditional average treatment effects. In *International Conference on Machine Learning*, 2023. URL https://api.semanticscholar.org/CorpusID:258564946.

Eric J. Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal inference. *Statistical Science*, 39(3):375–390, 2024. doi: 10.1214/23-STS911.

Aad W. van der Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer, 1996.

Yong Wu, Yanwei Fu, Shouyan Wang, and Xinwei Sun. Doubly robust proximal causal learning for continuous treatments. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=TjGJFkU3xL.

Liyuan Xu and Arthur Gretton. A neural mean embedding approach for back-door and front-door adjustment. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=rLguqxYvYHB.

Liyuan Xu, Yutian Chen, Siddarth Srinivasan, Nando de Freitas, Arnaud Doucet, and Arthur Gretton. Learning deep features in instrumental variable regression. In *International Conference on Learning Representations*, 2021a. URL https://openreview.net/forum?id=sy4Kg_ZQmS7.

Liyuan Xu, Heishiro Kanagawa, and Arthur Gretton. Deep proxy causal learning and its application to confounded bandit policy evaluation. In *Advances in Neural Information Processing Systems*, 2021b. URL https://openreview.net/forum?id=0FDxsIEv9G.

Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, and Aidong Zhang. Representation learning for treatment effect estimation from observational data. In *Advances in Neural Information Processing Systems*, 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/a50abba8132a77191791390c3eb19fe7-Paper.pdf.

Houssam Zenati, Judith Abécassis, Julie Josse, and Bertrand Thirion. Double debiased machine learning for mediation analysis with continuous treatments. In *International Conference on Artificial Intelligence and Statistics*, volume 258, pp. 4150–4158, 2025.

Bingxi Zhang, Tao Shen, and Yifan Cui. Neural estimation of treatment bridge functions for proximal causal inference. *Statistical Analysis and Data Mining: An ASA Data Science Journal*, 18(5): e70045, 2025. doi: https://doi.org/10.1002/sam.70045. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/sam.70045.

APPENDIX CONTENTS

A Interpretation and scope of the completeness assumptions **16**

B Treatment bridge identification of the heterogeneous dose-response **18**

C Doubly robust identification: dose, heterogeneous, and conditional response curves **19**

D Influence-function representations for discrete targets **20**

E Neural mean embedding-based proxy causal learning for dose-response **25**

E.1 Dose-response curve estimation: outcome bridge approach 25

E.2 Dose-response curve estimation: treatment bridge approach 30

E.3 Dose-response curve estimation: doubly robust approach 33

F Neural mean embedding-based proxy causal learning for heterogeneous dose-response **34**

F.1 Heterogeneous dose-response estimation: outcome bridge method 34

F.2 Heterogeneous dose-response estimation: treatment bridge method 35

F.3 Heterogeneous dose-response estimation: doubly robust method 37

G Neural mean embedding-based proxy causal learning for treatment-conditional dose response **38**

G.1 Conditional dose-response estimation: outcome bridge method 38

G.2 Conditional dose-response estimation: treatment bridge method 39

G.3 Conditional dose-response estimation: doubly robust method 41

H Rademacher-complexity-based error bounds for the proposed estimators **41**

H.1 Background on Rademacher complexity analysis 42

H.2 Outcome bridge consistency for dose-response estimation 43

H.3 From projected outcome-bridge consistency to dose-response consistency 47

H.4 Treatment bridge consistency for dose-response estimation 49

H.5 Dose-response consistency of TREATMENTNET 54

H.6 Dose-response consistency of doubly robust method 57

H.7 Self-normalized density-ratio-free heterogeneous dose-response 61

H.8 Treatment-bridge consistency for conditional dose-response 63

H.9 Conditional dose-response consistency of TREATMENTNET 65

H.10 Conditional dose-response consistency of the doubly robust method 67

I Supplementary on numerical experiments **69**

I.1 Cross-fitting 69

I.2 Vectorized computation of U-statistics with einsum 69

I.3 Treatment bridge normalization 70

I.4 Benchmark descriptions and data-generating processes 72
I.5 Sensitivity to proxy informativeness 77
I.6 Neural architectures 78
I.7 Compute resources. 81

A INTERPRETATION AND SCOPE OF THE COMPLETENESS ASSUMPTIONS

The conditions in Section 3 are standard sufficient conditions for proximal identification results (Miao et al., 2018; Cui et al., 2024; Tchetgen Tchetgen et al., 2024). Closely related injectivity conditions also underlie nonparametric instrumental-variable regression (Newey, 2013; Newey & Powell, 2003). PCL also does not remove causal assumptions: it retains well-defined interventions, no interference, consistency, latent exchangeability, and overlap. Relative to adjustment on observed covariates, it replaces exchangeability given X with exchangeability given (U, X), together with the proxy conditional independence conditions $Y \perp Z \mid (U, A, X)$ and $W \perp (A, Z) \mid (U, X)$. These restrictions are most credible when temporal ordering, design information, or scientific knowledge supports Z as a treatment-inducing proxy (negative-control exposure) and W as an outcome-inducing proxy (negative-control outcome). Because U is unobserved, neither the proxy restrictions nor latent exchangeability and overlap are generally verifiable from the observed law alone. Below, we give the discussion on the completeness assumptions.

Completeness is an information condition. For fixed (a, x) and $Q \in \{Z, W\}$, define

$$T_{Q,a,x}\ell(q) := \mathbb{E}\{\ell(U) \mid Q = q, A = a, X = x\}.$$

Assumption 2.2 states that $\ker(T_{Q,a,x}) = \{0\}$. Hence, no nonzero square-integrable direction of U is invisible to the relevant proxy after conditioning on (A, X). The two clauses of Assumption 2.2 have route-specific roles: Z-completeness transfers the outcome-bridge equation to the latent law, whereas W-completeness does so for the treatment bridge.

In the discrete case, it can be understood as a rank-condition. Within an (a, x)-stratum, U has K states and Q has L states with positive probabilities. Then Q-completeness is equivalent to Miao et al. (2018); Cui et al. (2024)

$$\operatorname{rank}[\mathbb{P}(Q = q_j \mid U = u_k, A = a, X = x)]_{j \le L,\, k \le K} = K.$$

For continuous proxies, a standard sufficient condition is that, within each (a, x)-stratum, the conditional law of U given $Q = q$ belongs to a sufficiently rich exponential family as described in Newey & Powell (2003)[Theorem 2.2]. Specifically, suppose that, relative to a common dominating measure,

$$p_{U|Q,A,X}(u \mid q, a, x) = b_{a,x}(u) \exp\left\{\eta_{a,x}(q)^\top \tau_{a,x}(u) - \Psi_{a,x}(\eta_{a,x}(q))\right\},$$

where $b_{a,x}(u) > 0$, the sufficient statistic $\tau_{a,x}$ is one-to-one in u, and the support of the natural parameter $\eta_{a,x}(Q)$, conditional on $A = a, X = x$, contains a nonempty open set. Under the usual finite-exponential-moment conditions, these assumptions imply Q-completeness (Newey & Powell, 2003). Intuitively, $\tau_{a,x}$ must distinguish different latent states, while variation in Q must move the conditional law of U through sufficiently many directions to reveal each of them. For example, it holds if $U \mid Z = z, A = a, X = x \sim \mathcal{N}(\alpha_{a,x} + \Gamma_{a,x} z, \Sigma_{a,x})$, $\Sigma_{a,x}$ is positive definite, $\Gamma_{a,x}$ has row rank d_U, and the support of Z contains an open set; see Newey & Powell (2003)[Theorem 2.3].

Sufficient conditions for bridge existence. We next give a Picard-type sufficient condition for the existence of the treatment bridge. A parallel result for outcome-bridge existence is given by Xu et al. (2021b, Appendix B.1). The proxy restrictions and Z-completeness ensure that the target derivative belongs to the closure of the range of the treatment-bridge operator; the Picard condition below ensures that it belongs to the range itself.

For fixed a, x, recall that the conditional expectation operators are defined as

$$T_h : L^2(\mathbb{P}_{W|A=a,X=x}) \longrightarrow L^2(\mathbb{P}_{Z|A=a,X=x}), \quad T_h b = \mathbb{E}[b(W) \mid A = a, X = x, Z = \cdot]$$
$$T_\varphi : L^2(\mathbb{P}_{Z|A=a,X=x}) \longrightarrow L^2(\mathbb{P}_{W|A=a,X=x}), \quad T_\varphi g = \mathbb{E}[g(Z) \mid A = a, X = x, W = \cdot].$$

Iterated expectation gives

$$\langle T_h b, c\rangle_{L^2(\mathbb{P}_{Z|a,x})} = \mathbb{E}\{b(W)c(Z) \mid A = a, X = x\} = \langle b, T_\varphi c\rangle_{L^2(\mathbb{P}_{W|a,x})},$$

Hence $T_\varphi = T_h^*$, and therefore $T_h = T_\varphi^*$.

Solving the outcome-bridge equation requires its target regression to belong to $\mathrm{ran}(T_h)$. Likewise, solving the treatment-bridge equation requires the relevant Radon–Nikodym derivative to belong to $\mathrm{ran}(T_\varphi)$.

For the population and treatment-conditional targets, define the latent Radon–Nikodym derivatives

$$r^U_{0,\mathrm{ATE}}(a,x,u) := \frac{d\mathbb{P}_{X,U}}{d\mathbb{P}_{X,U|A=a}}(x,u),$$
$$r^U_{0,\mathrm{ATT}}(a,a',x,u) := \frac{d\mathbb{P}_{X,U|A=a'}}{d\mathbb{P}_{X,U|A=a}}(x,u).$$

For $\tau \in \{\mathrm{ATE}, \mathrm{ATT}\}$, let r^U_τ denote the corresponding latent derivative and let r_τ denote, respectively, $r_{0,\mathrm{ATE}}(a,x,\cdot)$ or $r_{0,\mathrm{ATT}}(a,a',x,\cdot)$.

Proposition A.1 (Sufficient conditions for treatment-bridge existence)**.** *Suppose Assumption 2.1 and Assumption 2.2 (i) hold. Assume that the numerator law in the definition of r^U_τ is absolutely continuous with respect to its denominator and $r^U_\tau \in L^2(\mathbb{P}_{U|A=a,X=x})$. Suppose T_φ, equivalently $T_h = T_\varphi^*$, is compact, and let $\{(\mu_j, e_j, f_j)\}_{j\geq 1}$ be a singular system for T_φ, where $\mu_j > 0$ and*

$$T_\varphi e_j = \mu_j f_j, \qquad T_h f_j = \mu_j e_j.$$

If

$$\sum_{j\geq 1} \frac{|\langle r_\tau, f_j\rangle_{L^2(\mathbb{P}_{W|A=a,X=x})}|^2}{\mu_j^2} < \infty, \tag{11}$$

then there exists $\varphi_0^\tau \in L^2(\mathbb{P}_{Z|A=a,X=x})$ satisfying

$$T_\varphi \varphi_0^\tau = r_\tau,$$

and a solution is given by

$$\varphi_0^\tau = \sum_{j\geq 1} \frac{\langle r_\tau, f_j\rangle_{L^2(\mathbb{P}_{W|A=a,X=x})}}{\mu_j} e_j. \tag{12}$$

Proof. Because the source and target laws extend through the same conditional law of W given (U,X), the projection property of Radon–Nikodym derivatives and the proxy restrictions give

$$r_\tau(W) = \mathbb{E}\{r^U_\tau(U) \mid A=a, X=x, W\} \quad \mathbb{P}_{W|A=a,X=x}\text{-a.e.}$$

Thus $r_\tau \in L^2(\mathbb{P}_{W|A=a,X=x})$ by conditional Jensen. Let $b \in \ker(T_h)$ and set $\ell_b(U) := \mathbb{E}\{b(W) \mid A=a, X=x, U\}$. Conditional Jensen gives $\ell_b \in L^2(\mathbb{P}_{U|A=a,X=x})$. Moreover,

$$0 = (T_h b)(Z) = \mathbb{E}\{\ell_b(U) \mid A=a, X=x, Z\},$$

where the second equality follows from $W \perp Z \mid (U,A,X)$. Assumption 2.2 (i) therefore implies $\ell_b(U) = 0$ almost surely. Consequently,

$$\begin{aligned}\langle r_\tau, b\rangle_{L^2(\mathbb{P}_{W|A=a,X=x})} &= \mathbb{E}\{r^U_\tau(U) b(W) \mid A=a, X=x\} \\ &= \mathbb{E}\{r^U_\tau(U)\ell_b(U) \mid A=a, X=x\} = 0.\end{aligned}$$

Hence $r_\tau \in \ker(T_h)^\perp = \ker(T_\varphi^*)^\perp$. Equation 11 and Picard's theorem now imply $r_\tau \in \mathrm{ran}(T_\varphi)$ and give the solution in Eq. 12 by (Kress, 2013)[Theorem 15.18]. □

Any jointly measurable and suitably square-integrable ATE solution gives a CATE bridge through Proposition 3.4 whenever $0 < D_0(a,v) < \infty$.For ATT, suppose that, for almost every (a,a',x), a solution of the corresponding conditional bridge equation can be chosen so that the resulting map $(a,a',x,z) \mapsto \varphi_0^{\mathrm{ATT}}(a,a',x,z)$ is jointly measurable and square-integrable. Then it solves the joint surface equation $T_{\mathrm{ATT}}\varphi_0^{\mathrm{ATT}} = r_{0,\mathrm{ATT}}$.

B TREATMENT BRIDGE IDENTIFICATION OF THE HETEROGENEOUS DOSE-RESPONSE

In this appendix, we prove the treatment-bridge identification results for the heterogeneous dose-response $f_{\mathrm{CATE}}(a, v)$ stated in Theorem 3.2 and Proposition 3.4. The corresponding ATE and ATT statements follow from Bozkurt et al. (2025a). To match the notation of the main text, we write $X = (S, V)$ and express changes of measure through probability laws and Radon–Nikodym derivatives; density or mass-function quotients are used only as dominated-case interpretations.

Proof of the treatment bridge identification for CATE statement in Theorem 3.2. Fix $a \in \mathcal{A}$ and $v \in \mathcal{V}$, and write $\varphi_0 = \varphi_0^{\mathrm{CATE}}$ for brevity. By assumption, φ_0 satisfies

$$\mathbb{E}[\varphi_0(a, v, S, Z) \mid A = a, S, V = v, W] = r_{0,\mathrm{CATE}}(a, v, S, W) \qquad \mathbb{P}_{W|A=a,S,V=v}\text{-a.e.} \tag{13}$$

Define the latent derivative

$$r^U_{0,\mathrm{CATE}}(a, v, s, u) := \frac{d\mathbb{P}_{S,U|V=v}}{d\mathbb{P}_{S,U|A=a,V=v}}(s, u), \qquad r^U_{0,\mathrm{CATE}}(a, v, \cdot, \cdot) \in L^2(\mathbb{P}_{S,U|A=a,V=v}).$$

Using $W \perp A \mid U, X$ from Assumption 2.1, with $X = (S, V)$, and the projection property of Radon-Nikodym derivatives give

$$\mathbb{E}\{r^U_{0,\mathrm{CATE}}(a, v, S, U) \mid A = a, S, V = v, W\} = r_{0,\mathrm{CATE}}(a, v, S, W).$$

Comparing this identity with Eq. 13, followed by Assumption 2.2 (ii), yields

$$\mathbb{E}\{\varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, S, V = v, U\} = r^U_{0,\mathrm{CATE}}(a, v, S, U). \tag{14}$$

Now define $m_a(S, U, V) \coloneqq \mathbb{E}[Y \mid A = a, S, V, U]$. Then

$$\begin{aligned}
&\mathbb{E}[Y\varphi_0(a, v, S, Z) \mid A = a, V = v] = \mathbb{E}[\mathbb{E}[Y\varphi_0(a, v, S, Z) \mid A = a, S, V = v, U] \mid A = a, V = v] \\
&\quad = \mathbb{E}[\mathbb{E}[Y \mid A = a, S, V = v, U]\, \mathbb{E}[\varphi_0(a, v, S, Z) \mid A = a, S, V = v, U] \mid A = a, V = v] \\
&\quad = \mathbb{E}\big[m_a(S, U, v)\, r^U_{0,\mathrm{CATE}}(a, v, S, U) \,\big|\, A = a, V = v\big] = \mathbb{E}[m_a(S, U, v) \mid V = v],
\end{aligned}$$

where the second equality uses $Y \perp Z \mid U, A, X$ from Assumption 2.1, with $X = (S, V)$, the third equality uses Equation 14, and the last equality is the projection property of Radon-Nikodym derivative.

Finally, by consistency and latent exchangeability conditional on $(X, U) = (S, V, U)$,

$$\begin{aligned}
\mathbb{E}[m_a(S, U, v) \mid V = v] &= \mathbb{E}[Y^{(a)} \mid V = v] = f_{\mathrm{CATE}}(a, v) \\
&= \mathbb{E}[Y\varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, V = v].
\end{aligned}$$

□

We next prove our alternative identification of $\varphi_0^{\mathrm{CATE}}$ using the treatment dose-response bridge function.

Proof of Proposition 3.4. Let $X = (S, V)$. By the ATE bridge equation and iterated expectations,

$$\begin{aligned}
D_0(a, v) &= \mathbb{E}\{\varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a, V = v\} \\
&= \mathbb{E}\{r_{0,\mathrm{ATE}}(a, X, W) \mid A = a, V = v\} = \frac{d(\mathbb{P}_A \otimes \mathbb{P}_V)}{d\mathbb{P}_{A,V}}(a, v).
\end{aligned}$$

The last equality is the conditional-expectation property of Radon–Nikodym derivatives. The chain rule therefore yields

$$r_{0,\mathrm{ATE}}(a, (s, v), w) = D_0(a, v) r_{0,\mathrm{CATE}}(a, v, s, w) \quad \text{a.e.}$$

If $0 < D_0(a, v) < \infty$, taking conditional expectations shows that $\varphi_0^{\mathrm{ATE}}(a, (s, v), z)/D_0(a, v)$ solves Eq. 3. The treatment-bridge identification of CATE then gives the ratio formula in Proposition 3.4.

It remains to verify self-normalized double robustness. If $h = h_0$, the conditional outcome-bridge moment makes the residual numerator in the expression of $\Psi^{\mathrm{SN}}_{\mathrm{CATE}}(a, v; h, \varphi)$ equal to zero. If instead $\varphi = \varphi_0^{\mathrm{ATE}}$, division by $D_0(a, v)$ produces the valid CATE bridge above, so the native CATE doubly robust identity applies. Thus the self-normalized functional equals $f_{\mathrm{CATE}}(a, v)$ under either route. □

C Doubly robust identification: dose, heterogeneous, and conditional response curves

In this section, we prove the doubly robust identification results in Theorem 3.3. We show that each doubly robust formula remains valid if either the outcome bridge or the corresponding treatment bridge is correctly specified.

Proof. We prove the result separately for the ATE, CATE, and ATT. Throughout, we use the notation $X = (S, V)$, and we write $h_0(a, X, W)$ for the outcome bridge.

The ATE formula is the covariate-adjusted version of Bozkurt et al. (2025b)[Theorem 2.7]. So, below, we prove CATE and ATT extensions.

CATE. Consider

$$\Psi_{\mathrm{CATE}}(a, v; h, \varphi) = \mathbb{E}[\varphi(a, v, S, Z)\{Y - h(a, X, W)\} \mid A = a, V = v] + \mathbb{E}[h(a, X, W) \mid V = v].$$

If $h = h_0$, then

$$\begin{aligned}\Psi_{\mathrm{CATE}}(a, v; h_0, \varphi) &= \mathbb{E}[\varphi(a, v, S, Z)\,\mathbb{E}[Y - h_0(a, X, W) \mid A = a, S, V = v, Z]|A = a, V = v] \\ &\quad + \mathbb{E}[h_0(a, X, W) \mid V = v] \\ &= \mathbb{E}[h_0(a, X, W) \mid V = v] = f_{\mathrm{CATE}}(a, v),\end{aligned}$$

where the second equality again follows from Equation 1, since $X = (S, V)$, and the last equality follows from the outcome bridge identification of CATE.

If $\varphi = \varphi_0^{\mathrm{CATE}}$, then

$$\begin{aligned}&\mathbb{E}[\varphi_0^{\mathrm{CATE}}(a, v, S, Z)h(a, X, W) \mid A = a, V = v] \\ &\quad = \mathbb{E}\big[\mathbb{E}[\varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, S, V = v, W]\, h(a, X, W)\big|A = a, V = v\big] \\ &\quad = \mathbb{E}[r_{0,\mathrm{CATE}}(a, v, S, W)\, h(a, X, W)|A = a, V = v] = \mathbb{E}[h(a, X, W) \mid V = v].\end{aligned}$$

Hence,

$$\begin{aligned}\Psi_{\mathrm{CATE}}(a, v; h, \varphi_0^{\mathrm{CATE}}) &= \mathbb{E}[Y\varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, V = v] \\ &\quad - \mathbb{E}[\varphi_0^{\mathrm{CATE}}(a, v, S, Z)h(a, X, W) \mid A = a, V = v] \\ &\quad + \mathbb{E}[h(a, X, W) \mid V = v] \\ &= \mathbb{E}[Y\varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, V = v] = f_{\mathrm{CATE}}(a, v),\end{aligned}$$

where the last equality follows from the treatment bridge identification of CATE.

ATT. Consider

$$\Psi_{\mathrm{ATT}}(a, a'; h, \varphi) = \mathbb{E}[\varphi(a, a', X, Z)\{Y - h(a, X, W)\} \mid A = a] + \mathbb{E}[h(a, X, W) \mid A = a'].$$

If $h = h_0$, then

$$\begin{aligned}\Psi_{\mathrm{ATT}}(a, a'; h_0, \varphi) &= \mathbb{E}[\varphi(a, a', X, Z)\,\mathbb{E}[Y - h_0(a, X, W) \mid A = a, X, Z]|A = a] \\ &\quad + \mathbb{E}[h_0(a, X, W) \mid A = a'] \\ &= \mathbb{E}[h_0(a, X, W) \mid A = a'] = f_{\mathrm{ATT}}(a, a'),\end{aligned}$$

by Equation 1 and outcome bridge identification of ATT.

If $\varphi = \varphi_0^{\mathrm{ATT}}$, then

$$\mathbb{E}[\varphi_0^{\mathrm{ATT}}(a, a', X, Z)h(a, X, W) \mid A = a]$$

$$= \mathbb{E}\big[\mathbb{E}[\varphi_0^{\mathrm{ATT}}(a,a',X,Z) \mid A=a,X,W]\, h(a,X,W)\big|A=a\big]$$
$$= \mathbb{E}[r_{0,\mathrm{ATT}}(a,a',X,W)\, h(a,X,W)|A=a] = \mathbb{E}[h(a,X,W) \mid A=a'].$$

Therefore,

$$\begin{aligned}\Psi_{\mathrm{ATT}}(a,a';h,\varphi_0^{\mathrm{ATT}}) &= \mathbb{E}[Y\varphi_0^{\mathrm{ATT}}(a,a',X,Z) \mid A=a] \\ &\quad - \mathbb{E}[\varphi_0^{\mathrm{ATT}}(a,a',X,Z)h(a,X,W) \mid A=a] + \mathbb{E}[h(a,X,W) \mid A=a'] \\ &= \mathbb{E}[Y\varphi_0^{\mathrm{ATT}}(a,a',X,Z) \mid A=a] = f_{\mathrm{ATT}}(a,a'),\end{aligned}$$

where the last equality follows from treatment bridge identification of ATT.

This proves the double robustness of all three identification formulas. □

D Influence-function representations for discrete targets

In this section, we derive and formally establish the Efficient Influence Functions (EIFs) for the bridge-based target causal functions. We focus specifically on the discrete treatment setting, where A takes values in a finite set $\mathcal{A}$, and Y is scalar. The restriction to discrete A is important: pointwise dose-response functionals are pathwise differentiable in the usual semiparametric sense for atoms $A=a$, whereas continuous-treatment pointwise effects require a different local or smoothed efficiency theory (Kennedy et al., 2017; Colangelo & Lee, 2025; Wu et al., 2024; Zenati et al., 2025). For the CATE result below we also take V to be discrete, so that $\mathbb{P}(V=v)>0$.

Throughout this section, $O=(Y,A,X,Z,W)$, and for CATE we write $X=(S,V)$ and use $h_0(a,S,V,W)$ as shorthand for $h_0(a,X,W)$. Define

$$\pi_a := \mathbb{P}(A=a), \qquad \pi_v := \mathbb{P}(V=v), \qquad \pi_{a|v} := \mathbb{P}(A=a \mid V=v).$$

For positive-mass treatment and conditioning cells, the Radon–Nikodym derivatives defined in the main text can be written as

$$r_{0,\mathrm{ATE}}(a,\cdot,\cdot) = \frac{d\mathbb{P}_{X,W}}{d\mathbb{P}_{X,W|A=a}},$$
$$r_{0,\mathrm{CATE}}(a,v,\cdot,\cdot) = \frac{d\mathbb{P}_{S,W|V=v}}{d\mathbb{P}_{S,W|A=a,V=v}},$$
$$r_{0,\mathrm{ATT}}(a,a',\cdot,\cdot) = \frac{d\mathbb{P}_{X,W|A=a'}}{d\mathbb{P}_{X,W|A=a}}.$$

All expectations are taken under the true observed law unless otherwise stated.

Semiparametric regularity. In addition to the bridge existence, uniqueness, square-integrability, and positivity conditions stated below, we use the same tangent-space regularity condition as in (Cui et al., 2024) and (Bozkurt et al., 2025b). Namely, the relevant conditional expectation operators are assumed to have dense range; the stronger surjectivity condition on the operators and their adjoints used by (Cui et al., 2024) is sufficient. For the ATE and ATT formulas, one may take

$$Tg := \mathbb{E}[g(W,A,X) \mid Z,A,X], \qquad T^*r := \mathbb{E}[r(Z,A,X) \mid W,A,X],$$

with $T: L^2(\mathbb{P}_{W,A,X}) \to L^2(\mathbb{P}_{Z,A,X})$. For the CATE formula, with $X=(S,V)$, take

$$T_{\mathrm{C}}g := \mathbb{E}[g(W,A,S,V) \mid Z,A,S,V], \qquad T_{\mathrm{C}}^*r := \mathbb{E}[r(Z,A,S,V) \mid W,A,S,V].$$

This condition is used only to conclude that the influence functions derived below are canonical gradients, hence efficient.

Theorem D.1 (Efficient influence functions for discrete targets)**.** *Let Assumptions 2.1 and 2.2 hold. Suppose that there exists a unique outcome bridge h_0 satisfying Eq. 1, and unique treatment bridges $\varphi_0^{\mathrm{ATE}}, \varphi_0^{\mathrm{CATE}}, \varphi_0^{\mathrm{ATT}}$ satisfying the treatment bridge equations 2. Assume, for the targets under consideration, that $\mathbb{P}_{X,W} \ll \mathbb{P}_{X,W|A=a}$, $\mathbb{P}_{S,W|V=v} \ll \mathbb{P}_{S,W|A=a,V=v}$, and $\mathbb{P}_{X,W|A=a'} \ll \mathbb{P}_{X,W|A=a}$, respectively. In particular, $\pi_a>0$, $\pi_{a'}>0$, $\pi_v>0$, and $\pi_{a|v}>0$ whenever the corresponding target is considered. Assume also the semiparametric regularity condition above.*

Then $f_{\mathrm{ATE}}(a)$, $f_{\mathrm{ATT}}(a, a')$, and, for discrete V, $f_{\mathrm{CATE}}(a, v)$ are pathwise differentiable. Their EIFs are given by

(i) ATE. *For any $a \in \mathcal{A}$,*

$$\psi^{\mathrm{EIF}}_{\mathrm{ATE}}(O; a) = \frac{\mathbb{1}\{A = a\}}{\pi_a}\, \varphi_0^{\mathrm{ATE}}(a, X, Z)\big\{Y - h_0(A, X, W)\big\} + h_0(a, X, W) - f_{\mathrm{ATE}}(a). \quad (15)$$

(ii) ATT / conditional counterfactual mean. *For any $a, a' \in \mathcal{A}$,*

$$\psi^{\mathrm{EIF}}_{\mathrm{ATT}}(O; a, a') = \frac{\mathbb{1}\{A = a\}}{\pi_a}\, \varphi_0^{\mathrm{ATT}}(a, a', X, Z)\big\{Y - h_0(A, X, W)\big\} \quad (16)$$

$$+ \frac{\mathbb{1}\{A = a'\}}{\pi_{a'}}\big\{h_0(a, X, W) - f_{\mathrm{ATT}}(a, a')\big\}. \quad (17)$$

(iii) CATE, with discrete V. *For any $a \in \mathcal{A}$ and any v with $\pi_v > 0$,*

$$\psi^{\mathrm{EIF}}_{\mathrm{CATE}}(O; a, v) = \frac{\mathbb{1}\{V = v\}}{\pi_v}\left[\frac{\mathbb{1}\{A = a\}}{\pi_{a|v}}\, \varphi_0^{\mathrm{CATE}}(a, v, S, Z)\big\{Y - h_0(A, S, V, W)\big\}\right.$$

$$\left. + h_0(a, S, v, W) - f_{\mathrm{CATE}}(a, v)\right]. \quad (18)$$

Proof. The ATE formula is the covariate-adjusted version of Bozkurt et al. (2025b, Theorem B.4). It is also obtained from the EIF of Cui et al. (2024, Theorem 3.1) after replacing the classical treatment bridge q_0, which satisfies $\mathbb{E}[q_0(a, X, Z) \mid A = a, X, W] = \pi_a^{-1} r_{0,\mathrm{ATE}}(a, X, W)$, by the normalized bridge $\varphi_0^{\mathrm{ATE}}(a, X, Z) = \pi_a q_0(a, X, Z)$. We give the derivations for ATT and CATE because the conditioning events change the derivative and the normalization.

Let $\{\mathbb{P}_t : t \in (-\delta, \delta)\}$ be a regular parametric submodel through the true law $\mathbb{P}_0 = \mathbb{P}$, and let

$$\dot{\ell}(O) := \frac{\partial}{\partial t} \log p_t(O)\Big|_{t=0}$$

be its score. Let h_t denote a differentiable path of outcome bridges along the submodel, and write $\dot{h} := \partial h_t / \partial t|_{t=0}$.

ATT. Fix $a, a' \in \mathcal{A}$ and write

$$\psi := f_{\mathrm{ATT}}(a, a') = \mathbb{E}[h_0(a, X, W) \mid A = a'].$$

Define

$$N_t := \mathbb{E}_t[\mathbb{1}\{A = a'\} h_t(a, X, W)], \qquad D_t := \mathbb{P}_t(A = a'),$$

so that $\psi_t = N_t / D_t$. Differentiating the ratio at $t = 0$ gives

$$\frac{\partial \psi_t}{\partial t}\Big|_{t=0} = \frac{1}{\pi_{a'}}\left[\frac{\partial N_t}{\partial t}\Big|_{t=0} - \psi\, \frac{\partial D_t}{\partial t}\Big|_{t=0}\right]. \quad (19)$$

The denominator derivative is

$$\frac{\partial D_t}{\partial t}\Big|_{t=0} = \mathbb{E}[\mathbb{1}\{A = a'\}\dot{\ell}(O)].$$

For the numerator,

$$\frac{\partial N_t}{\partial t}\Big|_{t=0} = \mathbb{E}[\mathbb{1}\{A = a'\} h_0(a, X, W)\dot{\ell}(O)] + \mathbb{E}[\mathbb{1}\{A = a'\}\dot{h}(a, X, W)]. \quad (20)$$

It remains to express the second term in Eq. 20 as an inner product with the score. The bridge equation along the submodel is

$$\mathbb{E}_t[Y - h_t(A, X, W) \mid A, X, Z] = 0.$$

Differentiating this equation and using the conditional mean-zero bridge residual gives, on the event $\{A = a\}$,

$$\mathbb{E}[\dot{h}(a, X, W) \mid A = a, X, Z] = \mathbb{E}[\{Y - h_0(a, X, W)\}\dot{\ell}(O) \mid A = a, X, Z]. \tag{21}$$

The ATT treatment bridge satisfies

$$\mathbb{E}[\varphi_0^{\mathrm{ATT}}(a, a', X, Z) \mid A = a, X, W] = r_{0,\mathrm{ATT}}(a, a', X, W).$$

Therefore, for any measurable $g(X, W)$ for which the following expectations are finite,

$$\begin{aligned}\mathbb{E}[\mathbb{1}\{A = a\}\varphi_0^{\mathrm{ATT}}(a, a', X, Z)g(X, W)] &= \pi_a \mathbb{E}[\varphi_0^{\mathrm{ATT}}(a, a', X, Z)g(X, W) \mid A = a] \\ &= \pi_a \mathbb{E}[r_{0,\mathrm{ATT}}(a, a', X, W)g(X, W) \mid A = a] \\ &= \pi_a \mathbb{E}[g(X, W) \mid A = a'].\end{aligned} \tag{22}$$

Taking $g = \dot{h}(a, \cdot, \cdot)$ in Eq. 22 yields

$$\mathbb{E}[\mathbb{1}\{A = a'\}\dot{h}(a, X, W)] = \frac{\pi_{a'}}{\pi_a}\mathbb{E}[\mathbb{1}\{A = a\}\varphi_0^{\mathrm{ATT}}(a, a', X, Z)\dot{h}(a, X, W)].$$

Combining this identity with Eq. 21 gives

$$\frac{1}{\pi_{a'}}\mathbb{E}[\mathbb{1}\{A = a'\}\dot{h}(a, X, W)] = \mathbb{E}\left[\frac{\mathbb{1}\{A = a\}}{\pi_a}\varphi_0^{\mathrm{ATT}}(a, a', X, Z)\{Y - h_0(a, X, W)\}\dot{\ell}(O)\right].$$

Substituting this display and Eq. 20 into Eq. 19 gives

$$\left.\frac{\partial \psi_t}{\partial t}\right|_{t=0} = \mathbb{E}[\psi_{\mathrm{ATT}}^{\mathrm{EIF}}(O; a, a')\dot{\ell}(O)].$$

The candidate in Eq. 17 is mean zero because the first term has conditional mean zero given $(A = a, X, Z)$, and the second term is centered conditional on $A = a'$.

It remains to check efficiency. Let $\mathcal{F} = (Z, A, X)$ and $\varepsilon := Y - h_0(A, X, W)$. Under the bridge-restricted semiparametric model, the tangent space can be written as $\Lambda = \Lambda_1 + \Lambda_2$, where

$$\Lambda_1 = \{s_1(\mathcal{F}) \in L^2(\mathbb{P}) : \mathbb{E}[s_1] = 0\},$$

and

$$\Lambda_2 = \left\{s_2 \in L^2(\mathbb{P}) : \mathbb{E}[s_2 \mid \mathcal{F}] = 0,\ \mathbb{E}[\varepsilon s_2 \mid \mathcal{F}] \in \overline{\mathcal{R}(T)}\right\}.$$

Set $H_{\mathrm{ATT}}(X, W) := h_0(a, X, W) - \psi$, and decompose $G_{\mathrm{ATT}} := \psi_{\mathrm{ATT}}^{\mathrm{EIF}}$ as

$$G_{\mathrm{ATT},1} := \mathbb{E}[G_{\mathrm{ATT}} \mid \mathcal{F}] = \frac{\mathbb{1}\{A = a'\}}{\pi_{a'}}\mathbb{E}[H_{\mathrm{ATT}}(X, W) \mid \mathcal{F}], \qquad G_{\mathrm{ATT},2} := G_{\mathrm{ATT}} - G_{\mathrm{ATT},1}.$$

Then $G_{\mathrm{ATT},1} \in \Lambda_1$. Also $\mathbb{E}[G_{\mathrm{ATT},2} \mid \mathcal{F}] = 0$. Moreover $\mathbb{E}[\varepsilon G_{\mathrm{ATT},2} \mid \mathcal{F}]$ is a square-integrable $\mathcal{F}$-measurable function; by the dense-range condition for T, it belongs to $\overline{\mathcal{R}(T)}$. Hence $G_{\mathrm{ATT},2} \in \Lambda_2$, so $G_{\mathrm{ATT}} \in \Lambda$. Since it represents the pathwise derivative and belongs to the tangent space, it is the canonical gradient.

CATE. Fix $a \in \mathcal{A}$ and v with $\pi_v > 0$. Write

$$\psi := f_{\mathrm{CATE}}(a, v) = \mathbb{E}[h_0(a, S, v, W) \mid V = v].$$

Define

$$N_t := \mathbb{E}_t[\mathbb{1}\{V = v\}h_t(a, S, v, W)], \qquad D_t := \mathbb{P}_t(V = v),$$

so that $\psi_t = N_t/D_t$. The quotient rule gives

$$\left.\frac{\partial \psi_t}{\partial t}\right|_{t=0} = \frac{1}{\pi_v}\left[\left.\frac{\partial N_t}{\partial t}\right|_{t=0} - \psi \left.\frac{\partial D_t}{\partial t}\right|_{t=0}\right], \tag{23}$$

with

$$\left.\frac{\partial D_t}{\partial t}\right|_{t=0} = \mathbb{E}[\mathbb{1}\{V = v\}\dot{\ell}(O)]$$

and

$$\left.\frac{\partial N_t}{\partial t}\right|_{t=0} = \mathbb{E}[\mathbb{1}\{V=v\}h_0(a,S,v,W)\dot{\ell}(O)] + \mathbb{E}[\mathbb{1}\{V=v\}\dot{h}(a,S,v,W)]. \tag{24}$$

The bridge equation along the submodel is

$$\mathbb{E}_t[Y - h_t(A,S,V,W) \mid A,S,V,Z] = 0.$$

Differentiating at $t=0$ gives

$$\mathbb{E}[\dot{h}(A,S,V,W) \mid A,S,V,Z] = \mathbb{E}[\{Y - h_0(A,S,V,W)\}\dot{\ell}(O) \mid A,S,V,Z]. \tag{25}$$

The CATE treatment bridge satisfies

$$\mathbb{E}[\varphi_0^{\mathrm{CATE}}(a,v,S,Z) \mid A=a,S,V=v,W] = r_{0,\mathrm{CATE}}(a,v,S,W).$$

Consequently, for any measurable $g(S,W)$ for which the following expectations are finite,

$$\begin{aligned}
&\mathbb{E}[\mathbb{1}\{V=v\}\mathbb{1}\{A=a\}\varphi_0^{\mathrm{CATE}}(a,v,S,Z)g(S,W)] \\
&\quad = \pi_v\pi_{a|v}\mathbb{E}[r_{0,\mathrm{CATE}}(a,v,S,W)g(S,W) \mid A=a,V=v] \\
&\quad = \pi_v\pi_{a|v}\mathbb{E}[g(S,W) \mid V=v] \\
&\quad = \pi_{a|v}\mathbb{E}[\mathbb{1}\{V=v\}g(S,W)].
\end{aligned} \tag{26}$$

Taking $g = \dot{h}(a,\cdot,v,\cdot)$ in Eq. 26 gives

$$\mathbb{E}[\mathbb{1}\{V=v\}\dot{h}(a,S,v,W)] = \frac{1}{\pi_{a|v}}\mathbb{E}[\mathbb{1}\{V=v\}\mathbb{1}\{A=a\}\varphi_0^{\mathrm{CATE}}(a,v,S,Z)\dot{h}(a,S,v,W)].$$

Combining this equation with Eq. 25 yields

$$\frac{1}{\pi_v}\mathbb{E}[\mathbb{1}\{V=v\}\dot{h}(a,S,v,W)] = \mathbb{E}\left[\frac{\mathbb{1}\{V=v\}}{\pi_v}\frac{\mathbb{1}\{A=a\}}{\pi_{a|v}}\varphi_0^{\mathrm{CATE}}(a,v,S,Z)\{Y - h_0(A,S,V,W)\}\dot{\ell}(O)\right].$$

Substituting this identity and Eq. 24 into Eq. 23 gives

$$\left.\frac{\partial \psi_t}{\partial t}\right|_{t=0} = \mathbb{E}[\psi_{\mathrm{CATE}}^{\mathrm{EIF}}(O;a,v)\dot{\ell}(O)].$$

Thus Eq. 18 is an influence function. It is mean zero by the outcome bridge equation and by the centering of $h_0(a,S,v,W) - f_{\mathrm{CATE}}(a,v)$ conditional on $V=v$.

For efficiency, let $\mathcal{F}_{\mathrm{C}} := (Z,A,S,V)$ and $\varepsilon := Y - h_0(A,S,V,W)$. The tangent space is $\Lambda_{\mathrm{C}} = \Lambda_{\mathrm{C},1} + \Lambda_{\mathrm{C},2}$, where

$$\Lambda_{\mathrm{C},1} = \{s_1(\mathcal{F}_{\mathrm{C}}) \in L^2(\mathbb{P}) : \mathbb{E}[s_1] = 0\},$$

and

$$\Lambda_{\mathrm{C},2} = \left\{s_2 \in L^2(\mathbb{P}) : \mathbb{E}[s_2 \mid \mathcal{F}_{\mathrm{C}}] = 0,\ \mathbb{E}[\varepsilon s_2 \mid \mathcal{F}_{\mathrm{C}}] \in \overline{\mathcal{R}(T_{\mathrm{C}})}\right\}.$$

Set $H_{\mathrm{C}}(S,W) := h_0(a,S,v,W) - \psi$, and write $G_{\mathrm{C}} := \psi_{\mathrm{CATE}}^{\mathrm{EIF}}(O;a,v)$. Then

$$G_{\mathrm{C},1} := \mathbb{E}[G_{\mathrm{C}} \mid \mathcal{F}_{\mathrm{C}}] = \frac{\mathbb{1}\{V=v\}}{\pi_v}\mathbb{E}[H_{\mathrm{C}}(S,W) \mid \mathcal{F}_{\mathrm{C}}], \qquad G_{\mathrm{C},2} := G_{\mathrm{C}} - G_{\mathrm{C},1}.$$

The first component belongs to $\Lambda_{\mathrm{C},1}$. The second satisfies $\mathbb{E}[G_{\mathrm{C},2} \mid \mathcal{F}_{\mathrm{C}}] = 0$, and $\mathbb{E}[\varepsilon G_{\mathrm{C},2} \mid \mathcal{F}_{\mathrm{C}}]$ is a square-integrable $\mathcal{F}_{\mathrm{C}}$-measurable function. By the dense-range condition for T_{C}, this conditional moment belongs to $\overline{\mathcal{R}(T_{\mathrm{C}})}$. Hence $G_{\mathrm{C},2} \in \Lambda_{\mathrm{C},2}$, so $G_{\mathrm{C}} \in \Lambda_{\mathrm{C}}$. Since it represents the pathwise derivative and lies in the tangent space, it is the canonical gradient. This proves Eq. 18 and completes the proof. □

Remark D.2 (Relation to classical proximal ATE identification)**.** *In the discrete-treatment case, our ATE representation is a normalized version of Cui et al. (2024). For a fixed treatment level a, their treatment bridge may be written as*

$$\mathbb{E}[q_0^{\mathrm{ATE}}(a,X,Z) \mid A=a,X,W] = \frac{1}{\pi_a}r_{0,\mathrm{ATE}}(a,X,W).$$

Their corresponding doubly robust representation of the counterfactual mean is

$$f^{\mathrm{DR}}_{\mathrm{ATE}}(a; h_0, q_0^{\mathrm{ATE}}) = \mathbb{E}\big[\mathbb{1}\{A=a\} q_0^{\mathrm{ATE}}(a,X,Z)\{Y - h_0(A,X,W)\}\big] + \mathbb{E}[h_0(a,X,W)]. \quad (27)$$

Equivalently, its influence-function moment is

$$\psi^{\mathrm{EIF}}_{q_0^{\mathrm{ATE}}}(O;a) = \mathbb{1}\{A=a\} q_0^{\mathrm{ATE}}(a,X,Z)\{Y - h_0(A,X,W)\} + h_0(a,X,W) - f_{\mathrm{ATE}}(a).$$

Our bridge in Eq. 2 satisfies

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a,X,Z) \mid A=a, X, W] = r_{0,\mathrm{ATE}}(a,X,W).$$

Thus, when the bridge solution is unique, $\varphi_0^{\mathrm{ATE}}(a,X,Z) = \pi_a q_0^{\mathrm{ATE}}(a,X,Z)$. Substituting this into Eq. 27 gives

$$\mathbb{E}\big[\varphi_0^{\mathrm{ATE}}(a,X,Z)\{Y - h_0(a,X,W)\} \mid A=a\big] + \mathbb{E}[h_0(a,X,W)],$$

which is exactly Eq. 5. Therefore the normalized bridge does not change the discrete EIF; it rewrites the same moment as a conditional expectation given $A=a$. This form is the one used by our neural mean-embedding estimator and avoids an explicit sparse indicator function when the target is indexed by a continuous or structured treatment value.

Remark D.3 (Relation to the classical proximal ATT EIF)**.** *The ATT formula above is not obtained by multiplying the whole ATE EIF by $\mathbb{1}\{A=a'\}/\pi_{a'}$. The bridge-residual term remains on the event $\{A=a\}$, while the centering term is evaluated on the event $\{A=a'\}$.*

Consider the binary case $a=0$, $a'=1$, and define $\mu_0 := f_{\mathrm{ATT}}(0,1) = \mathbb{E}[Y^{(0)} \mid A=1]$. Cui et al.'s ATT bridge is restricted to the control arm and satisfies

$$\mathbb{E}[h_0^{\mathrm{ATT}}(X,W) \mid A=0, X, Z] = \mathbb{E}[Y \mid A=0, X, Z],$$
$$\mathbb{E}[q_0^{\mathrm{ATT}}(X,Z) \mid A=0, X, W] = \frac{\pi_1}{\pi_0} r_{0,\mathrm{ATT}}(0,1,X,W).$$

Their proof of the ATT EIF first derives the influence function for μ_0, namely

$$\psi^{\mathrm{EIF}}_{\mu_0}(O) = \frac{\mathbb{1}\{A=0\}}{\pi_1} q_0^{\mathrm{ATT}}(X,Z)\{Y - h_0^{\mathrm{ATT}}(X,W)\} + \frac{\mathbb{1}\{A=1\}}{\pi_1}\{h_0^{\mathrm{ATT}}(X,W) - \mu_0\}. \quad (28)$$

Equivalently,

$$\begin{aligned}\mu_0 &= \mathbb{E}\left[\frac{\mathbb{1}\{A=0\}}{\pi_1} q_0^{\mathrm{ATT}}(X,Z)\{Y - h_0^{\mathrm{ATT}}(X,W)\} + \frac{\mathbb{1}\{A=1\}}{\pi_1} h_0^{\mathrm{ATT}}(X,W)\right] \\ &= \mathbb{E}\left[\frac{\mathbb{1}\{A=0\}}{\pi_1} q_0^{\mathrm{ATT}}(X,Z)\{Y - h_0^{\mathrm{ATT}}(X,W)\}\right] + \mathbb{E}[h_0^{\mathrm{ATT}}(X,W) \mid A=1]. \quad (29)\end{aligned}$$

Our normalized ATT bridge satisfies

$$\mathbb{E}[\varphi_0^{\mathrm{ATT}}(0,1,X,Z) \mid A=0, X, W] = r_{0,\mathrm{ATT}}(0,1,X,W).$$

Thus, under uniqueness, $\varphi_0^{\mathrm{ATT}}(0,1,X,Z) = (\pi_0/\pi_1) q_0^{\mathrm{ATT}}(X,Z)$, and Eq.17 reduces to Eq. 28 with $h_0^{\mathrm{ATT}}(X,W) = h_0(0,X,W)$.

The usual binary ATT contrast is

$$\tau_{\mathrm{ATT}} := \mathbb{E}[Y^{(1)} - Y^{(0)} \mid A=1] = \mathbb{E}[Y \mid A=1] - f_{\mathrm{ATT}}(0,1).$$

Its EIF is therefore

$$\frac{\mathbb{1}\{A=1\}}{\pi_1}\{Y - \mathbb{E}[Y \mid A=1]\} - \psi^{\mathrm{EIF}}_{\mathrm{ATT}}(O;0,1).$$

Hence $\psi^{\mathrm{EIF}}_{\mathrm{ATT}}(O;a,a')$ in Theorem D.1 is the EIF for the conditional counterfactual mean $f_{\mathrm{ATT}}(a,a')$; the EIF for an ATT contrast is obtained by adding the observed-arm component and taking the difference.

Remark D.4 (Continuous and structured treatments)**.** *The EIFs in Theorem D.1 are discrete-treatment EIFs. They justify the corresponding discrete doubly robust moment equations. The identification*

formulas in Theorem 3.3 are still stated for continuous or structured A, provided the bridge equations are interpreted through regular conditional laws and the required positivity conditions hold. However, the pointwise semiparametric efficiency theory for continuous treatments is not obtained by formally replacing the discrete atom weight $\mathbb{1}\{A = a\}/\pi_a$ by $r_{0,\mathrm{ATE}}(a, X, W)$. Existing continuous-treatment DR theory typically works with localized or smoothed target functionals (Kennedy et al., 2017; Colangelo & Lee, 2025; Wu et al., 2024; Zenati et al., 2025). Our contribution in the continuous and high-dimensional setting is therefore an identification and estimation construction based on normalized treatment bridges and neural mean embeddings, not a closed-form pointwise continuous-treatment EIF.

Remark D.5 (Heterogeneous dose-response identification)**.** *Equation 5 extends the doubly robust proximal identification formula from population dose-response curves to the heterogeneous response function $f_{\mathrm{CATE}}(a, v) = \mathbb{E}[Y^{(a)} \mid V = v]$. This is distinct from the binary ATT case: the conditioning event is an observed context variable $V = v$, not the realized treatment arm.*

E Neural mean embedding-based proxy causal learning for dose-response

We estimate the dose-response curve using an outcome bridge, a treatment bridge, or their doubly robust combination. Both bridge estimators have the analogous architecture: a first-stage neural conditional mean embedding, a second-stage linear bridge head built on learned features, and a final causal function evaluation stage. This allows the representation itself to adapt to the data and makes the resulting estimators suitable for stochastic optimization in high-dimensional and structured settings in contrast to fixed RKHS features in Bozkurt et al. (2025b).

Our construction consists of three components:

1. **Outcome bridge network.** We begin with the Deep Feature Proxy Causal Learning (DFPCL) architecture of Xu et al. (2021b), which provides a neural parameterization of the outcome bridge estimator. We refine this component by incorporating proximal closed-form updates for the final linear layer (Galashov et al., 2025).
2. **Treatment bridge network.** We next introduce a neural estimator of the treatment bridge function in Section E.2. This construction follows similar principles to the outcome-side network: adaptive feature learning, proximal closed-form updates for the last linear layer.
3. **Neural doubly robust unification.** Finally, in Section E.3, we combine the outcome- and treatment-bridge components into a fully neural doubly robust estimator of the dose-response curve. This final stage is designed to leverage the complementary strengths of both bridge functions.

E.1 Dose-response curve estimation: outcome bridge approach

We first describe the outcome-bridge component, which builds on the DFPCL architecture of Xu et al. (2021b). We refer to this component as OutcomeNet. To distinguish it from the treatment-side construction introduced later in Section E.2, we attach the superscript (h) to all quantities associated with the outcome bridge.

The outcome bridge satisfies

$$\mathbb{E}[Y \mid A = a, X = x, Z = z] = \mathbb{E}[h(a, x, W) \mid A = a, X = x, Z = z].$$

Let $Q_h := (A, X, Z)$ and define

$$(T_h u)(a, x, z) := \mathbb{E}\{u(a, x, W) \mid A = a, X = x, Z = z\},$$
$$m_0(q_h) := \mathbb{E}(Y \mid Q_h = q_h) \in \mathbb{R}^{d_Y}.$$

Consider the population least-squares risk

$$\mathcal{R}_{h,2}(u) := \mathbb{E}\Big[\|Y - (T_h u)(Q_h)\|_2^2\Big]. \tag{30}$$

If h_0 is an outcome bridge, then $T_h h_0 = m_0$, and the conditional least-squares decomposition gives

$$\mathcal{R}_{h,2}(u) - \mathcal{R}_{h,2}(h_0) = \|T_h u - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})} = \|T_h(u - h_0)\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})}.$$

Thus, minimizing $\mathcal{R}_{h,2}$ targets the outcome-bridge equation in projected norm. Its empirical analogue cannot be formed directly from pointwise observations without estimating the unknown conditional mean $(T_h u)(a_i, x_i, z_i)$. Replacing it by the pointwise value $u(a_i, x_i, w_i)$ would instead target an ordinary regression of Y on (A, X, W), not the bridge equation.

Therefore, we parametrize the bridge and the relevant conditional proxy features as follows. Let $Y \in \mathbb{R}^{d_Y}$, so that $h(a, x, w) \in \mathbb{R}^{d_Y}$. If $d_{AX,2}^{(h)}$ and $d_{W,2}^{(h)}$ are the output dimensions of $\phi_{AX,2}^{(h)}$ and $\phi_{W,2}^{(h)}$, respectively, define $d_{h,2} := d_{AX,2}^{(h)} d_{W,2}^{(h)}$ and take $\boldsymbol{H}^{(h)} \in \mathbb{R}^{d_{h,2} \times d_Y}$. Then

$$h(a, x, w) = \left\{ \boldsymbol{H}^{(h)} \right\}^\top \left(\phi_{AX,2}^{(h)}(a, x) \otimes \phi_{W,2}^{(h)}(w) \right). \tag{31}$$

$$\begin{aligned} \mu_W^{(h)}(a, x, z) &= \mathbb{E}\left[\phi_{W,2}^{(h)}(W) \mid A = a, X = x, Z = z \right] \\ \hat{\mu}_W^{(h)}(a, x, z; \boldsymbol{V}^{(h)}) &= \left\{ \boldsymbol{V}^{(h)} \right\}^\top \phi_{AXZ,1}^{(h)}(a, x, z). \end{aligned} \tag{32}$$

For the parameterization in Eq. 31, linearity of conditional expectation gives

$$(T_h h)(a, x, z) = \left\{ \boldsymbol{H}^{(h)} \right\}^\top \left\{ \phi_{AX,2}^{(h)}(a, x) \otimes \mu_W^{(h)}(a, x, z) \right\}.$$

Stage 1 therefore estimates only the conditional feature mean $\mu_W^{(h)}$ through the neural mean embedding in Eq. 32. Stage 2 substitutes this estimate into the sample analogue of Eq. 30. The embedding is an estimation device for the projected bridge, rather than an additional identification assumption.

More specifically:

- $\phi_{AXZ,1}^{(h)} : \mathcal{A} \times \mathcal{X} \times \mathcal{Z} \to \mathbb{R}^{d_{h,1}}$ denotes the first-stage feature extractor, parameterized by $\theta_1^{(h)}$.
- $\phi_{AX,2}^{(h)}$ and $\phi_{W,2}^{(h)}$ denote the second-stage feature extractors for (A, X), and W, jointly parameterized by $\theta_2^{(h)}$.
- $\boldsymbol{V}^{(h)}$ is the first-stage linear operator mapping the learned features of (A, X, Z) to predicted proxy features in the representation space of W.
- $\boldsymbol{H}^{(h)} \in \mathbb{R}^{d_{h,2} \times d_Y}$ is the second-stage linear head which maps the tensorized representation of (A, X, W) to $h(a, x, w) \in \mathbb{R}^{d_Y}$.

This parameterization separates the two roles played by the neural model. The first stage learns how the information contained in (A, X, Z) predicts the proxy representation of W. The second stage then uses this learned proxy representation to fit the outcome bridge. In this way, the architecture preserves the logic of proximal bridge estimation while allowing the representation spaces themselves to be adapted during training.

The first stage learns the neural conditional mean embedding $\hat{\mu}_W^{(h)}(a, x, z)$ by minimizing the regularized least-squares objective

$$\hat{\mathcal{L}}_{h,1}(\theta_1^{(h)}, \boldsymbol{V}^{(h)}) = \frac{1}{|\mathcal{I}_1^{(h)}|} \sum_{i \in \mathcal{I}_1^{(h)}} \left\| \phi_{W,2}^{(h)}(w_i) - \left\{ \boldsymbol{V}^{(h)} \right\}^\top \phi_{AXZ,1}^{(h)}(a_i, x_i, z_i) \right\|_2^2 + \lambda_1^{(h)} \|\boldsymbol{V}^{(h)}\|_F^2, \tag{33}$$

where $\lambda_1^{(h)}$ is the regularization parameter for the first-stage linear layer and $\|\cdot\|_F$ denotes the Frobenius norm. Importantly, although Equation 33 depends on the second-stage proxy featurizer $\phi_{W,2}^{(h)}$, the parameters $\theta_{W,2}^{(h)}$ are not updated in this stage. Instead, $\phi_{W,2}^{(h)}(W)$ is treated as the target representation to be predicted from (A, X, Z).

In the second stage, the learned first-stage operator $\hat{\boldsymbol{V}}^{(h)}$ is used to construct predicted proxy features on $\mathcal{I}_2^{(h)}$. The outcome bridge is then estimated by minimizing

$$\hat{\mathcal{L}}_{h,2}(\theta_2^{(h)}, \boldsymbol{H}^{(h)}) = \frac{1}{|\mathcal{I}_2^{(h)}|} \sum_{i \in \mathcal{I}_2^{(h)}} \left\| y_i - \left\{ \boldsymbol{H}^{(h)} \right\}^\top \boldsymbol{q}_i^{(h)} \right\|_2^2 + \lambda_2^{(h)} \|\boldsymbol{H}^{(h)}\|_F^2, \tag{34}$$

where $\theta_2^{(h)} = \{\theta_{A,2}^{(h)}, \theta_{X,2}^{(h)}, \theta_{W,2}^{(h)}\}$ collects the parameters of the second-stage featurizers, and $\boldsymbol{q}_i^{(h)} \in \mathbb{R}^{d_{h,2}}$ is defined for $i \in \mathcal{I}_2^{(h)}$ by

$$\boldsymbol{q}_i^{(h)} = \phi_{AX,2}^{(h)}(a_i, x_i) \otimes \hat{\boldsymbol{V}}^{(h)}(\theta_{W,2}^{(h)})^\top \phi_{AXZ,1}^{(h)}(a_i, x_i, z_i).$$

A key difficulty is that Equation 34 depends on $\theta_{W,2}^{(h)}$ only implicitly, through the first-stage optimum $\hat{\boldsymbol{V}}^{(h)}(\theta_{W,2}^{(h)})$. Indeed, changing $\theta_{W,2}^{(h)}$ changes the target representation in the first stage, and therefore also changes the first-stage optimum itself. Following Xu et al. (2021b;a), this dependence is handled approximately by differentiating through the closed-form solution for $\hat{\boldsymbol{V}}^{(h)}(\theta_{W,2}^{(h)})$, while treating the first-stage feature representation $\phi_{AXZ,1}^{(h)}$ as fixed. This yields a tractable approximation to the full bi-level gradient.

The resulting learning procedure alternates between gradient-based updates of the neural feature extractors and analytical updates of the final linear layers.

Neural parameter update (stage 1). The first-stage feature extractor is updated by a gradient step on Equation 33 with learning rate $\eta_1^{(h)}$: $\theta_1^{(h)} \leftarrow \theta_1^{(h)} - \eta_1^{(h)} \frac{\partial \hat{\mathcal{L}}_{h,1}}{\partial \theta_1^{(h)}}$.

Closed-form update of the first-stage linear layer. For index set $\mathcal{I}$, define

$$\boldsymbol{F}_h(\mathcal{I}) = \left[\phi_{AXZ,1}^{(h)}(a_i, x_i, z_i) \right]_{i \in \mathcal{I}} \in \mathbb{R}^{d_{h,1} \times n_h}, \quad \boldsymbol{S}_h(\mathcal{I}) = \left[\phi_{W,2}^{(h)}(w_i) \right]_{i \in \mathcal{I}} \in \mathbb{R}^{d_{W,2}^{(h)} \times n_h}.$$

Here $d_{W,2}^{(h)}$ is the output dimension of $\phi_{W,2}^{(h)}$, and $d_{h,1}$ is the output dimension of $\phi_{AXZ,1}^{(h)}$.

Fixing the feature extractors, the first-stage operator $\boldsymbol{V}^{(h)}$ is updated analytically as a minimizer of Eq. 33 by

$$\hat{\boldsymbol{V}}^{(h)}(\theta_{W,2}^{(h)}) = \left(\boldsymbol{F}_h(\mathcal{I}_1^{(h)}) \boldsymbol{F}_h(\mathcal{I}_1^{(h)})^\top + |\mathcal{I}_1^{(h)}| \lambda_1^{(h)} \boldsymbol{I} \right)^{-1} \boldsymbol{F}_h(\mathcal{I}_1^{(h)}) \boldsymbol{S}_h(\mathcal{I}_1^{(h)})^\top, \tag{35}$$

Neural parameter update (stage 2). The second-stage featurizers are updated by a gradient step on Equation 34 with learning rate $\eta_2^{(h)}$: $\theta_2^{(h)} \leftarrow \theta_2^{(h)} - \eta_2^{(h)} \frac{\partial \hat{\mathcal{L}}_{h,2}}{\partial \theta_2^{(h)}}$. For $\theta_{W,2}^{(h)}$, the gradient is propagated through the closed-form expression in Equation 35.

Closed-form update of the second-stage linear layer. For $\mathcal{I}_2^{(h)}$, define

$$\boldsymbol{q}_i^{(h)} = \phi_{AX,2}^{(h)}(a_i, x_i) \otimes \hat{\boldsymbol{V}}^{(h)\top} \phi_{AXZ,1}^{(h)}(a_i, x_i, z_i)$$

Writing

$$\boldsymbol{Q}^{(h)}(\mathcal{I}_2^{(h)}) := \left[\boldsymbol{q}_i^{(h)} \right]_{i \in \mathcal{I}_2^{(h)}} \in \mathbb{R}^{d_{h,2} \times |\mathcal{I}_2^{(h)}|}, \qquad \boldsymbol{Y}(\mathcal{I}_2^{(h)}) := [y_i]_{i \in \mathcal{I}_2^{(h)}} \in \mathbb{R}^{d_Y \times |\mathcal{I}_2^{(h)}|},$$

the matrix head is updated analytically as the minimizer of Equation 34:

$$\hat{\boldsymbol{H}}^{(h)} = \left(\boldsymbol{Q}^{(h)}(\mathcal{I}_2^{(h)}) \boldsymbol{Q}^{(h)}(\mathcal{I}_2^{(h)})^\top + |\mathcal{I}_2^{(h)}| \lambda_2^{(h)} \boldsymbol{I} \right)^{-1} \boldsymbol{Q}^{(h)}(\mathcal{I}_2^{(h)}) \boldsymbol{Y}(\mathcal{I}_2^{(h)})^\top, \tag{36}$$

Following Xu et al. (2021b;a), it is helpful to perform several first-stage updates before each second-stage update. In practice, we carry out $T_{h,1}$ first-stage updates for every $T_{h,2}$ second-stage updates, typically with $T_{h,1} \geq T_{h,2}$, so that the neural conditional mean embedding remains sufficiently accurate throughout training.

Proximal closed-form updates. The analytical updates in Equations 35 and 36 are exact minimizers of the corresponding quadratic objectives on a fixed dataset. In stochastic mini-batch training, however, directly replacing the last-layer parameters by these batchwise optima can lead to unstable trajectories, since each batch induces a different local optimum. Following Galashov et al. (2025), we therefore replace the zero-centered ridge penalties by proximal penalties centered at the previous iterates. This keeps the batchwise closed-form updates close to the current parameter state and substantially stabilizes training.

In the present two-stage setting, there is one additional subtlety. The stage-2 loss depends on the first-stage optimum through the predicted proxy features. Consequently, in addition to the persistent first-stage iterate $\hat{\boldsymbol{V}}_t^{(h)}$, we introduce an auxiliary first-stage solution that is recomputed *on the fly* on the current second-stage batch. This auxiliary operator is not stored as a separate model parameter. Its only role is to evaluate the second-stage loss with the current proxy representation $\phi_{W,2}^{(h)}$. To define it, note that $\mathcal{I}_2^{(h)}$ must also retain the proxy observations w_i, even though they do not appear explicitly in the quadratic bridge loss.

Let $\mathcal{B}_{1,t}^{(h)} \subset \mathcal{I}_1^{(h)}$ and $\mathcal{B}_{2,t}^{(h)} \subset \mathcal{I}_2^{(h)}$ denote mini-batches sampled at iteration t. We define the first-stage proximal batch loss

$$\hat{\mathcal{L}}_{h,1,t}^{\mathrm{prox}}\Big(\theta_1^{(h)}, \boldsymbol{V}^{(h)}; \theta_2^{(h)}, \mathcal{B}_{1,t}^{(h)}\Big) = \frac{1}{|\mathcal{B}_{1,t}^{(h)}|} \sum_{i \in \mathcal{B}_{1,t}^{(h)}} \left\| \phi_{W,2}^{(h)}(w_i) - \Big(\boldsymbol{V}^{(h)}\Big)^\top \phi_{AXZ,1}^{(h)}(a_i, x_i, z_i) \right\|_2^2 + \lambda_1^{(h)} \left\| \boldsymbol{V}^{(h)} - \hat{\boldsymbol{V}}_t^{(h)} \right\|_F^2, \tag{37}$$

and, suppressing fixed neural parameters from its arguments, write the second-stage proximal loss as

$$\hat{\mathcal{L}}_{h,2,t}^{\mathrm{prox}}\Big(\theta_2^{(h)}, \boldsymbol{H}^{(h)}\Big) = \frac{1}{|\mathcal{B}_{2,t}^{(h)}|} \sum_{i \in \mathcal{B}_{2,t}^{(h)}} \left\| y_i - \Big\{\boldsymbol{H}^{(h)}\Big\}^\top \left(\phi_{AX,2}^{(h)}(a_i, x_i) \otimes \hat{\mu}_W^{(h)}(a_i, x_i, z_i; \check{\boldsymbol{V}}_t^{(h)})\right) \right\|_2^2 + \lambda_2^{(h)} \left\| \boldsymbol{H}^{(h)} - \hat{\boldsymbol{H}}_t^{(h)} \right\|_F^2, \tag{38}$$

where $\lambda_1^{(h)}$ and $\lambda_2^{(h)}$ are proximal regularization parameters, $\hat{\boldsymbol{V}}_t^{(h)}$ and $\hat{\boldsymbol{H}}_t^{(h)}$ are the previous iterates, and $\check{\boldsymbol{V}}_t^{(h)}$ denotes the auxiliary on-the-fly first-stage solution computed from the current second-stage batch.

More precisely, we define the persistent first-stage update by

$$\hat{\boldsymbol{V}}_{t+1}^{(h)}\Big(\theta_1^{(h)}; \theta_2^{(h)}, \mathcal{B}_{1,t}^{(h)}\Big) = \arg\min_{\boldsymbol{V}^{(h)}} \hat{\mathcal{L}}_{h,1,t}^{\mathrm{prox}}\Big(\theta_1^{(h)}, \boldsymbol{V}^{(h)}; \theta_2^{(h)}, \mathcal{B}_{1,t}^{(h)}\Big), \tag{39}$$

and the auxiliary on-the-fly first-stage operator by

$$\check{\boldsymbol{V}}_t^{(h)}\Big(\theta_1^{(h)}; \theta_2^{(h)}, \mathcal{B}_{2,t}^{(h)}\Big) = \arg\min_{\boldsymbol{V}^{(h)}} \hat{\mathcal{L}}_{h,1,t}^{\mathrm{prox}}\Big(\theta_1^{(h)}, \boldsymbol{V}^{(h)}; \theta_2^{(h)}, \mathcal{B}_{2,t}^{(h)}\Big). \tag{40}$$

The first quantity is the actual stage-1 parameter update, while the second is an auxiliary batchwise solve used only inside the stage-2 objective.

Given $\check{\boldsymbol{V}}_t^{(h)}$, the second-stage proximal update is

$$\hat{\boldsymbol{H}}_{t+1}^{(h)}\Big(\theta_1^{(h)}; \theta_2^{(h)}, \mathcal{B}_2^{(h)}\Big) = \arg\min_{\boldsymbol{H}^{(h)}} \hat{\mathcal{L}}_{h,2,t}^{\mathrm{prox}}\Big(\boldsymbol{H}^{(h)}\Big). \tag{41}$$

Expanding these minimizers yields the proximal closed-form updates. For any batch $\mathcal{B}$, we have

$$\boldsymbol{F}_h(\mathcal{B}) = \Big[\phi_{AXZ,1}^{(h)}(a_i, x_i, z_i)\Big]_{i \in \mathcal{B}}, \qquad \boldsymbol{S}_h(\mathcal{B}) = \Big[\phi_{W,2}^{(h)}(w_i)\Big]_{i \in \mathcal{B}},$$

where the tuples (a_i, x_i, z_i, w_i) are read from the corresponding batch. Then

$$\hat{\boldsymbol{V}}_{t+1}^{(h)}\Big(\theta_1^{(h)}; \theta_2^{(h)}, \mathcal{B}_{1,t}^{(h)}\Big) = \Big(\boldsymbol{F}_h(\mathcal{B}_{1,t}^{(h)}) \boldsymbol{F}_h(\mathcal{B}_{1,t}^{(h)})^\top + |\mathcal{B}_{1,t}^{(h)}| \lambda_1^{(h)} \boldsymbol{I}\Big)^{-1}.$$

$$\left(\boldsymbol{F}_h(\mathcal{B}_{1,t}^{(h)})\boldsymbol{S}_h(\mathcal{B}_{1,t}^{(h)})^\top + |\mathcal{B}_{1,t}^{(h)}|\lambda_1^{(h)}\hat{\boldsymbol{V}}_t^{(h)}\right), \tag{42}$$

$$\check{\boldsymbol{V}}_t^{(h)}\left(\theta_1^{(h)};\theta_2^{(h)},\mathcal{B}_2^{(h)}\right) = \left(\boldsymbol{F}_h(\mathcal{B}_{2,t}^{(h)})\boldsymbol{F}_h(\mathcal{B}_{2,t}^{(h)})^\top + |\mathcal{B}_{2,t}^{(h)}|\lambda_1^{(h)}\boldsymbol{I}\right)^{-1}\cdot \left(\boldsymbol{F}_h(\mathcal{B}_{2,t}^{(h)})\boldsymbol{S}_h(\mathcal{B}_{2,t}^{(h)})^\top + |\mathcal{B}_{2,t}^{(h)}|\lambda_1^{(h)}\hat{\boldsymbol{V}}_t^{(h)}\right), \tag{43}$$

The corresponding second-stage features on $\mathcal{B}_{2,t}^{(h)}$ are

$$\boldsymbol{q}_i^{(h)} := \phi_{AX,2}^{(h)}(a_i,x_i)\otimes\left\{\check{\boldsymbol{V}}_t^{(h)}\right\}^\top \phi_{AXZ,1}^{(h)}(a_i,x_i,z_i), \qquad i\in\mathcal{B}_{2,t}^{(h)}.$$

We collect them columnwise as

$$\boldsymbol{Q}^{(h)}(\mathcal{B}_{2,t}^{(h)}) := \left[\boldsymbol{q}_i^{(h)}\right]_{i\in\mathcal{B}_{2,t}^{(h)}} \in \mathbb{R}^{d_{h,2}\times|\mathcal{B}_{2,t}^{(h)}|}, \qquad \boldsymbol{Y}(\mathcal{B}_{2,t}^{(h)}) := [y_i]_{i\in\mathcal{B}_{2,t}^{(h)}} \in \mathbb{R}^{d_Y\times|\mathcal{B}_{2,t}^{(h)}|}.$$

Then, the second-stage proximal update is

$$\begin{aligned}\hat{\boldsymbol{H}}_{t+1}^{(h)} = &\left(\boldsymbol{Q}^{(h)}(\mathcal{B}_{2,t}^{(h)})\boldsymbol{Q}^{(h)}(\mathcal{B}_{2,t}^{(h)})^\top + |\mathcal{B}_{2,t}^{(h)}|\lambda_2^{(h)}\boldsymbol{I}\right)^{-1}\\ &\times\left(\boldsymbol{Q}^{(h)}(\mathcal{B}_{2,t}^{(h)})\boldsymbol{Y}(\mathcal{B}_{2,t}^{(h)})^\top + |\mathcal{B}_{2,t}^{(h)}|\lambda_2^{(h)}\hat{\boldsymbol{H}}_t^{(h)}\right).\end{aligned} \tag{44}$$

Equivalently, the batch-size factors may be absorbed into $\lambda_1^{(h)}$ and $\lambda_2^{(h)}$; we write them explicitly here only to match the averaged losses in Equations 37 and 38.

Finally, note that the stage-2 loss depends on $\theta_{W,2}^{(h)}$ through the auxiliary operator $\check{\boldsymbol{V}}_t^{(h)}$. In practice, gradients with respect to $\theta_{W,2}^{(h)}$ are therefore propagated through the closed-form mapping in Equation 43, while the first-stage representation $\phi_{AXZ,1}^{(h)}$ is treated as fixed, exactly as in the DFPCL update (Xu et al., 2021b).

Dose-response estimation. The third stage of OUTCOMENET is a sample average of the outcome bridge function. In particular, on an independent evaluation fold $\mathcal{I}_3^{(h)}$, OUTCOMENET estimates the dose-response curve by

$$\hat{f}_{\mathrm{ATE}}^{(h)}(a) = \frac{1}{|\mathcal{I}_3^{(h)}|}\sum_{i\in\mathcal{I}_3^{(h)}}\hat{h}(a,x_i,w_i)$$

For $\phi_{AX,2}^{(h)}(\cdot,\cdot) = \phi_{A,2}^{(h)}(\cdot)\otimes\phi_{X,2}^{(h)}(\cdot)$, define

$$\widehat{\mu}_{XW}^{(h)} := \frac{1}{|\mathcal{I}_3^{(h)}|}\sum_{i\in\mathcal{I}_3^{(h)}}\phi_{X,2}^{(h)}(x_i)\otimes\phi_{W,2}^{(h)}(w_i).$$

The estimator can then be written equivalently as $\hat{f}_{\mathrm{ATE}}^{(h)}(a) = \left\{\hat{\boldsymbol{H}}^{(h)}\right\}^\top\left[\phi_{A,2}^{(h)}(a)\otimes\widehat{\mu}_{XW}^{(h)}\right]$.

This empirical average is valid for continuous and discrete treatment because the learned bridge can be evaluated at any inquiry treatment level a; it is not a regression on the observed treatment.

The complete OutcomeNet training procedure is summarized in Algorithm 1.

Algorithm 1 Population dose-response estimation with OUTCOMENET.

Input: Observations $\mathcal{D} = \{o_i\}_{i=1}^n$, with $o_i = (a_i, x_i, w_i, z_i, y_i)$, and disjoint folds $\mathcal{I}_1^{(h)}, \mathcal{I}_2^{(h)}, \mathcal{I}_3^{(h)}$.
Parameters: Neural parameters $\theta_1^{(h)}, \theta_2^{(h)}$; linear layers $\boldsymbol{V}^{(h)}, \boldsymbol{H}^{(h)}$.
Hyperparameters: Learning rates $\eta_1^{(h)}, \eta_2^{(h)}$, regularization coefficients $\lambda_1^{(h)}, \lambda_2^{(h)}$, and update counts $T_{h,1}, T_{h,2}$.
Output: Fitted outcome bridge $\hat{h}$ and population dose-response estimate $\hat{f}_{\mathrm{ATE}}^{(h)}$.

repeat
  Sample mini-batches $\mathcal{B}_1^{(h)} \subset \mathcal{I}_1^{(h)}$ and $\mathcal{B}_2^{(h)} \subset \mathcal{I}_2^{(h)}$.
  for $s = 1, \ldots, T_{h,1}$ **do**
    Update $\theta_1^{(h)}$ by one gradient step on Equation 37, holding the proxy-target featurizer fixed.
    Update the persistent operator $\hat{\boldsymbol{V}}^{(h)}$ by Equation 42 on $\mathcal{B}_1^{(h)}$.
  end for
  for $s = 1, \ldots, T_{h,2}$ **do**
    Compute $\check{\boldsymbol{V}}^{(h)}$ on $\mathcal{B}_2^{(h)}$ using Equation 43.
    Update $\theta_2^{(h)}$ by one gradient step on Equation 38; differentiate through $\check{\boldsymbol{V}}^{(h)}$ for the proxy-side features.
    Update $\hat{\boldsymbol{H}}^{(h)}$ by Equation 44.
  end for
until convergence
Set $\hat{h}(a, x, w) = \{\hat{\boldsymbol{H}}^{(h)}\}^\top [\phi_{AX,2}^{(h)}(a, x) \otimes \phi_{W,2}^{(h)}(w)]$.
Set $\hat{f}_{\mathrm{ATE}}^{(h)}(a) = |\mathcal{I}_3^{(h)}|^{-1} \sum_{i \in \mathcal{I}_3^{(h)}} \hat{h}(a, x_i, w_i)$.
Return $\hat{h}$ and $\hat{f}_{\mathrm{ATE}}^{(h)}$.

E.2 DOSE-RESPONSE CURVE ESTIMATION: TREATMENT BRIDGE APPROACH

We now introduce the treatment bridge-based component of our neural mean embedding framework, which we refer to as TREATMENTNET. The target of this component is the treatment bridge $\varphi_0^{\mathrm{ATE}}(a, x, z)$, which satisfies the conditional moment relation

$$\begin{aligned} g_{\varphi,0}(a, x, w) &\coloneqq \mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a, X, W] \\ &= r_{0,\mathrm{ATE}}(a, X, W), \quad r_{0,\mathrm{ATE}}(a, x, w) = \frac{d\left(\mathbb{P}_A \otimes \mathbb{P}_{X,W}\right)}{d\mathbb{P}_{A,X,W}}. \end{aligned}$$

In contrast to the OUTCOMENET part of this method, we do not have access to the target of the treatment bridge method, density ratios, and the mainstream approach is to plug-in estimated density ratios into the objective function to learn the treatment bridge function. Recent work (Bozkurt et al., 2025a;b) provides a density-ratio-free objective to learn the treatment bridge function, and we follow this as characterized in Lemma 4.1 which is proved below.

Proof of Lemma 4.1. The ratio-free risk is defined in Eq. 8 as

$$\mathcal{R}_{\mathrm{ATE}}^{\mathrm{RF}}(u) \coloneqq \mathbb{E}\{u(A_1, X_1, W_1)^2\} - 2\mathbb{E}\{u(A_1, X_2, W_2)\}.$$

Let $\mathbb{P}_Q = \mathbb{P}_{A,X,W}$ and $\mathbb{P}_\otimes = \mathbb{P}_A \otimes \mathbb{P}_{X,W}$. Then, by change of measure

$$\begin{aligned} \mathcal{R}_{\mathrm{ATE}}^{\mathrm{RF}}(u) &= \int u^2 d\mathbb{P}_Q - 2\int u d\mathbb{P}_\otimes = \int \left(u^2 - 2u r_{0,\mathrm{ATE}}\right) d\mathbb{P}_Q \\ &= \|u - r_{0,\mathrm{ATE}}\|_{L^2(\mathbb{P}_Q)}^2 - \|r_{0,\mathrm{ATE}}\|_{L^2(\mathbb{P}_Q)}^2 \end{aligned}$$

Substituting $u = r_{0,\mathrm{ATE}}$ and subtracting proves the claim. □

The treatment bridge and the first-stage neural conditional mean embedding are parameterized as

$$\begin{aligned} \varphi(a, x, z) &= \boldsymbol{\varphi}^\top \left(\phi_{AX,2}^{(\varphi)}(a, x) \otimes \phi_{Z,2}^{(\varphi)}(z)\right), \\ \mu_Z^{(\varphi)}(a, x, w) &= \mathbb{E}[\phi_{Z,2}^{(\varphi)}(Z) \mid A = a, X = x, W = w] \\ \hat{\mu}_Z^{(\varphi)}(a, x, w; \boldsymbol{V}^{(\varphi)}) &= \left(\boldsymbol{V}^{(\varphi)}\right)^\top \phi_{AXW,1}^{(\varphi)}(a, x, w). \end{aligned}$$

Here:

- $\phi^{(\varphi)}_{AXW,1} : \mathcal{A} \times \mathcal{X} \times \mathcal{W} \to \mathbb{R}^{d_{\varphi,1}}$ is the first-stage feature extractor, parameterized by $\theta^{(\varphi)}_1$.
- $\phi^{(\varphi)}_{AX,2}$ and $\phi^{(\varphi)}_{Z,2}$ are the second-stage feature extractors, jointly parameterized by $\theta^{(\varphi)}_2$.
- $\boldsymbol{V}^{(\varphi)}$ is the first-stage linear operator mapping features of (A, X, W) to predicted proxy features.
- $\boldsymbol{\varphi}$ is the second-stage linear head mapping the tensorized representation of (A, X, Z) to the bridge value $\varphi(a, x, z)$.

- Consequently, the fitted projected bridge function is

$$\hat{g}_\varphi(a, x, w) = \boldsymbol{\varphi}^\top \left[\phi^{(\varphi)}_{AX,2}(a, x) \otimes \hat{\mu}^{(\varphi)}_Z(a, x, w; \boldsymbol{V}^{(\varphi)})\right]$$

As on the outcome side, TREATMENTNET is trained by a bi-level multi-stage procedure. For fixed current second-stage proxy features, define the first stage sample risk on first stage batch $\mathcal{B}_{1,t}$:

$$\hat{\mathcal{R}}_{\varphi,1}(\theta^{(\varphi)}_1, \boldsymbol{V}^{(\varphi)}; \theta^{\varphi}_2, \mathcal{B}_{1,t}) = \frac{1}{|\mathcal{B}^{(\varphi)}_{1,t}|} \sum_{i \in \mathcal{B}^{(\varphi)}_{1,t}} \left\| \phi^{(\varphi)}_{Z,2}(z_i) - \left(\boldsymbol{V}^{(\varphi)}\right)^\top \phi^{(\varphi)}_{AXW,1}(a_i, x_i, w_i) \right\|^2_2 \tag{45}$$

$$\hat{\mathcal{L}}_{\varphi,1,t}(\theta^{(\varphi)}_1, \boldsymbol{V}^{(\varphi)}; \theta^{\varphi}_2, \mathcal{B}_1) = \hat{\mathcal{R}}_{\varphi,1}(\theta^{(\varphi)}_1, \boldsymbol{V}^{(\varphi)}; \theta^{\varphi}_2, \mathcal{B}_1) + \lambda^{(\varphi)}_1 \left\| \boldsymbol{V}^{(\varphi)} - \hat{\boldsymbol{V}}^{(\varphi)}_t \right\|^2_F. \tag{46}$$

the targets $\phi^{(\varphi)}_{Z,2}(Z)$ depend on the second-stage proxy featurizer, but $\theta^{(\varphi)}_{Z,2}$ is not updated in this stage; the feature map acts only as the target representation.

On batch $\mathcal{B}$, define

$$F_\varphi(\mathcal{B}) = \left[\phi^{(\varphi)}_{AXW,1}(a_i, x_i, w_i)\right]_{i \in \mathcal{B}}, \quad S_\varphi(\mathcal{B}) = \left[\phi^{(\varphi)}_{Z,2}(z_i)\right]_{i \in \mathcal{B}}.$$

Then, the proximal first-stage minimizer is

$$\begin{aligned}\hat{\boldsymbol{V}}^{(\varphi)}_{t+1}\left(\theta^{(\varphi)}_1; \theta^{(\varphi)}_2, \mathcal{B}^{(\varphi)}_{1,t}\right) = &\left(F_\varphi(\mathcal{B}^{(\varphi)}_{1,t}) F_\varphi(\mathcal{B}^{(\varphi)}_{1,t})^\top + |\mathcal{B}^{(\varphi)}_{1,t}| \lambda^{(\varphi)}_1 \boldsymbol{I}\right)^{-1} \cdot \\ &\left(F_\varphi(\mathcal{B}^{(\varphi)}_{1,t}) S_\varphi(\mathcal{B}^{(\varphi)}_{1,t})^\top + |\mathcal{B}^{(\varphi)}_{1,t}| \lambda^{(\varphi)}_1 \hat{\boldsymbol{V}}^{(\varphi)}_t\right).\end{aligned} \tag{47}$$

The second stage is trained on a mini-batch $\mathcal{B}^{(\varphi)}_{2,t} \subset \mathcal{I}^{(\varphi)}_2$. As on the outcome side, the second-stage prediction depends on the current first-stage optimum. We therefore distinguish between:

- the *persistent* first-stage operator $\hat{\boldsymbol{V}}^{(\varphi)}_t$, which is updated on batches from $\mathcal{I}^{(\varphi)}_1$, and
- an *auxiliary on-the-fly* operator $\check{\boldsymbol{V}}^{(\varphi)}_t$, recomputed on the current second-stage batch $\mathcal{B}^{(\varphi)}_{2,t}$.

The auxiliary operator is not stored as a separate model parameter. Its only role is to evaluate the second-stage loss using the current proxy feature representation $\phi^{(\varphi)}_{Z,2}$, exactly as in the outcome-bridge formulation.

Formally, $\check{\boldsymbol{V}}^{(\varphi)}_t$ is defined as the minimizer of the first-stage proximal objective evaluated on $\mathcal{B}^{(\varphi)}_2$:

$$\begin{aligned}\check{\boldsymbol{V}}^{(\varphi)}_t\left(\theta^{(\varphi)}_1; \theta^{(\varphi)}_2, \mathcal{B}^{(\varphi)}_2\right) = &\left(F_\varphi(\mathcal{B}^{(\varphi)}_{2,t}) F_\varphi(\mathcal{B}^{(\varphi)}_{2,t})^\top + |\mathcal{B}^{(\varphi)}_{2,t}| \lambda^{(\varphi)}_1 \boldsymbol{I}\right)^{-1} \cdot \\ &\left(F_\varphi(\mathcal{B}^{(\varphi)}_{2,t}) S_\varphi(\mathcal{B}^{(\varphi)}_{2,t})^\top + |\mathcal{B}^{(\varphi)}_{2,t}| \lambda^{(\varphi)}_1 \hat{\boldsymbol{V}}^{(\varphi)}_t\right).\end{aligned} \tag{48}$$

Let $\mathcal{B}^{(\varphi)}_{2,t}$ be the second-stage batch sampled at iteration t. For every ordered pair (i, j) on $\mathcal{B}^{(\varphi)}_{2,t}$, define

$$\boldsymbol{u}^{(\varphi)}_{ij,t} = \phi^{(\varphi)}_{AX,2}(a_i, x_j) \otimes \left\{\check{\boldsymbol{V}}^{(\varphi)}_t\right\}^\top \phi^{(\varphi)}_{AXW,1}(a_i, x_j, w_j). \tag{49}$$

Then, the plug-in unbiased ratio-free loss at iteration t is

$$\hat{\mathcal{L}}^{\mathrm{RF,U}}_{\varphi,2,t}(\varphi) = \hat{\mathcal{R}}^{\mathrm{RF,U}}_{\varphi,2,t}(\varphi) + \lambda_2^{(\varphi)}\|\boldsymbol{\varphi} - \hat{\boldsymbol{\varphi}}_t\|_2^2 \tag{50}$$

$$\hat{\mathcal{R}}^{\mathrm{RF,U}}_{\varphi,2,t}(\varphi) = \frac{1}{|\mathcal{B}^{(\varphi)}_{2,t}|}\sum_{i\in\mathcal{B}^{(\varphi)}_{2,t}}\left(\boldsymbol{\varphi}^\top \boldsymbol{u}^{(\varphi)}_{ii,t}\right)^2 - \frac{2}{|\mathcal{B}^{(\varphi)}_{2,t}|(|\mathcal{B}^{(\varphi)}_{2,t}|-1)}\sum_{\substack{i,j\in\mathcal{B}^{(\varphi)}_{2,t}\\ i\neq j}}\boldsymbol{\varphi}^\top \boldsymbol{u}^{(\varphi)}_{ij,t}. \tag{51}$$

A difficulty, analogous to the one encountered in OutcomeNet, is that Equation 50 depends on $\theta^{(\varphi)}_{Z,2}$ only implicitly, through the first-stage optimum $\check{\boldsymbol{V}}^{(\varphi)}_t$. Following Xu et al. (2021b;a), we handle this dependence by differentiating through the closed-form mapping in Equation 48, while treating the first-stage feature matrix $\Phi^{(\varphi)}_{AXW,1}$ as fixed when computing gradients with respect to $\theta^{(\varphi)}_{Z,2}$.

Conditional on the separately fitted first-stage representations, the off-diagonal term is the complete U-statistic for the product-law expectation. For the head update, define

$$\boldsymbol{G}^{(\varphi)}_t = \frac{1}{|\mathcal{B}^{(\varphi)}_{2,t}|}\sum_{i\in\mathcal{B}^{(\varphi)}_{2,t}}\boldsymbol{u}^{(\varphi)}_{ii,t}\boldsymbol{u}^{(\varphi)\top}_{ii,t}, \quad \boldsymbol{b}^{(\varphi)}_{U,t} = \frac{1}{|\mathcal{B}^{(\varphi)}_{2,t}|(|\mathcal{B}^{(\varphi)}_{2,t}|-1)}\sum_{\substack{i,j\in\mathcal{B}^{(\varphi)}_{2,t}\\ i\neq j}}\boldsymbol{u}^{(\varphi)}_{ij,t}.$$

Then, the proximal second-stage minimizer is

$$\hat{\boldsymbol{\varphi}}_{t+1} = \left(\boldsymbol{G}^{(\varphi)}_t + \lambda_2^{(\varphi)}\boldsymbol{I}\right)^{-1}\left(\boldsymbol{b}^{(\varphi)}_{U,t} + \lambda_2^{(\varphi)}\hat{\boldsymbol{\varphi}}_t\right) \tag{52}$$

The featurizer $\phi^{(\varphi)}_{AX,2}$ is updated directly through Eq. 50.

Dose-response curve estimation with the treatment bridge. The preceding two-stage procedure yields an estimator $\hat{\boldsymbol{\varphi}}$ of the treatment bridge satisfying the identifying relation in Eq.2. To recover the final dose-response curve, we must still estimate the conditional mean $\mathbb{E}[Y\varphi_0^{\mathrm{ATE}}(a,X,Z)\mid A=a]$. We distinguish two cases of continuous treatment and discrete treatment domains.

- *Continuous treatment domain*:
  Unlike the outcome-bridge formulation, the target is itself a regression function of a. On the independent fold $\mathcal{I}^{(\varphi)}_3$, define

  $$\mathcal{D}^{(\varphi)}_3 := \{(a_i, y^{(\varphi)}_i)\}_{i\in\mathcal{I}^{(\varphi)}_3}, \qquad y^{(\varphi)}_i := y_i\hat{\varphi}(a_i,x_i,z_i)\in\mathbb{R}^{d_Y}.$$

  We fit the network $k^\varphi(\cdot;\theta^\varphi_3)$ by minimizing the empirical mean-squared error

  $$\widehat{\mathcal{L}}_{\varphi,3}(\theta^\varphi_3) := \frac{1}{|\mathcal{I}^{(\varphi)}_3|}\sum_{i\in\mathcal{I}^{(\varphi)}_3}\left\|y^{(\varphi)}_i - k^\varphi(a_i;\theta^\varphi_3)\right\|_2^2. \tag{53}$$

  Let $\hat{\theta}^\varphi_3 \in \arg\min_{\theta^\varphi_3}\widehat{\mathcal{L}}_{\varphi,3}(\theta^\varphi_3)$. Since the population minimizer of squared error is the conditional mean,

  $$k^\varphi(a;\hat{\theta}^\varphi_3) \approx \mathbb{E}\{Y\hat{\varphi}(a,X,Z)\mid A=a\}, \qquad \hat{f}^{(\varphi)}_{\mathrm{ATE}}(a) := k^\varphi(a;\hat{\theta}^\varphi_3).$$

- *Discrete treatment domain*:
  When $\mathcal{A}$ is finite, the MSE regression reduces to a within-treatment average. Writing $t_{\varphi,a} := \sum_{i\in\mathcal{I}^{(\varphi)}_3}\mathbf{1}\{a_i=a\}$,

  $$\hat{f}^{(\varphi)}_{\mathrm{ATE}}(a) = \frac{1}{t_{\varphi,a}}\sum_{\substack{i\in\mathcal{I}^{(\varphi)}_3\\ a_i=a}} y_i\hat{\varphi}(a,x_i,z_i), \qquad t_{\varphi,a}>0.$$

The complete training procedure for the treatment bridge function is summarized in Algorithm 2.

Algorithm 2 Population dose-response estimation with density-ratio-free TREATMENTNET. No density ratio is estimated or supplied as a training target.

Input: Observations $\mathcal{D} = \{o_i\}_{i=1}^n$ and disjoint folds $\mathcal{I}_1^{(\varphi)}, \mathcal{I}_2^{(\varphi)}, \mathcal{I}_3^{(\varphi)}$.
Parameters: Neural parameters $\theta_1^{(\varphi)}, \theta_2^{(\varphi)}, \theta_3^{\varphi}$; linear layers $\boldsymbol{V}^{(\varphi)}, \boldsymbol{\varphi}$.
Hyperparameters: Learning rates $\eta_1^{(\varphi)}, \eta_2^{(\varphi)}$, regularization coefficients $\lambda_1^{(\varphi)}, \lambda_2^{(\varphi)}$, and update counts $T_{\varphi,1}, T_{\varphi,2}$.
Output: Fitted treatment bridge $\hat{\varphi}^{\mathrm{ATE}}$ and population dose-response estimate $\hat{f}_{\mathrm{ATE}}^{(\varphi)}$.

repeat
  Sample mini-batches $\mathcal{B}_1^{(\varphi)} \subset \mathcal{I}_1^{(\varphi)}$ and $\mathcal{B}_2^{(\varphi)} \subset \mathcal{I}_2^{(\varphi)}$.
  for $s = 1, \ldots, T_{\varphi,1}$ **do**
    Update $\theta_1^{(\varphi)}$ by one gradient step on Equation 46, holding the proxy-target featurizer fixed.
    Update the persistent operator $\hat{\boldsymbol{V}}^{(\varphi)}$ by Equation 47.
  end for
  for $s = 1, \ldots, T_{\varphi,2}$ **do**
    Compute $\check{\boldsymbol{V}}^{(\varphi)}$ on $\mathcal{B}_2^{(\varphi)}$ using Equation 48.
    Update $\theta_2^{(\varphi)}$ by one gradient step on the fixed ratio-free U-objective in Equation 50, differentiating through $\check{\boldsymbol{V}}^{(\varphi)}$ for the proxy-side features.
    Recompute $\check{\boldsymbol{V}}^{(\varphi)}$, the projected features, $\boldsymbol{G}^{(\varphi)}$, and $\boldsymbol{b}_U^{(\varphi)}$ after the feature update.
    Update $\hat{\boldsymbol{\varphi}}$ by the proximal closed-form solve in Equation 52.
  end for
until convergence
Set $\hat{\varphi}^{\mathrm{ATE}}(a, x, z) = \hat{\boldsymbol{\varphi}}^\top [\phi_{AX,2}^{(\varphi)}(a, x) \otimes \phi_{Z,2}^{(\varphi)}(z)]$.
if A is continuous **then**
  For $i \in \mathcal{I}_3^{(\varphi)}$, set $y_i^{(\varphi)} = y_i \hat{\varphi}^{\mathrm{ATE}}(a_i, x_i, z_i)$, and fit k^{φ} by the MSE in Equation 53.
  Set $\hat{f}_{\mathrm{ATE}}^{(\varphi)}(a) = k^{\varphi}(a; \hat{\theta}_3^{\varphi})$.
else
  Set $\hat{f}_{\mathrm{ATE}}^{(\varphi)}(a)$ to the within-a average of $y_i \hat{\varphi}^{\mathrm{ATE}}(a, x_i, z_i)$ over $\mathcal{I}_3^{(\varphi)}$.
end if
Return $\hat{\varphi}^{\mathrm{ATE}}$ and $\hat{f}_{\mathrm{ATE}}^{(\varphi)}$.

E.3 DOSE-RESPONSE CURVE ESTIMATION: DOUBLY ROBUST APPROACH

We now combine the learned outcome and treatment bridges into a doubly robust estimator of the dose-response curve. Recall the identifying representation

$$\Psi_{\mathrm{ATE}}(a; h_0, \varphi_0^{\mathrm{ATE}}) = \mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, X, Z)\{Y - h_0(a, X, W)\} \mid A = a] + \mathbb{E}[h_0(a, X, W)].$$

Given estimators $\hat{h}$ from OUTCOMENET and $\hat{\varphi}$ from TREATMENTNET, the only remaining task is to estimate the conditional expectation terms appearing in this formula. Let $\mathcal{I}_3^{(\mathrm{DR})}$ be an evaluation fold independent of the bridge fits.

- *Continuous treatment domain:*
  Define the residual pseudo-outcome data

$$\mathcal{D}_3^{(\mathrm{DR})} := \{(a_i, y_i^{(\mathrm{DR})})\}_{i \in \mathcal{I}_3^{(\mathrm{DR})}},$$
$$y_i^{(\mathrm{DR})} := \hat{\varphi}(a_i, x_i, z_i)\{y_i - \hat{h}(a_i, x_i, w_i)\} \in \mathbb{R}^{d_Y}.$$

  We fit $k^{\mathrm{DR}}(\cdot; \theta_3^{\mathrm{DR}})$ by minimizing the empirical mean-squared error

$$\widehat{\mathcal{L}}_{\mathrm{DR},3}(\theta_3^{\mathrm{DR}}) := \frac{1}{|\mathcal{I}_3^{(\mathrm{DR})}|} \sum_{i \in \mathcal{I}_3^{(\mathrm{DR})}} \left\| y_i^{(\mathrm{DR})} - k^{\mathrm{DR}}(a_i; \theta_3^{\mathrm{DR}}) \right\|_2^2. \tag{54}$$

  Let $\hat{\theta}_3^{\mathrm{DR}} \in \arg\min_{\theta_3^{\mathrm{DR}}} \widehat{\mathcal{L}}_{\mathrm{DR},3}(\theta_3^{\mathrm{DR}})$. The population minimizer is

$$\mathbb{E}[\hat{\varphi}(a, X, Z)\{Y - \hat{h}(a, X, W)\} \mid A = a].$$

Combining this fitted correction with the outcome-bridge average, computed on the same evaluation fold, gives

$$\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) := \frac{1}{|\mathcal{I}_3^{(\mathrm{DR})}|} \sum_{i \in \mathcal{I}_3^{(\mathrm{DR})}} \hat{h}(a, x_i, w_i) + k^{\mathrm{DR}}(a; \hat{\theta}_3^{\mathrm{DR}}). \tag{55}$$

- *Discrete treatment domain*:

  When A is finite, the correction regression reduces to a saturated within-treatment average. Let $t_{\mathrm{DR},a} := \sum_{i \in \mathcal{I}_3^{(\mathrm{DR})}} \mathbf{1}\{a_i = a\}$. Then, for $t_{\mathrm{DR},a} > 0$,

$$\begin{aligned}\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) = &\frac{1}{|\mathcal{I}_3^{(\mathrm{DR})}|} \sum_{i \in \mathcal{I}_3^{(\mathrm{DR})}} \hat{h}(a, x_i, w_i) \\ &+ \frac{1}{t_{\mathrm{DR},a}} \sum_{\substack{i \in \mathcal{I}_3^{(\mathrm{DR})} \\ a_i = a}} \hat{\varphi}(a, x_i, z_i)\{y_i - \hat{h}(a, x_i, w_i)\}.\end{aligned}$$

Algorithm 3 Population dose-response estimation with DR-PCL-NET.

Input: Observations $\mathcal{D} = \{o_i\}_{i=1}^n$, bridge-fitting folds, and an independent evaluation fold $\mathcal{I}_3^{(\mathrm{DR})}$.
Subroutines: The bridge-fitting stages of Algorithms 1 and 2.
Output: Doubly robust population dose-response estimate $\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}}$.

1: Fit $\hat{h}$ and $\hat{\varphi}^{\mathrm{ATE}}$ on the bridge-fitting folds with Algorithms 1 and 2, respectively.
2: For each $i \in \mathcal{I}_3^{(\mathrm{DR})}$, form

$$y_i^{(\mathrm{DR})} = \hat{\varphi}^{\mathrm{ATE}}(a_i, x_i, z_i)\{y_i - \hat{h}(a_i, x_i, w_i)\}.$$

3: **if** A is continuous **then**
4: Fit $k^{\mathrm{DR}}(\cdot; \theta_3^{\mathrm{DR}})$ by the MSE in Equation 54.
5: Set

$$\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) = \frac{1}{|\mathcal{I}_3^{(\mathrm{DR})}|} \sum_{i \in \mathcal{I}_3^{(\mathrm{DR})}} \hat{h}(a, x_i, w_i) + k^{\mathrm{DR}}(a; \hat{\theta}_3^{\mathrm{DR}}).$$

6: **else**
7: Replace the correction regression by the within-a average of $y_i^{(\mathrm{DR})}$, and add the outcome-bridge average over $\mathcal{I}_3^{(\mathrm{DR})}$.
8: **end if**
9: **Return** $\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}}$.

F NEURAL MEAN EMBEDDING-BASED PROXY CAUSAL LEARNING FOR HETEROGENEOUS DOSE-RESPONSE

In this section, we derive the doubly robust algorithm for heterogeneous dose-response curve. In particular, similar to dose-response counterpart, we first derive the outcome bridge-based algorithm. Then, we derive the analogous treatment bridge algorithm, and then combine both into doubly robust algorithm.

F.1 HETEROGENEOUS DOSE-RESPONSE ESTIMATION: OUTCOME BRIDGE METHOD

Write $X = (S, V)$, and use $h(a, v, s, w) := h(a, (s, v), w)$ when displaying the effect modifier separately. No CATE-specific outcome bridge is required: we learn $\hat{h}(a, X, W)$ exactly as in Appendix E.1, with V included in X. Outcome-bridge identification gives

$$f_{\mathrm{CATE}}(a, v) = \mathbb{E}[h_0(a, v, S, W) \mid V = v]. \tag{56}$$

For this target, we use the factorized second-stage architecture

$$\hat{h}(a,v,s,w) = \left\{\hat{\boldsymbol{H}}^{(h)}\right\}^{\top}\left[\phi_{A,2}^{(h)}(a)\otimes\phi_{V,2}^{(h)}(v)\otimes c^{(h)}(s,w)\right], \quad c^{(h)}(s,w) := \phi_{S,2}^{(h)}(s)\otimes\phi_{W,2}^{(h)}(w), \tag{57}$$

where $\hat{\boldsymbol{H}}^{(h)} \in \mathbb{R}^{d_{h,2}\times d_Y}$ is the fitted linear head; it is a column vector when $d_Y = 1$. Conditional on the data used to fit $\hat{h}$, let

$$\mu_{SW|V}^{(h)}(v) := \mathbb{E}[c^{(h)}(S,W) \mid V = v].$$

Linearity of the fitted head gives

$$\mathbb{E}[\hat{h}(a,v,S,W) \mid V = v] = \left\{\hat{\boldsymbol{H}}^{(h)}\right\}^{\top}\left[\phi_{A,2}^{(h)}(a)\otimes\phi_{V,2}^{(h)}(v)\otimes\mu_{SW|V}^{(h)}(v)\right].$$

Let $\mathcal{I}_3^{(h)}$ be an independent third-stage split of size $n_{h,3}$. We estimate the single conditional mean embedding $\mu_{SW|V}^{(h)}$ by regressing the joint feature responses $c_i^{(h)} := c^{(h)}(s_i,w_i)$ on v_i. With a regression network $m^{(h)}(\cdot;\theta_3^{(h)})$, let

$$\hat{\theta}_3^{(h)} \in \arg\min_{\theta} \frac{1}{n_{h,3}} \sum_{i\in\mathcal{I}_3^{(h)}} \left\| c_i^{(h)} - m^{(h)}(v_i;\theta)\right\|_2^2$$

and write $\hat{\mu}_{SW|V}^{(h)}(v) := m^{(h)}(v;\hat{\theta}_3^{(h)})$. The resulting estimator is

$$\hat{f}_{\mathrm{CATE}}^{(h)}(a,v) := \left\{\hat{\boldsymbol{H}}^{(h)}\right\}^{\top}\left[\phi_{A,2}^{(h)}(a)\otimes\phi_{V,2}^{(h)}(v)\otimes\hat{\mu}_{SW|V}^{(h)}(v)\right]. \tag{58}$$

Because the learned embedding depends only on v, one third-stage fit can be reused for every intervention query a; no retraining over treatment levels is required. At the oracle bridge and population embedding, this expression equals $f_{\mathrm{CATE}}(a,v)$.

Algorithm 4 Heterogeneous dose-response estimation with OUTCOMENET.

Input: Observations $\mathcal{D} = \{(a_i,v_i,s_i,w_i,z_i,y_i)\}_{i=1}^n$ and disjoint folds $\mathcal{I}_1^{(h)},\mathcal{I}_2^{(h)},\mathcal{I}_3^{(h)}$.
Subroutine: The bridge-fitting stages of Algorithm 1.
Output: Fitted outcome bridge $\hat{h}$ and heterogeneous response estimate $\hat{f}_{\mathrm{CATE}}^{(h)}$.

1: Set $X = (S,V)$ and fit $\hat{h}$ with the factorized representation in Equation 57.
2: For each $i \in \mathcal{I}_3^{(h)}$, form
$$c_i^{(h)} = \phi_{S,2}^{(h)}(s_i)\otimes\phi_{W,2}^{(h)}(w_i).$$
3: Fit $m^{(h)}(\cdot;\theta_3^{(h)})$ by minimizing
$$\frac{1}{|\mathcal{I}_3^{(h)}|}\sum_{i\in\mathcal{I}_3^{(h)}} \|c_i^{(h)} - m^{(h)}(v_i;\theta_3^{(h)})\|_2^2,$$
and set $\hat{\mu}_{SW|V}^{(h)}(v) = m^{(h)}(v;\hat{\theta}_3^{(h)})$.
4: Construct $\hat{f}_{\mathrm{CATE}}^{(h)}(a,v)$ using Equation 58.
5: **Return** $\hat{h}$ and $\hat{f}_{\mathrm{CATE}}^{(h)}$.

F.2 HETEROGENEOUS DOSE-RESPONSE ESTIMATION: TREATMENT BRIDGE METHOD

We now extend TREATMENTNET to the heterogeneous dose-response function. As in the outcome-bridge subsection, we write $X = (S,V)$, where V denotes the effect modifiers of interest and S collects the remaining observed covariates.

For a raw CATE bridge candidate $\varphi(a,v,s,z)$, define

$$g_{\varphi}^{\mathrm{CATE}}(a,v,s,w) := (T_{\varphi}^{\mathrm{CATE}}\varphi)(a,v,s,w)$$

$$:= \mathbb{E}[\varphi(a, v, s, Z) \mid A = a, V = v, S = s, W = w].$$

The native bridge satisfies

$$T_\varphi^{\mathrm{CATE}} \varphi_0^{\mathrm{CATE}} = r_{0,\mathrm{CATE}}, \qquad f_{\mathrm{CATE}}(a, v) = \mathbb{E}[Y \varphi_0^{\mathrm{CATE}}(a, v, S, Z) \mid A = a, V = v].$$

To characterize a density-ratio-free objective, let $\mathbb{P}_{\mathrm{CATE}} := \mathbb{P}_{A,V,S,W}$ and define the conditionally decoupled law

$$\mathbb{Q}_{\mathrm{CATE}}(da, dv, ds, dw) := \mathbb{P}_V(dv)\, \mathbb{P}_{A|V}(da \mid v)\, \mathbb{P}_{S,W|V}(ds, dw \mid v). \tag{59}$$

Assume $\mathbb{Q}_{\mathrm{CATE}} \ll \mathbb{P}_{\mathrm{CATE}}$, and set

$$r_{0,\mathrm{CATE}} := \frac{d\mathbb{Q}_{\mathrm{CATE}}}{d\mathbb{P}_{\mathrm{CATE}}} \in L^2(\mathbb{P}_{\mathrm{CATE}}).$$

Under a common dominating measure, this derivative admits the version

$$r_{0,\mathrm{CATE}}(a, v, s, w) = \frac{p_{A|V}(a \mid v)}{p_{A|V,S,W}(a \mid v, s, w)} \quad \text{on the denominator support.}$$

For $u \in L^2(\mathbb{P}_{\mathrm{CATE}})$, define

$$\mathcal{R}^{\mathrm{RF}}_{\mathrm{CATE}}(u) := \mathbb{E}_{\mathbb{P}_{\mathrm{CATE}}}[u(A, V, S, W)^2] - 2\mathbb{E}_{\mathbb{Q}_{\mathrm{CATE}}}[u(A, V, S, W)]. \tag{60}$$

The linear expectation has the conditional-product representation

$$\mathbb{E}_{\mathbb{Q}_{\mathrm{CATE}}}[u] = \mathbb{E}_V \left[\mathbb{E}_{A^\dagger|V} \left\{ \mathbb{E}_{S,W|V} \left[u(A^\dagger, V, S, W) \right] \right\} \right],$$

where $A^\dagger \perp (S, W) \mid V$. Since $\mathbb{E}_{\mathbb{Q}_{\mathrm{CATE}}}[u] = \mathbb{E}_{\mathbb{P}_{\mathrm{CATE}}}[r_{0,\mathrm{CATE}} u]$, completing the square gives

$$\mathcal{R}^{\mathrm{RF}}_{\mathrm{CATE}}(u) - \mathcal{R}^{\mathrm{RF}}_{\mathrm{CATE}}(r_{0,\mathrm{CATE}}) = \|u - r_{0,\mathrm{CATE}}\|^2_{L^2(\mathbb{P}_{\mathrm{CATE}})}. \tag{61}$$

Thus, a native density-ratio-free CATE risk exists, but its linear term requires draws from the conditional product $\mathbb{P}_{A|V} \otimes \mathbb{P}_{S,W|V}$ at the same value of V. A global permutation of A instead samples from $\mathbb{P}_A$ and is generally biased for Equation 60. When V is discrete, within-V pairing is possible in sufficiently populated strata. For continuous or high-dimensional V, exact matching is unavailable and a conditional sampler, smoothing method, or additional conditional regression is required. We avoid this extra nuisance by reusing the density-ratio-free ATE bridge through Proposition 3.4.

Recall that CATE is identified with treatment bridge function via

$$f_{\mathrm{CATE}}(a, v) = \frac{\mathbb{E}\{Y \varphi_0^{\mathrm{ATE}}(a, v, S, Z) \mid A = a, V = v\}}{\mathbb{E}\{\varphi_0^{\mathrm{ATE}}(a, v, S, Z) \mid A = a, V = v\}}.$$

The denominator is the oracle normalizer:

$$D_0(a, v) := \mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, v, S, Z) \mid A = a, V = v] = \frac{d(\mathbb{P}_A \otimes \mathbb{P}_V)}{d\mathbb{P}_{A,V}}(a, v).$$

The last equality follows from the conditional-expectation property of Radon–Nikodym derivatives; see Proposition 3.4. Under a common dominating measure, it reduces to $p_A(a)/p_{A|V}(a \mid v)$ on the denominator support. We learn one $\hat{\varphi}^{\mathrm{ATE}}$ using all $X = (S, V)$, as described in Appendix E.2; no native CATE bridge is fitted.

On an independent third-stage split $\mathcal{I}_3^{(\varphi)}$, define

$$\hat{\varphi}_k^{\mathrm{ATE}} := \hat{\varphi}^{\mathrm{ATE}}(a_k, v_k, s_k, z_k), \qquad \gamma_{N,k}^{(\varphi)} := y_k \hat{\varphi}_k^{\mathrm{ATE}}, \qquad \gamma_{D,k}^{(\varphi)} := \hat{\varphi}_k^{\mathrm{ATE}}.$$

We fit a shared network with vector-valued numerator head $N^{(\varphi)}$ and scalar denominator head $D^{(\varphi)}$ by minimizing the joint empirical mean-squared error

$$\frac{1}{|\mathcal{I}_3^{(\varphi)}|} \sum_{k \in \mathcal{I}_3^{(\varphi)}} \left[\|\gamma_{N,k}^{(\varphi)} - N^{(\varphi)}(a_k, v_k)\|_2^2 + \{\gamma_{D,k}^{(\varphi)} - D^{(\varphi)}(a_k, v_k)\}^2 \right].$$

Denote the fitted heads by $\hat{N}^{(\varphi)}$ and $\hat{D}^{(\varphi)}$. Their population MSE targets are

$$N_{\hat{\varphi}}(a,v) := \mathbb{E}[Y\hat{\varphi}^{\mathrm{ATE}}(a,v,S,Z) \mid A=a, V=v],$$
$$D_{\hat{\varphi}}(a,v) := \mathbb{E}[\hat{\varphi}^{\mathrm{ATE}}(a,v,S,Z) \mid A=a, V=v].$$

For a small $\epsilon_D > 0$, the stabilized TREATMENTNET estimator is

$$\hat{f}^{(\varphi)}_{\mathrm{CATE}}(a,v) := \frac{\hat{N}^{(\varphi)}(a,v)}{\max\{\hat{D}^{(\varphi)}(a,v), \epsilon_D\}}. \tag{62}$$

Algorithm 5 Self-normalized heterogeneous dose-response estimation with TREATMENTNET, reusing the ATE-scale treatment bridge.

Input: Observations $\mathcal{D} = \{(a_i, v_i, s_i, w_i, z_i, y_i)\}_{i=1}^n$, disjoint folds $\mathcal{I}_1^{(\varphi)}, \mathcal{I}_2^{(\varphi)}, \mathcal{I}_3^{(\varphi)}$, and $\epsilon_D > 0$.
Subroutine: The bridge-fitting stages of Algorithm 2.
Output: Self-normalized heterogeneous response estimate $\hat{f}^{(\varphi)}_{\mathrm{CATE}}$.

1: Set $X = (S, V)$ and fit the single ATE-scale bridge $\hat{\varphi}^{\mathrm{ATE}}$; do not fit a native CATE bridge.
2: For each $i \in \mathcal{I}_3^{(\varphi)}$, form

$$\hat{\varphi}_i^{\mathrm{ATE}} = \hat{\varphi}^{\mathrm{ATE}}(a_i, v_i, s_i, z_i), \qquad \gamma_{N,i}^{(\varphi)} = y_i\hat{\varphi}_i^{\mathrm{ATE}}, \qquad \gamma_{D,i}^{(\varphi)} = \hat{\varphi}_i^{\mathrm{ATE}}.$$

3: Jointly fit a shared network with vector head $N^{(\varphi)}$ and scalar head $D^{(\varphi)}$ by minimizing

$$\frac{1}{|\mathcal{I}_3^{(\varphi)}|} \sum_{i \in \mathcal{I}_3^{(\varphi)}} \left[\|\gamma_{N,i}^{(\varphi)} - N^{(\varphi)}(a_i, v_i)\|_2^2 + \{\gamma_{D,i}^{(\varphi)} - D^{(\varphi)}(a_i, v_i)\}^2 \right].$$

4: Set

$$\hat{f}^{(\varphi)}_{\mathrm{CATE}}(a,v) = \frac{\hat{N}^{(\varphi)}(a,v)}{\max\{\hat{D}^{(\varphi)}(a,v), \epsilon_D\}}$$

as in Equation 62.
5: **Return** $\hat{f}^{(\varphi)}_{\mathrm{CATE}}$.

F.3 HETEROGENEOUS DOSE-RESPONSE ESTIMATION: DOUBLY ROBUST METHOD

Let h be an outcome-bridge candidate and φ an ATE-scale treatment-bridge candidate. For finite expectations and $D_\varphi(a,v) \neq 0$, define

$$M_h(a,v) := \mathbb{E}[h(a,v,S,W) \mid V=v],$$
$$K_{h,\varphi}(a,v) := \mathbb{E}[\varphi(a,v,S,Z)\{Y - h(a,v,S,W)\} \mid A=a, V=v],$$
$$D_\varphi(a,v) := \mathbb{E}[\varphi(a,v,S,Z) \mid A=a, V=v].$$

The self-normalized functional is

$$\Psi^{\mathrm{SN}}_{\mathrm{CATE}}(a,v;h,\varphi) := M_h(a,v) + \frac{K_{h,\varphi}(a,v)}{D_\varphi(a,v)}. \tag{63}$$

It equals $f_{\mathrm{CATE}}(a,v)$ if either $h = h_0$, or $\varphi = \varphi_0^{\mathrm{ATE}}$.

For estimation, reuse $\hat{f}^{(h)}_{\mathrm{CATE}}$ from Appendix F.1. On a third-stage fold $\mathcal{I}_3^{(\mathrm{DR})}$, held out from bridge fitting, set

$$\hat{\varphi}_i^{\mathrm{ATE}} := \hat{\varphi}^{\mathrm{ATE}}(a_i, v_i, s_i, z_i), \qquad \hat{h}_i := \hat{h}(a_i, v_i, s_i, w_i),$$
$$\gamma_{K,i}^{\mathrm{DR}} := \hat{\varphi}_i^{\mathrm{ATE}}(y_i - \hat{h}_i), \qquad \gamma_{D,i}^{\mathrm{DR}} := \hat{\varphi}_i^{\mathrm{ATE}}.$$

As in Appendix F.2, jointly regress these paired labels on (a_i, v_i) by the same joint MSE, with $\gamma_{K,i}^{\mathrm{DR}}$ replacing $\gamma_{N,i}^{(\varphi)}$, obtaining the vector-valued correction $\hat{K}^{\mathrm{DR}}$ and scalar normalizer $\hat{D}^{\mathrm{DR}}$. Thus, only the numerator label differs from the TREATMENTNET third stage. The estimator is

$$\hat{f}^{\mathrm{DR}}_{\mathrm{CATE}}(a,v) := \hat{f}^{(h)}_{\mathrm{CATE}}(a,v) + \frac{\hat{K}^{\mathrm{DR}}(a,v)}{\max\{\hat{D}^{\mathrm{DR}}(a,v), \epsilon_D\}}, \tag{64}$$

with the same denominator stabilization as in Equation 62.

Algorithm 6 Heterogeneous dose-response estimation with DR-PCL-NET.

Input: Observations $\mathcal{D} = \{(a_i, v_i, s_i, w_i, z_i, y_i)\}_{i=1}^n$, bridge-fitting folds, an independent fold $\mathcal{I}_3^{(\mathrm{DR})}$, and $\epsilon_D > 0$.
Subroutines: Algorithms 4 and 2.
Output: Doubly robust heterogeneous response estimate $\hat{f}_{\mathrm{CATE}}^{\mathrm{DR}}$.

1: Fit $\hat{h}$ and $\hat{f}_{\mathrm{CATE}}^{(h)}$ with Algorithm 4.
2: With $X = (S, V)$, fit the ATE-scale bridge $\hat{\varphi}^{\mathrm{ATE}}$ using the bridge stages of Algorithm 2.
3: For each $i \in \mathcal{I}_3^{(\mathrm{DR})}$, set

$$\hat{h}_i = \hat{h}(a_i, v_i, s_i, w_i), \qquad \hat{\varphi}_i^{\mathrm{ATE}} = \hat{\varphi}^{\mathrm{ATE}}(a_i, v_i, s_i, z_i),$$

$$\gamma_{K,i}^{\mathrm{DR}} = \hat{\varphi}_i^{\mathrm{ATE}}(y_i - \hat{h}_i), \qquad \gamma_{D,i}^{\mathrm{DR}} = \hat{\varphi}_i^{\mathrm{ATE}}.$$

4: Jointly regress $(\gamma_{K,i}^{\mathrm{DR}}, \gamma_{D,i}^{\mathrm{DR}})$ on (a_i, v_i) by squared Euclidean and scalar MSE, obtaining $\hat{K}^{\mathrm{DR}}$ and $\hat{D}^{\mathrm{DR}}$.
5: Set

$$\hat{f}_{\mathrm{CATE}}^{\mathrm{DR}}(a, v) = \hat{f}_{\mathrm{CATE}}^{(h)}(a, v) + \frac{\hat{K}^{\mathrm{DR}}(a, v)}{\max\{\hat{D}^{\mathrm{DR}}(a, v), \epsilon_D\}}$$

as in Equation 64.
6: **Return** $\hat{f}_{\mathrm{CATE}}^{\mathrm{DR}}$.

G NEURAL MEAN EMBEDDING-BASED PROXY CAUSAL LEARNING FOR TREATMENT-CONDITIONAL DOSE RESPONSE

We next derive our proposed neural mean embedding algorithms for treatment-conditional dose response estimation. We first present the outcome bridge route. Then, we derive the treatment bridge route and combine it with the outcome bridge method to derive the doubly robust estimator.

G.1 CONDITIONAL DOSE-RESPONSE ESTIMATION: OUTCOME BRIDGE METHOD

For $f_{\mathrm{ATT}}(a, a') = \mathbb{E}\{Y^{(a)} \mid A = a'\}$, a is the intervention level and a' indexes the factual reference population. Outcome-bridge identification gives

$$f_{\mathrm{ATT}}(a, a') = \mathbb{E}[h_0(a, X, W) \mid A = a'].$$

Thus, no ATT-specific outcome bridge is required: we reuse the OUTCOMENET estimate $\hat{h}$ from Appendix E.1. For the factorized second-stage representation, write

$$\hat{h}(a, x, w) = \left\{\hat{\boldsymbol{H}}^{(h)}\right\}^\top \left[\phi_{A,2}^{(h)}(a) \otimes c_{XW}^{(h)}(x, w)\right], \qquad c_{XW}^{(h)}(x, w) := \phi_{X,2}^{(h)}(x) \otimes \phi_{W,2}^{(h)}(w).$$

Conditional on the fitted bridge, define $\mu_{XW|A}^{(h)}(a') := \mathbb{E}[c_{XW}^{(h)}(X, W) \mid A = a']$.

On an independent third-stage split $\mathcal{I}_3^{(h)}$, estimate this conditional mean by the empirical MSE regression

$$\hat{\mu}_{XW|A}^{(h)} \in \arg\min_m \frac{1}{|\mathcal{I}_3^{(h)}|} \sum_{i \in \mathcal{I}_3^{(h)}} \|c_{XW}^{(h)}(x_i, w_i) - m(a_i)\|_2^2.$$

The resulting estimator is

$$\hat{f}_{\mathrm{ATT}}^{(h)}(a, a') := \left\{\hat{\boldsymbol{H}}^{(h)}\right\}^\top \left[\phi_{A,2}^{(h)}(a) \otimes \hat{\mu}_{XW|A}^{(h)}(a')\right]. \tag{65}$$

The same third-stage fit yields the entire (a, a')-surface; no retraining over either treatment argument is required. For finite discrete A, the MSE regression reduces to the within-level average of $c_{XW}^{(h)}(X, W)$; for continuous A, m is a neural network. At the oracle bridge and population conditional mean embedding, Equation 65 equals $f_{\mathrm{ATT}}(a, a')$.

Algorithm 7 Treatment-conditional dose-response estimation with OUTCOMENET.

Input: Observations $\mathcal{D} = \{o_i\}_{i=1}^n$ and disjoint folds $\mathcal{I}_1^{(h)}, \mathcal{I}_2^{(h)}, \mathcal{I}_3^{(h)}$.
Subroutine: The bridge-fitting stages of Algorithm 1.
Output: Treatment-conditional response surface $\hat{f}_{\mathrm{ATT}}^{(h)}$.

1: Fit the population outcome bridge $\hat{h}$; no ATT-specific outcome bridge is required.
2: For each $i \in \mathcal{I}_3^{(h)}$, form

$$c_{XW,i}^{(h)} = \phi_{X,2}^{(h)}(x_i) \otimes \phi_{W,2}^{(h)}(w_i).$$

3: Fit $m^{(h)}(\cdot; \theta_3^{(h)})$ by minimizing

$$\frac{1}{|\mathcal{I}_3^{(h)}|} \sum_{i \in \mathcal{I}_3^{(h)}} \|c_{XW,i}^{(h)} - m^{(h)}(a_i; \theta_3^{(h)})\|_2^2,$$

and set $\hat{\mu}_{XW|A}^{(h)}(a') = m^{(h)}(a'; \hat{\theta}_3^{(h)})$.
4: Construct $\hat{f}_{\mathrm{ATT}}^{(h)}(a, a')$ using Equation 65.
5: **Return** the full surface $\hat{f}_{\mathrm{ATT}}^{(h)}$.

G.2 CONDITIONAL DOSE-RESPONSE ESTIMATION: TREATMENT BRIDGE METHOD

For the conditional dose-response surface $f_{\mathrm{ATT}}(a, a') = \mathbb{E}\{Y^{(a)} \mid A = a'\}$, the first argument a is the intervention level, whereas a' indexes the factual reference population. Define the probability laws

$$\begin{aligned}\mathbb{P}_{\mathrm{ATT}}(da, da', dx, dw) &:= \mathbb{P}_A(da)\mathbb{P}_A(da')\mathbb{P}_{X,W|A=a}(dx, dw),\\ \mathbb{Q}_{\mathrm{ATT}}(da, da', dx, dw) &:= \mathbb{P}_A(da)\mathbb{P}_A(da')\mathbb{P}_{X,W|A=a'}(dx, dw).\end{aligned} \tag{66}$$

Assume $\mathbb{Q}_{\mathrm{ATT}} \ll \mathbb{P}_{\mathrm{ATT}}$ on the treatment pairs of interest, and let

$$r_{0,\mathrm{ATT}} := \frac{d\mathbb{Q}_{\mathrm{ATT}}}{d\mathbb{P}_{\mathrm{ATT}}} \in L^2(\mathbb{P}_{\mathrm{ATT}}).$$

For a raw bridge candidate $\varphi : \mathcal{A} \times \mathcal{A} \times \mathcal{X} \times \mathcal{Z} \to \mathbb{R}$, define its conditional projection

$$(T_{\mathrm{ATT}}\varphi)(a, a', x, w) =: g_\varphi^{\mathrm{ATT}}(a, a', x, w) := \mathbb{E}[\varphi(a, a', x, Z) \mid A = a, X = x, W = w].$$

A valid ATT bridge satisfies $T_{\mathrm{ATT}}\varphi_0^{\mathrm{ATT}} = r_{0,\mathrm{ATT}}$.

Let (A_1, X_1, W_1) and (A_2, X_2, W_2) be independent observational units. For $u \in L^2(\mathbb{P}_{\mathrm{ATT}})$, consider

$$\mathcal{R}_{\mathrm{ATT}}^{\mathrm{RF}}(u) := \mathbb{E}\big[u(A_1, A_2, X_1, W_1)^2 - 2u(A_1, A_2, X_2, W_2)\big]. \tag{67}$$

The same change-of-measure argument as in Lemma 4.1 yields

$$\mathcal{R}_{\mathrm{ATT}}^{\mathrm{RF}}(u) - \mathcal{R}_{\mathrm{ATT}}^{\mathrm{RF}}(r_{0,\mathrm{ATT}}) = \|u - r_{0,\mathrm{ATT}}\|_{L^2(\mathbb{P}_{\mathrm{ATT}})}^2. \tag{68}$$

Thus, the reference level can be treated as an input: one bridge estimates the entire (a, a')-surface without retraining at each a'.

The first-stage conditional mean embedding is unchanged from Appendix E.2. We use the compact parameterization

$$\begin{aligned}\mu_Z^{(\varphi)}(a, x, w) &:= \mathbb{E}[\phi_{Z,2}^{(\varphi)}(Z) \mid A = a, X = x, W = w],\\ \hat{\mu}_Z^{(\varphi)}(a, x, w) &= \big(\hat{\boldsymbol{V}}^{(\varphi)}\big)^\top \phi_{AXW,1}^{(\varphi)}(a, x, w),\\ \varphi(a, a', x, z) &= \boldsymbol{\varphi}^\top \left[\phi_{AA'X,2}^{(\varphi)}(a, a', x) \otimes \phi_{Z,2}^{(\varphi)}(z)\right],\\ \hat{g}_\varphi^{\mathrm{ATT}}(a, a', x, w) &= \boldsymbol{\varphi}^\top \left[\phi_{AA'X,2}^{(\varphi)}(a, a', x) \otimes \hat{\mu}_Z^{(\varphi)}(a, x, w)\right].\end{aligned}$$

Here $\phi_{AA'X,2}^{(\varphi)}$ is one joint featurizer of (a, a', x). For a second-stage batch of size M, its complete off-diagonal objective is

$$\widehat{\mathcal{R}}_{\mathrm{ATT},2}^{\mathrm{RF,U}}(\varphi) = \frac{1}{M(M-1)} \sum_{i \neq j} \left\{\hat{g}_\varphi^{\mathrm{ATT}}(a_i, a_j, x_i, w_i)^2 - 2\hat{g}_\varphi^{\mathrm{ATT}}(a_i, a_j, x_j, w_j)\right\}. \tag{69}$$

We apply the same alternating feature updates and proximal closed-form head updates as for TREATMENTNET. Crucially, the transported term evaluates the embedding at $\hat{\mu}_Z^{(\varphi)}(a_i, x_j, w_j)$: i supplies the source treatment, while j supplies the intact anchor block (a_j, x_j, w_j). Replacing a_i by a_j changes the transported law.

Finally, freeze the bridge and use an independent third-stage split of size t_{ATT}. For each ordered pair $i \neq j$, form

$$\gamma_{ij}^{\mathrm{ATT}} := y_i \hat{\varphi}(a_i, a_j, x_i, z_i).$$

Fit a two-input regression $k^{\varphi}(\cdot, \cdot; \theta_3^{\varphi})$ by minimizing the pairwise empirical MSE

$$\frac{1}{t_{\mathrm{ATT}}(t_{\mathrm{ATT}} - 1)} \sum_{i \neq j} \left\| \gamma_{ij}^{\mathrm{ATT}} - k^{\varphi}(a_i, a_j; \theta_3^{\varphi}) \right\|_2^2, \qquad \hat{f}_{\mathrm{ATT}}^{(\varphi)}(a, a') := k^{\varphi}(a, a'; \hat{\theta}_3^{\varphi}).$$

Indeed, at the oracle bridge,

$$\mathbb{E}[Y_1 \varphi_0^{\mathrm{ATT}}(A_1, A_2, X_1, Z_1) \mid A_1 = a, A_2 = a'] = f_{\mathrm{ATT}}(a, a').$$

Algorithm 8 Treatment-conditional dose-response estimation with density-ratio-free TREATMENTNET.

Input: Observations $\mathcal{D} = \{o_i\}_{i=1}^n$, disjoint bridge-fitting folds $\mathcal{I}_1^{(\varphi)}, \mathcal{I}_2^{(\varphi)}$, and an independent final-stage fold $\mathcal{I}_3^{(\varphi)}$.
Parameters: Neural parameters $\theta_1^{(\varphi)}, \theta_2^{(\varphi)}, \theta_3^{\varphi}$, first-stage operator $\boldsymbol{V}^{(\varphi)}$, bridge head $\boldsymbol{\varphi}$, and final network k^{φ}.
Hyperparameters: Learning rates $\eta_1^{(\varphi)}, \eta_2^{(\varphi)}$, proximal coefficients $\lambda_1^{(\varphi)}, \lambda_2^{(\varphi)}$, and update counts $T_{\varphi,1}, T_{\varphi,2}$.
Output: Fitted joint-anchor bridge $\hat{\varphi}^{\mathrm{ATT}}$ and treatment-conditional response surface $\hat{f}_{\mathrm{ATT}}^{(\varphi)}$.

repeat
  Sample mini-batches $\mathcal{B}_1^{(\varphi)} \subset \mathcal{I}_1^{(\varphi)}$ and $\mathcal{B}_2^{(\varphi)} \subset \mathcal{I}_2^{(\varphi)}$.
  for $s = 1, \ldots, T_{\varphi,1}$ **do**
    Update $\theta_1^{(\varphi)}$ by one gradient step on Equation 46, holding the proxy-target featurizer fixed.
    Update the persistent operator $\hat{\boldsymbol{V}}^{(\varphi)}$ by Equation 47.
  end for
  for $s = 1, \ldots, T_{\varphi,2}$ **do**
    Compute $\check{\boldsymbol{V}}^{(\varphi)}$ on $\mathcal{B}_2^{(\varphi)}$ using Equation 48.
    For each ordered $i \neq j$ in $\mathcal{B}_2^{(\varphi)}$, form the matched and transported projected features

$$q_{ij}^{\mathrm{mat}} = \phi_{AA'X,2}^{(\varphi)}(a_i, a_j, x_i) \otimes \{\check{\boldsymbol{V}}^{(\varphi)}\}^{\top} \phi_{AXW,1}^{(\varphi)}(a_i, x_i, w_i),$$
$$q_{ij}^{\mathrm{tr}} = \phi_{AA'X,2}^{(\varphi)}(a_i, a_j, x_j) \otimes \{\check{\boldsymbol{V}}^{(\varphi)}\}^{\top} \phi_{AXW,1}^{(\varphi)}(a_i, x_j, w_j).$$

    Update $\theta_2^{(\varphi)}$ by one gradient step on the fixed ratio-free U-objective in Equation 69, differentiating through $\hat{\boldsymbol{V}}^{(\varphi)}$ for the proxy-side features.
    Recompute the projected features and update $\hat{\boldsymbol{\varphi}}$ by the exact proximal closed-form minimizer of

$$\widehat{\mathcal{R}}_{\mathrm{ATT},2}^{\mathrm{RF,U}}(\boldsymbol{\varphi}) + \lambda_2^{(\varphi)} \|\boldsymbol{\varphi} - \hat{\boldsymbol{\varphi}}_{\mathrm{old}}\|_2^2,$$

    where $\hat{\boldsymbol{\varphi}}_{\mathrm{old}}$ is the head before this update.
  end for
until convergence
Set

$$\hat{\varphi}^{\mathrm{ATT}}(a, a', x, z) = \hat{\boldsymbol{\varphi}}^{\top} [\phi_{AA'X,2}^{(\varphi)}(a, a', x) \otimes \phi_{Z,2}^{(\varphi)}(z)].$$

For all ordered $i \neq j$ in $\mathcal{I}_3^{(\varphi)}$, form

$$\gamma_{ij}^{\mathrm{ATT}} = y_i \hat{\varphi}^{\mathrm{ATT}}(a_i, a_j, x_i, z_i).$$

Fit $k^{\varphi}(a_i, a_j; \theta_3^{\varphi})$ by the ordered-pair MSE displayed above and set $\hat{f}_{\mathrm{ATT}}^{(\varphi)}(a, a') = k^{\varphi}(a, a'; \hat{\theta}_3^{\varphi})$.
Return $\hat{\varphi}^{\mathrm{ATT}}$ and the full surface $\hat{f}_{\mathrm{ATT}}^{(\varphi)}$.

G.3 CONDITIONAL DOSE-RESPONSE ESTIMATION: DOUBLY ROBUST METHOD

Recall the doubly robust functional

$$\Psi_{\mathrm{ATT}}(a, a'; h, \varphi) = \mathbb{E}[\varphi(a, a', X, Z)\{Y - h(a, X, W)\} \mid A = a] + \mathbb{E}[h(a, X, W) \mid A = a'],$$

which equals $f_{\mathrm{ATT}}(a, a')$ if either $h = h_0$ or $\varphi = \varphi_0^{\mathrm{ATT}}$, by Theorem 3.3. Conditional on the fitted bridges, let O_1 and O_2 be independent final-stage observations and define

$$\Gamma_{12}^{\mathrm{ATT,DR}} := \hat{\varphi}^{\mathrm{ATT}}(A_1, A_2, X_1, Z_1)\{Y_1 - \hat{h}(A_1, X_1, W_1)\} + \hat{h}(A_1, X_2, W_2).$$

Independence gives

$$\mathbb{E}[\Gamma_{12}^{\mathrm{ATT,DR}} \mid A_1 = a, A_2 = a'] = \Psi_{\mathrm{ATT}}(a, a'; \hat{h}, \hat{\varphi}^{\mathrm{ATT}}).$$

We therefore regress this ordered-pair pseudo-outcome on (A_1, A_2), as summarized in Algorithm 9; the exact remainder is analyzed in Appendix H.10.

Algorithm 9 Treatment-conditional dose-response estimation with DR-PCL-NET.

Input: Observations $\mathcal{D} = \{(a_i, x_i, w_i, z_i, y_i)\}_{i=1}^n$, bridge-fitting folds, and an independent final-stage fold $\mathcal{I}_3^{(\mathrm{DR})}$.
Subroutines: Algorithms 1 and 8.
Output: Doubly robust treatment-conditional response surface $\hat{f}_{\mathrm{ATT}}^{\mathrm{DR}}$.

1: Fit the population outcome bridge $\hat{h}$ and the joint-anchor treatment bridge $\hat{\varphi}^{\mathrm{ATT}}$.
2: For every ordered pair $i \neq j$ in $\mathcal{I}_3^{(\mathrm{DR})}$, form

$$\gamma_{ij}^{\mathrm{ATT,DR}} = \hat{\varphi}^{\mathrm{ATT}}(a_i, a_j, x_i, z_i)\{y_i - \hat{h}(a_i, x_i, w_i)\} + \hat{h}(a_i, x_j, w_j).$$

3: Fit a two-input network $k^{\mathrm{DR}}(\cdot, \cdot; \theta_3^{\mathrm{DR}})$ by minimizing

$$\frac{1}{t_{\mathrm{DR}}(t_{\mathrm{DR}} - 1)} \sum_{i \neq j} \|\gamma_{ij}^{\mathrm{ATT,DR}} - k^{\mathrm{DR}}(a_i, a_j; \theta_3^{\mathrm{DR}})\|_2^2, \qquad t_{\mathrm{DR}} := |\mathcal{I}_3^{(\mathrm{DR})}|.$$

4: Set $\hat{f}_{\mathrm{ATT}}^{\mathrm{DR}}(a, a') = k^{\mathrm{DR}}(a, a'; \hat{\theta}_3^{\mathrm{DR}})$.
5: **Return** the full surface $\hat{f}_{\mathrm{ATT}}^{\mathrm{DR}}$.

H RADEMACHER-COMPLEXITY-BASED ERROR BOUNDS FOR THE PROPOSED ESTIMATORS

We derive Rademacher-complexity-based error bounds for the projected bridges and the induced causal functions. We assume independent splits of the dataset for each regression stage. The analysis allows stage-specific approximate empirical minimization.

Assumption H.1. *For every empirical-risk stage j considered below, let $\widehat{\mathcal{R}}_j$ denote its normalized empirical criterion and $\mathcal{V}_j$ its candidate class. We assume that the fitted estimator $\hat{\vartheta}_j$ is measurable and satisfies*

$$\widehat{\mathcal{R}}_j(\hat{\vartheta}_j) \leq \inf_{\vartheta \in \mathcal{V}_j} \widehat{\mathcal{R}}_j(\vartheta) + \xi_j, \qquad \xi_j \geq 0,$$

where the possibly data-dependent gap ξ_j is measured on the same normalized criterion scale as $\widehat{\mathcal{R}}_j$.

Exact empirical minimization is the special case $\xi_j = 0$. Note that this assumption is standard in the analysis of empirical risk minimization estimators (van der Vaart & Wellner, 1996, see section 3.2). For the profiled bridge first stages, the corresponding inequality is required uniformly over the candidate proxy-feature map. Because the projected errors below are squared excess risks, the gaps enter their bounds additively. We do not derive algorithm-specific rates for the implemented nonconvex optimizer; the guarantees apply to any fitted estimators satisfying the displayed empirical suboptimality conditions.

H.1 Background on Rademacher complexity analysis

The resulting error bounds depend on the complexity of the function classes used at each regression stage. Therefore, we first provide a review of Rademacher complexity analysis here.

Let $\mathcal{O}$ be a measurable space, let $\mathcal{F}$ be a class of measurable functions $f : \mathcal{O} \to [0, B_{\mathcal{F}}]$, and let $\boldsymbol{O}_N = (O_1, \dots, O_N)$. Write

$$\mathbb{P}f := \mathbb{E}\{f(O)\}, \qquad \mathbb{P}_N f := \frac{1}{N}\sum_{i=1}^{N} f(O_i).$$

Definition H.2 (Empirical Rademacher complexity)**.** *The empirical Rademacher complexity of $\mathcal{F}$ is*

$$\widehat{\mathfrak{R}}_{\boldsymbol{O}_N}(\mathcal{F}) := \mathbb{E}_\sigma\left[\sup_{f\in\mathcal{F}} \frac{1}{N}\sum_{i=1}^{N} \sigma_i f(O_i) \,\middle|\, \boldsymbol{O}_N\right],$$

where $\sigma_1, \dots, \sigma_N$ are independent Rademacher variables.

Proposition H.3 (Uniform Rademacher deviation; Mohri et al. (2018))**.** *If $O_1, \dots, O_N \overset{\text{iid}}{\sim} \mathbb{P}$, then, for every $\delta \in (0,1)$, each of the following statements holds with probability at least $1-\delta$, uniformly over $f \in \mathcal{F}$:*

$$\mathbb{P}f \le \mathbb{P}_N f + 2\widehat{\mathfrak{R}}_{\boldsymbol{O}_N}(\mathcal{F}) + 3B_{\mathcal{F}}\sqrt{\frac{\log(2/\delta)}{2N}}.$$

$$\mathbb{P}_N f \le \mathbb{P}f + 2\widehat{\mathfrak{R}}_{\boldsymbol{O}_N}(\mathcal{F}) + 3B_{\mathcal{F}}\sqrt{\frac{\log(2/\delta)}{2N}}.$$

We also need the corresponding bound for an order-two U-process. Let $O_1, \dots, O_m \overset{\text{iid}}{\sim} \mathbb{P}$, write $\boldsymbol{O}_m = (O_1, \dots, O_m)$, let O, O' be independent draws from $\mathbb{P}$, and let $\mathcal{L}$ be a class of measurable kernels $\ell : \mathcal{O}^2 \to \mathbb{R}$. Define

$$\mathbb{P}^{\otimes 2}\ell := \mathbb{E}\{\ell(O, O')\}, \qquad \mathbb{U}_m\ell := \frac{1}{m(m-1)}\sum_{i\neq j} \ell(O_i, O_j).$$

The kernels need not be symmetric. Indeed, with $\ell^{\mathrm{sym}}(o, o') := \{\ell(o, o') + \ell(o', o)\}/2$,

$$\mathbb{U}_m\ell = \binom{m}{2}^{-1}\sum_{i<j} \ell^{\mathrm{sym}}(O_i, O_j), \qquad \mathbb{P}^{\otimes 2}\ell^{\mathrm{sym}} = \mathbb{P}^{\otimes 2}\ell.$$

Set $k_m := \lfloor m/2 \rfloor$. Let Π be a uniformly random permutation of $\{1, \dots, m\}$, independent of the sample, and let $\sigma_1, \dots, \sigma_{k_m}$ be independent Rademacher variables.

Definition H.4 (Pairwise empirical Rademacher complexity)**.** *The random-pairing complexity of $\mathcal{L}$ is*

$$\widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L}) := \mathbb{E}_{\Pi,\sigma}\left[\sup_{\ell\in\mathcal{L}} \frac{1}{k_m}\sum_{r=1}^{k_m} \sigma_r \ell\big(O_{\Pi(2r-1)}, O_{\Pi(2r)}\big) \,\middle|\, \boldsymbol{O}_m\right].$$

The random-pairing reduction is classical; see Clémençon et al. (2008, Appendix A, Lemma A.1).

Proposition H.5 (Uniform U-process Rademacher deviation)**.** *Suppose $m \ge 2$, the relevant suprema are measurable, and*

$$M_{\mathcal{L}} := \sup_{\ell\in\mathcal{L}} \sup_{o,o'\in\mathcal{O}} |\ell(o, o')| < \infty.$$

For every $\delta \in (0,1)$, each of the following statements holds with probability at least $1-\delta$, uniformly over $\ell \in \mathcal{L}$:

$$\mathbb{P}^{\otimes 2}\ell \le \mathbb{U}_m\ell + 2\widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L}) + 3M_{\mathcal{L}}\sqrt{\frac{8\log(2/\delta)}{m}},$$

$$\mathbb{U}_m\ell \le \mathbb{P}^{\otimes 2}\ell + 2\widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L}) + 3M_{\mathcal{L}}\sqrt{\frac{8\log(2/\delta)}{m}}.$$

Proof. For a permutation Π, define the disjoint-pair average

$$V_{m,\Pi}\ell := \frac{1}{k_m}\sum_{r=1}^{k_m}\ell\big(O_{\Pi(2r-1)}, O_{\Pi(2r)}\big)\,.$$

Each ordered pair has the same probability with each paired position. Hence $\mathbb{E}_\Pi[V_{m,\Pi}\ell \mid \boldsymbol{O}_m] = \mathbb{U}_m\ell$. For

$$Z_+ := \sup_{\ell\in\mathcal{L}}\{\mathbb{U}_m\ell - \mathbb{P}^{\otimes 2}\ell\}, \qquad Z_- := \sup_{\ell\in\mathcal{L}}\{\mathbb{P}^{\otimes 2}\ell - \mathbb{U}_m\ell\},$$

Jensen's inequality and ordinary symmetrization of the k_m independent pairs give, for either sign,

$$\mathbb{E}Z_\pm \le 2\mathbb{E}\Big[\widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L})\Big]\,.$$

For Z_-, apply the same symmetrization to $-\mathcal{L}$; symmetry of the Rademacher signs gives $\widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(-\mathcal{L}) = \widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L})$. Replacing one observation changes Z_+ or Z_- by at most $4M_\mathcal{L}/m$. It changes the pairwise empirical complexity by at most the same amount: a fixed matching changes by at most $2M_\mathcal{L}/k_m$, and the replaced observation is included with probability $2k_m/m$. Therefore, by McDiarmid's inequality, for a fixed choice of sign, with probability at least $1-\delta$,

$$Z_\pm \le \mathbb{E}Z_\pm + t_\delta, \qquad \mathbb{E}\Big[\widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L})\Big] \le \widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L}) + t_\delta,$$

simultaneously, where $t_\delta = M_\mathcal{L}\sqrt{8\log(2/\delta)/m}$. Thus $Z_\pm \le 2\widehat{\mathfrak{R}}^{(2),U}_{\boldsymbol{O}_m}(\mathcal{L}) + 3t_\delta$. The two choices of sign give the two stated inequalities separately. □

Although $\mathbb{U}_m\ell$ averages $m(m-1)$ dependent terms, its effective sample size in this first-order bound is $k_m = \lfloor m/2\rfloor$, which yields the usual $m^{-1/2}$ scale.

H.2 Outcome bridge consistency for dose-response estimation

We establish projected consistency of OutcomeNet for $Y \in \mathbb{R}^{d_Y}$. Let

$$Q_h := (A, X, Z) \subset \mathcal{Q}_h, \qquad m_0(q_h) := \mathbb{E}[Y \mid Q_h = q_h] \in \mathbb{R}^{d_Y}.$$

For $u : \mathcal{A}\times\mathcal{X}\times\mathcal{W} \to \mathbb{R}^{d_Y}$, define

$$(T_h u)(a, x, z) := \mathbb{E}[u(a, x, W) \mid A = a, X = x, Z = z].$$

Its projected seminorm is

$$\|u\|_{T_h} := \|T_h u\|_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})}.$$

Let the independent stage samples be $D_1^{(h)} = \{(\bar{q}_{h,i}, \bar{w}_i)\}_{i\in\mathcal{I}_1^{(h)}}$, $D_2^{(h)} = \{(\tilde{y}_i, \tilde{q}_{h,i})\}_{i\in\mathcal{I}_2^{(h)}}$, with $|\mathcal{I}_1^{(h)}| = n_h$ and $|\mathcal{I}_2^{(h)}| = m_h$. The first sample is i.i.d. from $\mathbb{P}_{Q_h,W}$; the second is i.i.d. from $\mathbb{P}_{Y,Q_h}$. We consider a learned-feature bridge class. Let $\nu_h = (n_h, m_h)$ and define

$$\Phi_{\nu_h} \subset \{\phi : \mathcal{W} \to \mathbb{R}^{d_{\nu_h}}\}, \qquad \mathcal{M}_{\nu_h} \subset \{\eta : \mathcal{Q}_h \to \mathbb{R}^{d_{\nu_h}}\},$$

for the class of learned outcome-proxy feature maps and the first-stage conditional mean regressors, respectively. We represent the multi-output head by a matrix

$$\Theta_{\nu_h} \subset \{\theta : \mathcal{A}\times\mathcal{X} \to \mathbb{R}^{d_Y\times d_{\nu_h}}\},$$

and the corresponding bridge class is $\mathcal{H}_{\nu_h} := \{h_{\theta,\phi}(a, x, w) = \theta(a, x)\phi(w) : \theta \in \Theta_{\nu_h},\ \phi \in \Phi_{\nu_h}\}$. The usual tensor-product architecture is subsumed in this notation: its final later is absorbed into $\theta(a, x)$.

For each $\phi \in \Phi_{\nu_h}$, define the population conditional feature mean $\mu_\phi^{(h)}(q_h) := \mathbb{E}[\phi(W) \mid Q_h = q_h]$. Then, for $q_h = (a, x, z)$, $(T_h h_{\theta,\phi})(q_h) = \langle\theta(a, x), \mu_\phi^{(h)}(q_h)\rangle$.

The population and empirical first-stage risks are

$$\mathcal{R}_{h,1}(\phi, \eta) := \mathbb{E}\|\phi(W) - \eta(Q_h)\|_2^2,$$

$$\widehat{\mathcal{R}}_{h,1}(\phi,\eta) := \frac{1}{n_h}\sum_{i=1}^{n_h} \|\phi(\bar{w}_i) - \eta(\bar{q}_{h,i})\|_2^2.$$

For every ϕ, let $\widehat{\mu}_\phi^{(h)} \in \mathcal{M}_{\nu_h}$ be a measurable approximate empirical minimizer, that is,

$$\widehat{\mathcal{R}}_{h,1}(\phi,\widehat{\mu}_\phi^{(h)}) \le \inf_{\eta\in\mathcal{M}_{\nu_h}} \widehat{\mathcal{R}}_{h,1}(\phi,\eta) + \xi_{h,1,\nu_h}, \qquad \xi_{h,1,\nu_h} \ge 0,$$

uniformly over $\phi \in \Phi_{\nu_h}$; the exact-minimizer case is $\xi_{h,1,\nu_h} = 0$.

Define the population and plug-in projected predictions

$$m_{\theta,\phi}^{(h)}(q_h) := \langle \theta(a,x), \mu_\phi^{(h)}(q_h)\rangle,$$
$$\widehat{m}_{\theta,\phi}^{(h)}(q_h) := \langle \theta(a,x), \widehat{\mu}_\phi^{(h)}(q_h)\rangle.$$

The profiled second-stage risk is

$$\widehat{\mathcal{R}}_{h,2}(\theta,\phi) := \frac{1}{m_h}\sum_{i\in\mathcal{I}_2^h} \left\| y_i - \widehat{m}_{\theta,\phi}^{(h)}(q_{h,i})\right\|_2^2.$$

Let $(\hat{\theta}^{(h)},\hat{\phi}^{(h)}) \in \Theta_{\nu_h}\times\Phi_{\nu_h}$ be a measurable approximate empirical minimizer of the second-stage risk,

$$\widehat{\mathcal{R}}_{h,2}(\hat{\theta}^{(h)},\hat{\phi}^{(h)}) \le \inf_{\theta\in\Theta_{\nu_h},\,\phi\in\Phi_{\nu_h}} \widehat{\mathcal{R}}_{h,2}(\theta,\phi) + \xi_{h,2,\nu_h}, \qquad \xi_{h,2,\nu_h}\ge 0, \qquad \hat{h} := h_{\hat{\theta}^{(h)},\hat{\phi}^{(h)}}.$$

Assumption H.6 (Outcome-side boundedness)**.** *There are finite constants $B_Y, B_\phi, B_\mu, B_\theta$, uniform in ν_h, such that*

$$\|Y\|_2 \le B_Y \quad a.s., \qquad \sup_{\phi\in\Phi_{\nu_h}}\sup_{w\in\mathcal{W}} \|\phi(w)\|_2 \le B_\phi, \qquad \sup_{\eta\in\mathcal{M}_{\nu_h}}\sup_{q_h\in\mathcal{Q}_h} \|\eta(q_h)\|_2 \le B_\mu,$$

and

$$\sup_{\theta\in\Theta_{\nu_h}}\sup_{(a,x)\in\mathcal{A}\times\mathcal{X}} \|\theta(a,x)\|_{\mathrm{op}} \le B_\theta.$$

Set

$$M_{h,1} := (B_\phi + B_\mu)^2, \qquad M_{h,2} := (B_Y + B_\theta B_\mu)^2.$$

Define the approximation errors

$$\kappa_{h,1,\nu_h} := \sup_{\phi\in\Phi_{\nu_h}} \inf_{\eta\in\mathcal{M}_{\nu_h}} \|\eta - \mu_\phi^{(h)}\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_{\nu_h}})},$$
$$\kappa_{h,2,\nu_h} := \inf_{\theta\in\Theta_{\nu_h},\,\phi\in\Phi_{\nu_h}} \|m_0 - m_{\theta,\phi}^{(h)}\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})}.$$

The two scalar loss classes are

$$\mathcal{L}_{h,1,\nu_h} := \left\{(q_h,w)\mapsto \|\phi(w)-\eta(q_h)\|_2^2 : \phi\in\Phi_{\nu_h},\ \eta\in\mathcal{M}_{\nu_h}\right\},$$
$$\mathcal{L}_{h,2,\nu_h} := \left\{(y,q_h)\mapsto \|y-\langle\theta(a,x),\eta(q_h)\rangle\|_2^2 : \theta\in\Theta_{\nu_h},\ \eta\in\mathcal{M}_{\nu_h}\right\}.$$

Their envelopes are $M_{h,1}$ and $M_{h,2}$, respectively.

To analyze the first- and second-stage errors of outcome bridge network, we will apply Proposition H.3. To this end, define

$$\Delta_{h,1,\nu_h}(\delta) := 2\widehat{\mathfrak{R}}_{D_1^{(h)}}(\mathcal{L}_{h,1,\nu_h}) + 3M_{h,1}\sqrt{\frac{\log(2/\delta)}{2n_h}},$$
$$\Delta_{h,2,\nu_h}(\delta) := 2\widehat{\mathfrak{R}}_{D_2^{(h)}}(\mathcal{L}_{h,2,\nu_h}) + 3M_{h,2}\sqrt{\frac{\log(2/\delta)}{2m_h}}.$$

Lemma H.7 (Uniform OUTCOMENET first-stage error)**.** *Let*

$$\epsilon_{h,1,\nu_h} := \sup_{\phi \in \Phi_{\nu_h}} \|\widehat{\mu}_\phi^{(h)} - \mu_\phi^{(h)}\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_{\nu_h}})}.$$

With probability at least $1 - 2\delta$*,*

$$\epsilon_{h,1,\nu_h} \le \kappa_{h,1,\nu_h} + 2\Delta_{h,1,\nu_h}(\delta) + \xi_{h,1,\nu_h}.$$

Proof. For fixed ϕ, the vector-valued projection identity gives

$$\mathcal{R}_{h,1}(\phi,\eta) - \mathcal{R}_{h,1}(\phi,\mu_\phi^{(h)}) = \|\eta - \mu_\phi^{(h)}\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_{\nu_h}})}.$$

Indeed,

$$\begin{aligned}\mathcal{R}_{h,1}(\phi,\eta) &= \mathbb{E}\|\phi(W) - \mu_\phi^{(h)}(Q) + \mu_\phi^{(h)}(Q) - \eta(Q)\|_2^2 \\ &= \mathcal{R}_{h,1}(\phi,\eta) + \|\eta - \mu_\phi^{(h)}\|^2_{L^2(P_Q;\mathbb{R}^{d_{\nu_h}})} + 2\mathbb{E}\langle \phi(W) - \mu_\phi^{(h)}(Q), \mu_\phi^{(h)}(Q) - \eta(Q)\rangle\end{aligned}$$

The cross term vanishes after conditioning on Q_h. Define

$$\Delta_{1,h,\nu_h} := \sup_{\phi \in \Phi_{\nu_h},\, \eta \in \mathcal{M}_{\nu_h}} \left| \mathcal{R}_{h,1}(\phi,\eta) - \widehat{\mathcal{R}}_{h,1}(\phi,\eta) \right|.$$

By approximate empirical risk minimization, $\widehat{\mathcal{R}}_{h,1}(\phi,\widehat{\mu}_\phi^{(h)}) \le \widehat{\mathcal{R}}_{h,1}(\phi,\eta) + \xi_{h,1,\nu_h}, \quad \forall \eta \in \mathcal{M}_{\nu_h}.$

Therefore, for every $\eta \in \mathcal{M}_{\nu_h}$,

$$\begin{aligned}\mathcal{R}_{h,1}(\phi,\widehat{\mu}_\phi^{(h)}) - \mathcal{R}_{h,1}(\phi,\mu_\phi^{(h)}) &\le \widehat{\mathcal{R}}_{h,1}(\phi,\widehat{\mu}_\phi^{(h)}) - \mathcal{R}_{h,1}(\phi,\mu_\phi^{(h)}) + \Delta_{1,\nu_h} \\ &\le \widehat{\mathcal{R}}_{h,1}(\phi,\eta) - \mathcal{R}_{h,1}(\phi,\mu_\phi^{(h)}) + \Delta_{1,\nu_h} + \xi_{h,1,\nu_h} \\ &\le \mathcal{R}_{h,1}(\phi,\eta) - \mathcal{R}_{h,1}(\phi,\mu_\phi^{(h)}) + 2\Delta_{1,\nu_h} + \xi_{h,1,\nu_h}.\end{aligned}$$

Using the projection identity, $\|\widehat{\mu}_\phi^{(h)} - \mu_\phi^{(h)}\|^2_{L^2(\mathbb{P}_Q;\mathbb{R}^{d_{\nu_h}})} \le \|\eta - \mu_\phi^{(h)}\|^2_{L^2(\mathbb{P}_Q;\mathbb{R}^{d_{\nu_h}})} + 2\Delta_{1,\nu_h} + \xi_{h,1,\nu_h}.$ Taking the infimum over $\eta \in \mathcal{M}_{\nu_h}$ and then the supremum over $\phi \in \Phi_{\nu_h}$ yields

$$\epsilon_{1,h,\nu_h} \le \kappa_{1,h,\nu_h} + 2\Delta_{1,h,\nu_h} + \xi_{h,1,\nu_h}.$$

The gap enters once, additively, because the optimization inequality is applied a single time in the comparison above and the criterion scale is that of the squared projected error.

Finally, applying Proposition H.3 to the bounded loss class $\mathcal{L}_{1,h,\nu_h}$ gives

$$\Delta_{1,h,\nu_h} \le \Delta_{1,h,\nu_h}(\delta)$$

with probability at least $1 - 2\delta$.

□

Theorem H.8 (Projected OUTCOMENET error bound)**.** *Under Assumption H.6, let both stages use measurable approximate empirical minimizers with gaps* $\xi_{h,1,\nu_h}, \xi_{h,2,\nu_h} \ge 0$*. For every* $\delta \in (0, 1/4)$*, with probability at least* $1 - 4\delta$*,*

$$\begin{aligned}\|T_h\hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})} \le\ & 4\kappa_{h,2,\nu_h} + 6B_\theta^2\kappa_{h,1,\nu_h} + 12B_\theta^2\Delta_{h,1,\nu_h}(\delta) + 4\Delta_{h,2,\nu_h}(\delta) \\ & + 6B_\theta^2\xi_{h,1,\nu_h} + 2\xi_{h,2,\nu_h}.\end{aligned}$$

Setting $\xi_{h,1,\nu_h} = \xi_{h,2,\nu_h} = 0$ *recovers the exact-minimization bound. If a square-integrable outcome bridge* h_0 *satisfies* $T_h h_0 = m_0$*, the same bound controls*

$$\|\hat{h} - h_0\|^2_{T_h} = \|T_h\hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})}.$$

Proof. Conditional on $D_1^{(h)}$, every $\widehat{\mu}_\phi^{(h)}$ is fixed and belongs to $\mathcal{M}_{\nu_h}$. Hence the conditional plug-in loss class is a subclass of $\mathcal{L}_{h,2,\nu_h}$. Define

$$\mathcal{R}^{\mathrm{plug}}_{h,2}(\theta,\phi) := \mathbb{E}\|Y - \widehat{m}^{(h)}_{\theta,\phi}(Q_h)\|_2^2.$$

This expectation is over an independent second-stage draw (Y, Q_h), conditional on $D_1^{(h)}$. Since $m_0(Q_h) = \mathbb{E}[Y \mid Q_h]$,

$$\mathcal{R}_{h,2}^{\mathrm{plug}}(\theta, \phi) = \mathbb{E}\|Y - m_0(Q_h)\|_2^2 + \|\widehat{m}_{\theta,\phi}^{(h)} - m_0\|_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})}^2.$$

Let

$$\Delta_{2,h,\nu_h} := \sup_{\theta\in\Theta_{\nu_h},\, \eta\in\mathcal{M}_{\nu_h}} \Bigg| \mathbb{E}\left[\{Y - \langle\theta(A,X), \eta(Q_h)\rangle\}^2\right] - \frac{1}{m_h}\sum_{i\in\mathcal{I}_2^{(h)}} \{y_i - \langle\theta(a_i,x_i), \eta(q_{h,i})\rangle\}^2 \Bigg|.$$

Because $\widehat{\mu}_\phi^{(h)} \in \mathcal{M}_{\nu_h}$, this uniform deviation controls the plug-in class:

$$\sup_{\theta,\phi} \left| \mathcal{R}_{h,2}^{\mathrm{plug}}(\theta,\phi) - \widehat{\mathcal{R}}_{h,2}^{\mathrm{plug}}(\theta,\phi) \right| \le \Delta_{2,\nu_h}.$$

By approximate empirical risk minimization, $\widehat{\mathcal{R}}_{h,2}^{\mathrm{plug}}(\hat\theta,\hat\phi) \le \widehat{\mathcal{R}}_{h,2}^{\mathrm{plug}}(\theta,\phi) + \xi_{h,2,\nu_h} \ \ \forall(\theta,\phi) \in \Theta_{\nu_h} \times \Phi_{\nu_h}$. Since the plug-in excess risk equals the squared projected error, adding and subtracting the uniform deviation twice and applying this inequality once gives

$$\|\widehat{m}_{\hat\theta^{(h)},\hat\phi^{(h)}}^{(h)} - m_0\|_{L^2}^2 \le \inf_{\theta,\phi} \|\widehat{m}_{\theta,\phi}^{(h)} - m_0\|_{L^2}^2 + 2\Delta_{h,2,\nu_h} + \xi_{h,2,\nu_h} \tag{70}$$

The operator-norm bound gives, uniformly in (θ,ϕ),

$$\|\widehat{m}_{\theta,\phi}^{(h)} - m_{\theta,\phi}^{(h)}\|_{L^2}^2 = \mathbb{E}\left\|\theta(A,X)\{\widehat{\mu}_\phi^{(h)}(Q_h) - \mu_\phi^{(h)}(Q_h)\}\right\|_2^2 \le B_\theta^2 \epsilon_{h,1,\nu_h}.$$

Thus

$$\inf_{\theta,\phi} \|\widehat{m}_{\theta,\phi}^{(h)} - m_0\|_{L^2}^2 \le 2\kappa_{h,2,\nu_h} + 2B_\theta^2\epsilon_{h,1,\nu_h}.$$

Combining this inequality with the preceding ERM bound gives

$$\|\widehat{m}_{\hat\theta^{(h)},\hat\phi^{(h)}}^{(h)} - m_0\|_{L^2}^2 \le 2\kappa_{h,2,\nu_h} + 2B_\theta^2\epsilon_{h,1,\nu_h} + 2\Delta_{h,2,\nu_h} + \xi_{h,2,\nu_h}.$$

Finally, $T_h\hat h = m_{\hat\theta^{(h)},\hat\phi^{(h)}}^{(h)}$. Hence, using $\|u+v\|^2 \le 2\|u\|^2 + 2\|v\|^2$, which doubles every term of the previous display and contributes a further $2B_\theta^2\epsilon_{h,1,\nu_h}$,

$$\begin{aligned}\|T_h\hat h - m_0\|_{L^2}^2 &\le 2\|\widehat{m}_{\hat\theta^{(h)},\hat\phi^{(h)}}^{(h)} - m_0\|_{L^2}^2 + 2B_\theta^2\epsilon_{h,1,\nu_h} \\ &\le 4\kappa_{h,2,\nu_h} + 6B_\theta^2\epsilon_{h,1,\nu_h} + 4\Delta_{h,2,\nu_h} + 2\xi_{h,2,\nu_h}.\end{aligned}$$

By Lemma H.7, with probability at least $1-2\delta$

$$\epsilon_{h,1,\nu_h} \le \kappa_{1,h,\nu_h} + 2\Delta_{1,h,\nu_h}(\delta) + \xi_{h,1,\nu_h},$$

and substituting this into the coefficient $6B_\theta^2\epsilon_{h,1,\nu_h}$ produces the term $6B_\theta^2\xi_{h,1,\nu_h}$.

By Proposition H.3 to the bounded loss class, with probability $1-2\delta$, $\mathcal{L}_{2,h,\nu_h}$ gives $\Delta_{2,h,\nu_h} \le \Delta_{2,h,\nu_h}(\delta)$. Hence, the union bond gives the stated result with probability at least $1-4\delta$.

Finally, if h_0 is a valid outcome bridge, then $T_h h_0 = m_0$. Therefore,

$$\|\hat h - h_0\|_{T_h}^2 = \|T_h(\hat h - h_0)\|_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_{\nu_h}})}^2 = \|T_h\hat h - m_0\|_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_{\nu_h}})}^2$$

□

Corollary H.9 (Projected OUTCOMENET consistency)**.** *For $\delta \in (0, 1/4)$, define*

$$\rho_{h,\nu_h}(\delta) := 4\kappa_{h,2,\nu_h} + 6B_\theta^2\kappa_{h,1,\nu_h} + 12B_\theta^2\Delta_{h,1,\nu_h}(\delta) + 4\Delta_{h,2,\nu_h}(\delta).$$

Define the optimization-aware projected rate

$$\tilde{\rho}_{h,\nu_h}(\delta) := \rho_{h,\nu_h}(\delta) + 6B_\theta^2 \xi_{h,1,\nu_h} + 2\xi_{h,2,\nu_h}.$$

Then, with probability at least $1 - 4\delta$*,*

$$\|T_h \hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})} \le \tilde{\rho}_{h,\nu_h}(\delta).$$

Let $\delta_{h,\nu_h} \downarrow 0$ *satisfy* $\log(1/\delta_{h,\nu_h})/\min\{n_h, m_h\} \to 0$*, and abbreviate* $\rho_{h,\nu_h} := \rho_{h,\nu_h}(\delta_{h,\nu_h})$ *and* $\tilde{\rho}_{h,\nu_h} := \tilde{\rho}_{h,\nu_h}(\delta_{h,\nu_h})$*. If*

$$\kappa_{h,1,\nu_h} \to 0, \qquad \kappa_{h,2,\nu_h} \to 0, \qquad \widehat{\mathfrak{R}}_{D_j^{(h)}}(\mathcal{L}_{h,j,\nu_h}) \xrightarrow{p} 0, \quad j \in \{1,2\},$$

$$\xi_{h,1,\nu_h} \xrightarrow{p} 0, \qquad \xi_{h,2,\nu_h} \xrightarrow{p} 0,$$

as $n_h, m_h \to \infty$*, then* $\tilde{\rho}_{h,\nu_h} \xrightarrow{p} 0$ *and*

$$\|T_h \hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})} = O_p(\tilde{\rho}_{h,\nu_h}), \qquad \|T_h \hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})} \xrightarrow{p} 0.$$

If $T_h h_0 = m_0$*, then* $\|\hat{h} - h_0\|^2_{T_h} \to 0$ *in probability.*

Remark H.10 (Projected rather than strong bridge recovery)**.** *The theorem controls the conditional-moment error* $T_h \hat{h} - m_0$*. It does not imply*

$$\|\hat{h} - h_0\|_{L^2(\mathbb{P}_{A,X,W};\mathbb{R}^{d_Y})} \to 0.$$

Indeed, since h_0 *is the solution to an ill-posed inverse problem, obtaining strong norm convergence for* $\hat{h}$ *requires controlling sieve measure of ill-posedness, which generally diverges as sieve dimension grows (Blundell et al., 2007; Chen & Reiss, 2011; Chen & Christensen, 2018; Meunier et al., 2024; Shen et al., 2025; Kim et al., 2025). In the severely ill-posed case, strong-norm convergence can be arbitrarily slow. The result below avoids this difficulty for the outcome-bridge plug-in estimator: using the dual treatment bridge, the dose-response functional is well-posed with respect to the weak residual* $T_h \hat{h} - m_0$*, so projected convergence is sufficient for consistency (Deaner, 2023). This weak-norm control is then used later in the doubly robust analysis.*

Remark H.11 (Outcome dimension)**.** *For fixed* d_Y*, no coordinatewise union bound is needed. If* d_Y *grows, the dependence of* B_Y*,* B_θ*, and* $\widehat{\mathfrak{R}}(\mathcal{L}_{h,2,\nu_h})$ *on* d_Y *must be tracked.*

H.3 From projected outcome-bridge consistency to dose-response consistency

We now transfer the projected conditional-moment rate for $\hat{h}$ to the dose-response curve. This is the ATE analogue of the well-posedness inequality in Deaner (2023, Theorem 4.1(i)). Corollary H.9 gives

$$\|T_h \hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})} = O_p(\tilde{\rho}_{h,\nu_h}).$$

The result below shows that the dose-response functional is continuous in this projected seminorm; strong recovery of h_0 is not required.

For a candidate bridge h, define the population plug-in dose-response curve

$$\bar{f}_h(a) := \mathbb{E}[h(a, X, W)] \in \mathbb{R}^{d_Y}.$$

Our goal is to control $\|\bar{f}_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A)}$ Assume the treatment-bridge identification route in Theorem 3.2. In particular, suppose an ATE treatment bridge φ_0^{ATE} exists and satisfies

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a, X, W] = r_{0,\mathrm{ATE}}(a, X, W).$$

This equality holds for $\mathbb{P}_A$-almost every a and $\mathbb{P}_{X,W|A=a}$-almost every (x, w). For $\mathbb{P}_A$-almost every a, set

$$c_{\varphi,\mathrm{ATE}}(a) := \left(\mathbb{E}[\{\varphi_0^{\mathrm{ATE}}(a, X, Z)\}^2 \mid A = a]\right)^{1/2}, \qquad C_{\varphi,\mathrm{ATE}} := \operatorname*{ess\,sup}_{a \sim \mathbb{P}_A} c_{\varphi,\mathrm{ATE}}(a),$$

and assume $C_{\varphi,\mathrm{ATE}} < \infty$.

Lemma H.12 (Projected residual controls dose-response error)**.** *For any $h \in \mathcal{H}_{\nu_h}$ and $\mathbb{P}_A$-almost every a,*

$$\left\|f_{\mathrm{ATE}}(a) - \bar{f}_h(a)\right\|_2 \le c_{\varphi,\mathrm{ATE}}(a)\, \|\mathbb{E}[Y - h(a,X,W) \mid A=a, X, Z]\|_{L^2(\mathbb{P}_{X,Z\mid A=a};\mathbb{R}^{d_Y})}.$$

Proof. By treatment-bridge identification,

$$f_{\mathrm{ATE}}(a) = \mathbb{E}[Y\varphi_0^{\mathrm{ATE}}(a,X,Z) \mid A=a].$$

By iterated expectations and the treatment-bridge equation,

$$\begin{aligned}\mathbb{E}[h(a,X,W)\varphi_0^{\mathrm{ATE}}(a,X,Z) \mid A=a] &= \mathbb{E}[h(a,X,W) r_{0,\mathrm{ATE}}(a,X,W) \mid A=a] \\ &= \mathbb{E}[h(a,X,W)] = \bar{f}_h(a).\end{aligned}$$

The second equality is a change of measure: disintegration of $r_{0,\mathrm{ATE}}$ gives $r_{0,\mathrm{ATE}}(a,\cdot) = \frac{d\mathbb{P}_{X,W}}{d\mathbb{P}_{X,W\mid A=a}}$ for $\mathbb{P}_A$-almost every a. Therefore,

$$f_{\mathrm{ATE}}(a) - \bar{f}_h(a) = \mathbb{E}[\varphi_0^{\mathrm{ATE}}(a,X,Z)\{Y - h(a,X,W)\} \mid A=a].$$

Conditioning on (X,Z),

$$f_{\mathrm{ATE}}(a) - \bar{f}_h(a) = \mathbb{E}\left[\varphi_0^{\mathrm{ATE}}(a,X,Z)\mathbb{E}[Y - h(a,X,W) \mid A=a,X,Z] \mid A=a\right].$$

Then Cauchy–Schwarz inequality yileds

$$\left\|f_{\mathrm{ATE}}(a) - \bar{f}_h(a)\right\|_2 \le c_{\varphi,\mathrm{ATE}}(a)\, \|\mathbb{E}[Y - h(a,X,W) \mid A=a, X, Z]\|_{L^2(\mathbb{P}_{X,Z\mid A=a};\mathbb{R}^{d_Y})}.$$

□

Theorem H.13 (Outcome-bridge dose-response consistency)**.** *Suppose*

$$\|T_h\hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})} = O_p(\tilde{\rho}_{h,\nu_h}), \tag{71}$$

and suppose the treatment-bridge conditions above hold. Then

$$\|\bar{f}_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le C^2_{\varphi,\mathrm{ATE}}\|T_h\hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})}.$$

Consequently,

$$\|\bar{f}_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p(\tilde{\rho}_{h,\nu_h}).$$

In particular, if $\tilde{\rho}_{h,\nu_h} \to 0$, then $\bar{f}_{\hat{h}} \to f_{\mathrm{ATE}}$ in $L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})$ in probability.

Proof. For $h = \hat{h}$, the previous lemma gives, for every a,

$$\left\|f_{\mathrm{ATE}}(a) - \bar{f}_{\hat{h}}(a)\right\|_2^2 \le C_\varphi(a)^2 \left\|\mathbb{E}[Y - \hat{h}(a,X,W) \mid A=a, X, Z]\right\|^2_{L^2(\mathbb{P}_{X,Z\mid A=a})}.$$

Integrating over $a \sim \mathbb{P}_A$,

$$\begin{aligned}\|\bar{f}_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A)} &= \int \left\|\bar{f}_{\hat{h}}(a) - f_{\mathrm{ATE}}(a)\right\|_2^2 d\mathbb{P}_A(a) \\ &\le \int C_\varphi(a)^2 \left[\int \left\|\mathbb{E}[Y - \hat{h}(a,X,W) \mid A=a, X=x, Z=z]\right\|_2^2 d\mathbb{P}_{X,Z\mid A=a}(x,z)\right] d\mathbb{P}_A(a) \\ &\le C_\varphi^2 \int\int \left\|\mathbb{E}[Y - \hat{h}(a,X,W) \mid A=a, X=x, Z=z]\right\|_2^2 d\mathbb{P}_{X,Z\mid A=a}(x,z) d\mathbb{P}_A(a) \\ &= C_\varphi^2 \int \left\|\mathbb{E}[Y - \hat{h}(A,X,W) \mid A, X, Z]\right\|_2^2 dP_{A,X,Z}(a,x,z).\end{aligned}$$

Since

$$\mathbb{E}[Y - \hat{h}(A,X,W) \mid Q_h] = m_0(Q_h) - T_h\hat{h}(Q_h),$$

we obtain

$$\|\bar{f}_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le C^2_{\varphi,\mathrm{ATE}}\|T_h\hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})}.$$

Equation 71 completes the proof.

□

Theorem H.14 (Empirical outcome-bridge averaging error)**.** *Let $\hat{h}$ be trained on a sample independent of an evaluation sample $\mathcal{D}_3^{(h)} = \{(x_i, w_i)\}_{i \in \mathcal{I}_3^{(h)}}$, where $(x_i, w_i) \overset{i.i.d.}{\sim} \mathbb{P}_{X,W}$ and $|\mathcal{I}_3^{(h)}| = t_h$. Define*

$$\mu_{\hat{h}}(a) := \bar{f}_{\hat{h}}(a) = \int \hat{h}(a, x, w)\, d\mathbb{P}_{X,W}(x, w), \qquad \hat{\mu}_h(a) := \frac{1}{t_h} \sum_{i \in \mathcal{I}_3^{(h)}} \hat{h}(a, x_i, w_i).$$

Under Assumption H.6,

$$\sup_{a,x,w} \|\hat{h}(a, x, w)\|_2 \le B_{h,\infty} := B_\theta B_\phi \quad \textit{almost surely}.$$

Then

$$\mathcal{E}_{\mu,h} := \|\hat{\mu}_h - \mu_{\hat{h}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p(t_h^{-1}).$$

Consequently, if $\|\mu_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p(\tilde{\rho}_{h,\nu_h})$, then

$$\|\hat{\mu}_h - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p(\tilde{\rho}_{h,\nu_h} + t_h^{-1}).$$

Proof. Condition on the trained bridge $\hat{h}$. For every fixed a, the empirical average $\hat{\mu}_h(a)$ is unbiased for $\mu_{\hat{h}}(a)$, because the evaluation sample is independent of the training sample:

$$\mathbb{E}[\hat{\mu}_h(a) \mid \hat{h}] = \int \hat{h}(a, x, w)\, d\mathbb{P}_{X,W}(x, w) = \mu_{\hat{h}}(a).$$

Set $\xi_i(a) := \hat{h}(a, x_i, w_i) - \mu_{\hat{h}}(a)$. Conditional on $\hat{h}$, the vectors $\xi_i(a)$ are independent and centered. Hence the cross inner products vanish and

$$\mathbb{E}\Big[\|\hat{\mu}_h(a) - \mu_{\hat{h}}(a)\|_2^2 \mid \hat{h}\Big] = \frac{1}{t_h^2} \sum_{i \in \mathcal{I}_3^{(h)}} \mathbb{E}[\|\xi_i(a)\|_2^2 \mid \hat{h}] \le \frac{B_{h,\infty}^2}{t_h}.$$

Here we used $\mathbb{E}[\|\xi_i(a)\|_2^2 \mid \hat{h}] = \mathbb{E}[\|\hat{h}(a, X, W)\|_2^2 \mid \hat{h}] - \|\mu_{\hat{h}}(a)\|_2^2 \le B_{h,\infty}^2$. Integrating over $a \sim \mathbb{P}_A$ gives

$$\mathbb{E}\Big[\mathcal{E}_{\mu,h} \mid \hat{h}\Big] = \int \mathbb{E}\Big[\|\hat{\mu}_h(a) - \mu_{\hat{h}}(a)\|_2^2 \mid \hat{h}\Big]\, d\mathbb{P}_A(a) \le \frac{B_{h,\infty}^2}{t_h}.$$

Therefore, by Markov's inequality, for every $\delta > 0$,

$$\mathbb{P}\left(\mathcal{E}_{\mu,h} > \frac{B_{h,\infty}^2}{\delta t_h} \,\middle|\, \hat{h}\right) \le \delta,$$

and hence $\mathcal{E}_{\mu,h} = O_p(t_h^{-1})$.

Finally, by writing $\hat{\mu}_h - f_{\mathrm{ATE}} = (\hat{\mu}_h - \mu_{\hat{h}}) + (\mu_{\hat{h}} - f_{\mathrm{ATE}})$, we obtain

$$\|\hat{\mu}_h - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le 2\|\hat{\mu}_h - \mu_{\hat{h}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} + 2\|\mu_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

Combining the empirical averaging bound with $\|\mu_{\hat{h}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p(\tilde{\rho}_{h,\nu_h})$ gives the stated rate.

□

H.4 Treatment bridge consistency for dose-response estimation

We establish projected consistency of TreatmentNet. The argument is the counterpart of the corresponding analysis of the OutcomeNet analysis in Appendix H.2, but its second stage is an order-two U-statistic.

Let $Q_\varphi := (A, X, W) \subset \mathcal{Q}_\varphi$. For a measurable $u : \mathcal{A} \times \mathcal{X} \times \mathcal{Z} \to \mathbb{R}$, define

$$(T_\varphi u)(a, x, w) := \mathbb{E}[u(a, x, Z) \mid A = a, X = x, W = w], \qquad \|u\|_{T_\varphi} := \|T_\varphi u\|_{L^2(\mathbb{P}_{Q_\varphi})}.$$

Let the independent first- and second-stage samples be $D_1^{(\varphi)} = \{(q_{\varphi,i}, z_i)\}_{i\in\mathcal{I}_1^{(\varphi)}}$, $D_2^{(\varphi)} = \{q_{\varphi,i}\}_{i\in\mathcal{I}_2^{(\varphi)}}$, with $|\mathcal{I}_1^{(\varphi)}| = n_\varphi$ and $|\mathcal{I}_2^{(\varphi)}| = m_\varphi$. The first sample is i.i.d. from $\mathbb{P}_{Q_\varphi,Z}$, the second is i.i.d. from $\mathbb{P}_{Q_\varphi}$. We consider the learned-feature bridge class. Let $\nu_\varphi = (n_\varphi, m_\varphi)$ and define

$$\Phi_{\nu_\varphi} \subset \{\phi : \mathcal{Z} \to \mathbb{R}^{d_{\nu_\varphi}}\}, \qquad \mathcal{M}_{\nu_\varphi} \subset \{\eta : \mathcal{Q}_\varphi \to \mathbb{R}^{d_{\nu_\varphi}}\},$$

for the class of learned treatment-proxy feature maps and the first-stage conditional mean regressors, respectively. We define the class of second-stage heads by

$$\Theta_{\nu_\varphi} \subset \{\theta : \mathcal{A}\times\mathcal{X} \to \mathbb{R}^{d_{\nu_\varphi}}\}.$$

The bridge class is

$$\mathcal{H}_{\nu_\varphi} := \left\{\varphi_{\theta,\phi}(a,x,z) = \langle\theta(a,x), \phi(z)\rangle : \theta\in\Theta_{\nu_\varphi},\ \phi\in\Phi_{\nu_\varphi}\right\}.$$

This compact class includes the tensor-product TREATMENTNET: its final linear layer and the second-stage (A, X)-features are absorbed into $\theta(a, x)$.

For each $\phi \in \Phi_{\nu_\varphi}$, define the conditional feature mean

$$\mu_\phi^{(\varphi)}(q_\varphi) := \mathbb{E}[\phi(Z) \mid Q_\varphi = q_\varphi].$$

Then, for $q_\varphi = (a, x, w)$,

$$(T_\varphi\varphi_{\theta,\phi})(q_\varphi) = \langle\theta(a,x), \mu_\phi^{(\varphi)}(q_\varphi)\rangle.$$

For each candidate feature map $\phi \in \Phi_{\nu_\varphi}$, the population and empirical first-stage risks are

$$\mathcal{R}_{\varphi,1}(\phi,\eta) := \mathbb{E}\|\phi(Z) - \eta(Q_\varphi)\|_2^2,$$
$$\widehat{\mathcal{R}}_{\varphi,1}(\phi,\eta) := \frac{1}{n_\varphi}\sum_{i\in\mathcal{I}_1^{(\varphi)}} \|\phi(z_i) - \eta(q_{\varphi,i})\|_2^2.$$

For every ϕ, let $\widehat{\mu}_\phi^{(\varphi)} \in \mathcal{M}_{\nu_\varphi}$ be a measurable approximate empirical minimizer, that is,

$$\widehat{\mathcal{R}}_{\varphi,1}(\phi, \widehat{\mu}_\phi^{(\varphi)}) \le \inf_{\eta\in\mathcal{M}_{\nu_\varphi}} \widehat{\mathcal{R}}_{\varphi,1}(\phi,\eta) + \xi_{\varphi,1,\nu_\varphi}, \qquad \xi_{\varphi,1,\nu_\varphi} \ge 0,$$

uniformly over $\phi \in \Phi_{\nu_\varphi}$; the exact-minimizer case is $\xi_{\varphi,1,\nu_\varphi} = 0$.

Define the population and plug-in projected bridges

$$g_{\theta,\phi}^{(\varphi)}(q_\varphi) := \langle\theta(a,x), \mu_\phi^{(\varphi)}(q_\varphi)\rangle = (T_\varphi\varphi_{\theta,\phi})(q_\varphi),$$
$$\widehat{g}_{\theta,\phi}^{(\varphi)}(q_\varphi) := \langle\theta(a,x), \widehat{\mu}_\phi^{(\varphi)}(q_\varphi)\rangle.$$

For $q = (a, x, w)$ and $q' = (a', x', w')$, define

$$q \diamond q' := (a, x', w').$$

For independent copies $Q_{\varphi,1}$ and $Q_{\varphi,2}$ of Q_φ, define, for $u \in L^2(\mathbb{P}_{Q_\varphi})$,

$$\mathcal{R}_{\varphi,2}^{\mathrm{RF}}(u) := \mathbb{E}\{u(Q_{\varphi,1})^2\} - 2\mathbb{E}\{u(Q_{\varphi,1}\diamond Q_{\varphi,2})\}.$$

By Lemma 4.1,

$$\mathcal{R}_{\varphi,2}^{\mathrm{RF}}(u) - \mathcal{R}_{\varphi,2}^{\mathrm{RF}}(r_0) = \|u - r_0\|_{L^2(\mathbb{P}_{Q_\varphi})}^2. \tag{72}$$

Conditional on $D_1^{(\varphi)}$, define the plug-in kernel

$$\hat{\ell}_{\theta,\phi}^{\mathrm{RF}}(q,q') := \{\hat{g}_{\theta,\phi}^{(\varphi)}(q)\}^2 - 2\hat{g}_{\theta,\phi}^{(\varphi)}(q\diamond q'). \tag{73}$$

The empirical second-stage objective is the complete ordered U-statistic

$$\widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\theta,\phi) := \frac{1}{m_\varphi(m_\varphi-1)}\sum_{i\neq j}\hat{\ell}^{\mathrm{RF}}_{\theta,\phi}(\tilde{q}_{\varphi,i},\tilde{q}_{\varphi,j})$$
$$= \frac{1}{m_\varphi}\sum_{i=1}^{m_\varphi}\{\hat{g}^{(\varphi)}_{\theta,\phi}(\tilde{q}_{\varphi,i})\}^2 - \frac{2}{m_\varphi(m_\varphi-1)}\sum_{i\neq j}\hat{g}^{(\varphi)}_{\theta,\phi}(\tilde{q}_{\varphi,i}\diamond\tilde{q}_{\varphi,j}).$$

Let $(\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)})\in\Theta_{\nu_\varphi}\times\Phi_{\nu_\varphi}$ be a measurable approximate empirical minimizer of the product-law objective,

$$\widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}) \le \inf_{\theta\in\Theta_{\nu_\varphi},\,\phi\in\Phi_{\nu_\varphi}}\widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\theta,\phi)+\xi_{\varphi,2,\nu_\varphi},\qquad \xi_{\varphi,2,\nu_\varphi}\ge 0,\qquad \hat{\varphi}:=\varphi_{\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}}.$$

Assumption H.15 (Treatment-side boundedness)**.** *There are finite constants B_ϕ, B_μ, B_θ, uniform in ν_φ, such that*

$$\sup_{\phi\in\Phi_{\nu_\varphi}}\sup_{z\in\mathcal{Z}}\|\phi(z)\|_2\le B_\phi,\qquad \sup_{\eta\in\mathcal{M}_{\nu_\varphi}}\sup_{q_\varphi\in\mathcal{Q}_\varphi}\|\eta(q_\varphi)\|_2\le B_\mu,$$

and

$$\sup_{\theta\in\Theta_{\nu_\varphi}}\sup_{(a,x)\in\mathcal{A}\times\mathcal{X}}\|\theta(a,x)\|_2\le B_\theta.$$

Here B_μ bounds the first-stage class $\mathcal{M}_{\nu_\varphi}$; the corresponding plug-in projected bridges are bounded by $B_\theta B_\mu$. Set

$$M_{\varphi,1} := (B_\phi+B_\mu)^2,\qquad M^{\mathrm{RF}}_{\varphi,2} := B_\theta^2 {B_\mu}^2 + 2B_\theta B_\mu.$$

Define the approximation errors

$$\kappa_{\varphi,1,\nu_\varphi} := \sup_{\phi\in\Phi_{\nu_\varphi}}\inf_{\eta\in\mathcal{M}_{\nu_\varphi}}\|\eta-\mu^{(\varphi)}_\phi\|^2_{L^2(\mathbb{P}_{Q_\varphi};\mathbb{R}^{d_{\nu_\varphi}})},$$
$$\kappa_{\varphi,2,\nu_\varphi} := \inf_{\theta\in\Theta_{\nu_\varphi},\,\phi\in\Phi_{\nu_\varphi}}\|g^{(\varphi)}_{\theta,\phi}-r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})}.$$

The scalar loss classes are

$$\mathcal{L}_{\varphi,1,\nu_\varphi} := \left\{(q_\varphi,z)\mapsto\|\phi(z)-\eta(q_\varphi)\|_2^2 : \phi\in\Phi_{\nu_\varphi},\ \eta\in\mathcal{M}_{\nu_\varphi}\right\},$$
$$\mathcal{L}^{\mathrm{RF}}_{\varphi,2,\nu_\varphi} := \left\{(q,q')\mapsto\langle\theta(a,x),\eta(q)\rangle^2 - 2\langle\theta(a,x'),\eta(q\diamond q')\rangle : \theta\in\Theta_{\nu_\varphi},\ \eta\in\mathcal{M}_{\nu_\varphi}\right\}.$$

Their envelopes are $M_{\varphi,1}$ and $M^{\mathrm{RF}}_{\varphi,2}$, respectively.

Define

$$\Delta_{\varphi,1,\nu_\varphi}(\delta) := 2\widehat{\mathfrak{R}}_{D_1^{(\varphi)}}(\mathcal{L}_{\varphi,1,\nu_\varphi}) + 3M_{\varphi,1}\sqrt{\frac{\log(2/\delta)}{2n_\varphi}},$$
$$\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) := 2\widehat{\mathfrak{R}}^{(2),U}_{D_2^{(\varphi)}}(\mathcal{L}^{\mathrm{RF}}_{\varphi,2,\nu_\varphi}) + 3M^{\mathrm{RF}}_{\varphi,2}\sqrt{\frac{8\log(2/\delta)}{m_\varphi}}.$$

The first term uses Proposition H.3. The second uses the U-process bound in Proposition H.5.

Lemma H.16 (Uniform TREATMENTNET first-stage error)**.** *Let*

$$\epsilon_{\varphi,1,\nu_\varphi} := \sup_{\phi\in\Phi_{\nu_\varphi}}\|\hat{\mu}^{(\varphi)}_\phi-\mu^{(\varphi)}_\phi\|^2_{L^2(\mathbb{P}_{Q_\varphi};\mathbb{R}^{d_{\nu_\varphi}})}.$$

With probability at least $1-2\delta$,

$$\epsilon_{\varphi,1,\nu_\varphi}\le\kappa_{\varphi,1,\nu_\varphi}+2\Delta_{\varphi,1,\nu_\varphi}(\delta)+\xi_{\varphi,1,\nu_\varphi}.$$

Proof. For fixed ϕ, the vector-valued projection identity gives

$$\mathcal{R}_{\varphi,1}(\phi,\eta)-\mathcal{R}_{\varphi,1}(\phi,\mu^{(\varphi)}_\phi) = \|\eta-\mu^{(\varphi)}_\phi\|^2_{L^2(\mathbb{P}_{Q_\varphi};\mathbb{R}^{d_{\nu_\varphi}})}.$$

By Proposition H.3, each of

$$\mathcal{R}_{\varphi,1}(\phi,\eta) \le \widehat{\mathcal{R}}_{\varphi,1}(\phi,\eta) + \Delta_{\varphi,1,\nu_\varphi}(\delta)$$

and

$$\widehat{\mathcal{R}}_{\varphi,1}(\phi,\eta) \le \mathcal{R}_{\varphi,1}(\phi,\eta) + \Delta_{\varphi,1,\nu_\varphi}(\delta)$$

holds uniformly with failure probability at most δ. On their intersection, approximate empirical optimality yields

$$\mathcal{R}_{\varphi,1}(\phi,\hat{\mu}_\phi^{(\varphi)}) \le \inf_{\eta\in\mathcal{M}_{\nu_\varphi}} \mathcal{R}_{\varphi,1}(\phi,\eta) + 2\Delta_{\varphi,1,\nu_\varphi}(\delta) + \xi_{\varphi,1,\nu_\varphi}.$$

Subtract the projection risk, take the infimum over η, and then the supremum over ϕ. □

Conditional on $D_1^{(\varphi)}$, the definitions of the population and plug-in projected bridges give, for every (θ,ϕ),

$$\hat{g}_{\theta,\phi}^{(\varphi)}(q_\varphi) - g_{\theta,\phi}^{(\varphi)}(q_\varphi) = \left\langle \theta(a,x), \hat{\mu}_\phi^{(\varphi)}(q_\varphi) - \mu_\phi^{(\varphi)}(q_\varphi) \right\rangle.$$

Therefore, Cauchy–Schwarz, Assumption H.15, and the definition of $\epsilon_{\varphi,1,\nu_\varphi}$ imply

$$\begin{aligned} &\sup_{\substack{\theta\in\Theta_{\nu_\varphi}\\ \phi\in\Phi_{\nu_\varphi}}} \|\hat{g}_{\theta,\phi}^{(\varphi)} - g_{\theta,\phi}^{(\varphi)}\|^2_{L^2(\mathbb{P}_{Q_\varphi})} \\ &\qquad \le B_\theta^2 \sup_{\phi\in\Phi_{\nu_\varphi}} \|\hat{\mu}_\phi^{(\varphi)} - \mu_\phi^{(\varphi)}\|^2_{L^2(\mathbb{P}_{Q_\varphi};\mathbb{R}^{d_{\nu_\varphi}})} = B_\theta^2 \epsilon_{\varphi,1,\nu_\varphi}. \end{aligned} \tag{74}$$

This plug-in stability bound transfers the first-stage conditional-mean error to the projected treatment bridge. It therefore connects the empirical second-stage risk based on $\hat{g}_{\theta,\phi}^{(\varphi)}$ to the population projection $g_{\theta,\phi}^{(\varphi)} = T_\varphi \varphi_{\theta,\phi}$.

Theorem H.17 (Projected TREATMENTNET error bound)**.** *Under the preceding setup and Assumption H.15, let both stages use measurable approximate empirical minimizers with gaps $\xi_{\varphi,1,\nu_\varphi}, \xi_{\varphi,2,\nu_\varphi} \ge 0$. For every $\delta \in (0,1/4)$, with probability at least $1-4\delta$,*

$$\begin{aligned} \|T_\varphi\hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})} \le\ & 4\kappa_{\varphi,2,\nu_\varphi} + 6B_\theta^2\kappa_{\varphi,1,\nu_\varphi} + 12B_\theta^2\Delta_{\varphi,1,\nu_\varphi}(\delta) + 4\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) \\ &+ 6B_\theta^2\xi_{\varphi,1,\nu_\varphi} + 2\xi_{\varphi,2,\nu_\varphi}. \end{aligned}$$

Setting both gaps to zero recovers the exact-minimization bound. If a square-integrable ATE treatment bridge φ_0^{ATE} satisfies $T_\varphi\varphi_0^{\mathrm{ATE}} = r_0$, the same bound controls

$$\|\hat{\varphi} - \varphi_0^{\mathrm{ATE}}\|^2_{T_\varphi} = \|T_\varphi\hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})}.$$

Proof. Conditional on $D_1^{(\varphi)}$, every $\hat{\mu}_\phi^{(\varphi)}$ is fixed and belongs to $\mathcal{M}_{\nu_\varphi}$. Hence the plug-in kernel class is a subclass of $\mathcal{L}^{\mathrm{RF}}_{\varphi,2,\nu_\varphi}$. Moreover, for every (θ,ϕ),

$$\mathbb{P}_{Q_\varphi}^{\otimes 2}\hat{\ell}^{\mathrm{RF}}_{\theta,\phi} = \mathcal{R}^{\mathrm{RF}}_{\varphi,2}(\hat{g}^{(\varphi)}_{\theta,\phi}), \qquad \mathbb{U}_{m_\varphi}\hat{\ell}^{\mathrm{RF}}_{\theta,\phi} = \widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\theta,\phi).$$

Proposition H.5 gives each of the two uniform inequalities

$$\begin{aligned} \mathcal{R}^{\mathrm{RF}}_{\varphi,2}(\hat{g}^{(\varphi)}_{\theta,\phi}) &\le \widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\theta,\phi) + \Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta), \\ \widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\theta,\phi) &\le \mathcal{R}^{\mathrm{RF}}_{\varphi,2}(\hat{g}^{(\varphi)}_{\theta,\phi}) + \Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) \end{aligned}$$

with conditional failure probability at most δ per direction. Indeed, on their intersection, for every $(\theta,\phi) \in \Theta_{\nu_\varphi} \times \Phi_{\nu_\varphi}$,

$$\begin{aligned} \mathcal{R}^{\mathrm{RF}}_{\varphi,2}(\hat{g}^{(\varphi)}_{\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}}) &\le \widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}) + \Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) \\ &\le \widehat{\mathcal{R}}^{\mathrm{RF},U}_{\varphi,2}(\theta,\phi) + \Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) + \xi_{\varphi,2,\nu_\varphi} \\ &\le \mathcal{R}^{\mathrm{RF}}_{\varphi,2}(\hat{g}^{(\varphi)}_{\theta,\phi}) + 2\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) + \xi_{\varphi,2,\nu_\varphi}. \end{aligned}$$

The middle inequality follows from approximate empirical optimality, which inserts $\xi_{\varphi,2,\nu_\varphi}$. Subtracting $\mathcal{R}^{\mathrm{RF}}_{\varphi,2}(r_0)$, applying Equation 72 to both sides, and taking the infimum over (θ,ϕ) give

$$\|\hat{g}^{(\varphi)}_{\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}} - r_0\|^2_{L^2} \le \inf_{\theta,\phi} \|\hat{g}^{(\varphi)}_{\theta,\phi} - r_0\|^2_{L^2} + 2\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) + \xi_{\varphi,2,\nu_\varphi}.$$

Here and below, the displayed L^2 norm is over $\mathbb{P}_{Q_\varphi}$. By Equation 74,

$$\begin{aligned}\inf_{\theta,\phi} \|\hat{g}^{(\varphi)}_{\theta,\phi} - r_0\|^2_{L^2} &\le 2\inf_{\theta,\phi} \|g^{(\varphi)}_{\theta,\phi} - r_0\|^2_{L^2} + 2B^2_\theta \epsilon_{\varphi,1,\nu_\varphi} \\ &= 2\kappa_{\varphi,2,\nu_\varphi} + 2B^2_\theta \epsilon_{\varphi,1,\nu_\varphi}.\end{aligned}$$

Therefore,

$$\|\hat{g}^{(\varphi)}_{\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}} - r_0\|^2_{L^2} \le 2\kappa_{\varphi,2,\nu_\varphi} + 2B^2_\theta \epsilon_{\varphi,1,\nu_\varphi} + 2\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) + \xi_{\varphi,2,\nu_\varphi}.$$

Finally, $T_\varphi\hat{\varphi} = g^{(\varphi)}_{\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}}$. A second application of Equation 74 gives

$$\begin{aligned}\|T_\varphi\hat{\varphi} - r_0\|^2_{L^2} &\le 2\|\hat{g}^{(\varphi)}_{\hat{\theta}^{(\varphi)},\hat{\phi}^{(\varphi)}} - r_0\|^2_{L^2} + 2B^2_\theta \epsilon_{\varphi,1,\nu_\varphi} \\ &\le 4\kappa_{\varphi,2,\nu_\varphi} + 6B^2_\theta \epsilon_{\varphi,1,\nu_\varphi} + 4\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) + 2\xi_{\varphi,2,\nu_\varphi}.\end{aligned}$$

The factor 2 on $\xi_{\varphi,2,\nu_\varphi}$ comes from the $\|u+v\|^2 \le 2\|u\|^2 + 2\|v\|^2$ step above. Apply Lemma H.16 , whose bound contributes $6B^2_\theta \xi_{\varphi,1,\nu_\varphi}$ through the coefficient $6B^2_\theta \epsilon_{\varphi,1,\nu_\varphi}$. Its two first-stage events and the two conditional U-process events fail with total probability at most 4δ. □

Corollary H.18 (Projected TREATMENTNET consistency)**.** *For $\delta \in (0, 1/4)$, define*

$$\rho_{\varphi,\nu_\varphi}(\delta) := 4\kappa_{\varphi,2,\nu_\varphi} + 6B^2_\theta \kappa_{\varphi,1,\nu_\varphi} + 12B^2_\theta \Delta_{\varphi,1,\nu_\varphi}(\delta) + 4\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta).$$

Define the optimization-aware projected rate

$$\tilde{\rho}_{\varphi,\nu_\varphi}(\delta) := \rho_{\varphi,\nu_\varphi}(\delta) + 6B^2_\theta \xi_{\varphi,1,\nu_\varphi} + 2\xi_{\varphi,2,\nu_\varphi}.$$

Then, with probability at least $1 - 4\delta$,

$$\|T_\varphi\hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})} \le \tilde{\rho}_{\varphi,\nu_\varphi}(\delta).$$

Let $\delta_{\varphi,\nu_\varphi} \downarrow 0$ satisfy $\log(1/\delta_{\varphi,\nu_\varphi})/\min\{n_\varphi, m_\varphi\} \to 0$, and abbreviate $\rho_{\varphi,\nu_\varphi} := \rho_{\varphi,\nu_\varphi}(\delta_{\varphi,\nu_\varphi})$ and $\tilde{\rho}_{\varphi,\nu_\varphi} := \tilde{\rho}_{\varphi,\nu_\varphi}(\delta_{\varphi,\nu_\varphi})$. If

$$\kappa_{\varphi,1,\nu_\varphi} \to 0, \qquad \kappa_{\varphi,2,\nu_\varphi} \to 0, \qquad \widehat{\mathfrak{R}}_{D_1^{(\varphi)}}(\mathcal{L}_{\varphi,1,\nu_\varphi}) \xrightarrow{p} 0, \qquad \widehat{\mathfrak{R}}^{(2),U}_{D_2^{(\varphi)}}(\mathcal{L}^{\mathrm{RF}}_{\varphi,2,\nu_\varphi}) \xrightarrow{p} 0,$$

$$\xi_{\varphi,1,\nu_\varphi} \xrightarrow{p} 0, \qquad \xi_{\varphi,2,\nu_\varphi} \xrightarrow{p} 0,$$

as $n_\varphi, m_\varphi \to \infty$, then $\tilde{\rho}_{\varphi,\nu_\varphi} \xrightarrow{p} 0$ and

$$\|T_\varphi\hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})} = O_p(\tilde{\rho}_{\varphi,\nu_\varphi}), \qquad \|T_\varphi\hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})} \xrightarrow{p} 0.$$

If $T_\varphi\varphi_0^{\mathrm{ATE}} = r_0$, then $\|\hat{\varphi} - \varphi_0^{\mathrm{ATE}}\|^2_{T_\varphi} \to 0$ in probability.

Remark H.19 (Projected rather than strong bridge recovery)**.** *The theorem controls the projected treatment-bridge error $T_\varphi\hat{\varphi} - r_0$. It does not imply*

$$\|\hat{\varphi} - \varphi_0^{\mathrm{ATE}}\|_{L^2(\mathbb{P}_{A,X,Z})} \to 0.$$

Strong bridge recovery requires an additional inverse-stability condition. The next subsection uses only the projected error to establish dose-response consistency.

H.5 DOSE-RESPONSE CONSISTENCY OF TREATMENTNET

We now convert the projected treatment-bridge result into consistency of the ATE dose-response curve. This step does not require strong convergence of $\hat{\varphi}$ in $L^2(\mathbb{P}_{A,X,Z})$. Instead, it uses the outcome bridge as the dual object and controls the population dose-response error through the projected treatment-bridge residual.

Let $Y \in \mathbb{R}^{d_Y}$. The treatment-bridge representation of the ATE dose-response curve is

$$f_{\mathrm{ATE}}(a) = \mathbb{E}[Y\varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a] \in \mathbb{R}^{d_Y},$$

where $T_\varphi \varphi_0^{\mathrm{ATE}} = r_0$. For any scalar candidate treatment bridge u, define

$$\bar{f}_u^{(\varphi)}(a) := \mathbb{E}[Yu(a, X, Z) \mid A = a] \in \mathbb{R}^{d_Y}.$$

The learned-bridge population curve is obtained by setting $u = \hat{\varphi}$.

Assumption H.20 (Dual outcome bridge and boundedness)**.** *Assume that:*

- *There exists an $\mathbb{R}^{d_Y}$-valued outcome bridge h_0 satisfying*
$$\mathbb{E}[Y - h_0(a, X, W) \mid A = a, X, Z] = 0_{\mathbb{R}^{d_Y}}.$$
- *For $\mathbb{P}_A$-almost every a, set*
$$c_h(a) := \big(\mathbb{E}[\|h_0(a, X, W)\|_2^2 \mid A = a]\big)^{1/2}, \qquad C_h := \operatorname*{ess\,sup}_{a \sim \mathbb{P}_A} c_h(a),$$
and assume $C_h < \infty$.

Lemma H.21 (Projected treatment residual controls the population curve)**.** *For any square-integrable scalar candidate treatment bridge $u(a, X, Z)$ for which $\bar{f}_u^{(\varphi)}$ exists, and for $\mathbb{P}_A$-almost every a,*

$$\left\| \bar{f}_u^{(\varphi)}(a) - f_{\mathrm{ATE}}(a) \right\|_2 \le c_h(a) \, \|(T_\varphi u)(a, X, W) - r_0(a, X, W)\|_{L^2(\mathbb{P}_{X,W \mid A=a})},$$

Consequently,

$$\|\bar{f}_u^{(\varphi)} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A; \mathbb{R}^{d_Y})} \le C_h^2 \|T_\varphi u - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})}.$$

Proof. For any u,

$$\bar{f}_u^{(\varphi)}(a) - f_{\mathrm{ATE}}(a) = \mathbb{E}[Y\{u(a, X, Z) - \varphi_0^{\mathrm{ATE}}(a, X, Z)\} \mid A = a].$$

Since $u(a, X, Z) - \varphi_0^{\mathrm{ATE}}(a, X, Z)$ is scalar and measurable with respect to (X, Z), the outcome-bridge equation gives

$$\mathbb{E}[\{Y - h_0(a, X, W)\}\{u(a, X, Z) - \varphi_0^{\mathrm{ATE}}(a, X, Z)\} \mid A = a] = 0_{\mathbb{R}^{d_Y}}.$$

Therefore,

$$\bar{f}_u^{(\varphi)}(a) - f_{\mathrm{ATE}}(a) = \mathbb{E}[h_0(a, X, W)\{u(a, X, Z) - \varphi_0^{\mathrm{ATE}}(a, X, Z)\} \mid A = a].$$

Conditioning on (X, W),

$$\bar{f}_u^{(\varphi)}(a) - f_{\mathrm{ATE}}(a) = \mathbb{E}\big[h_0(a, X, W)\mathbb{E}[u(a, X, Z) - \varphi_0^{\mathrm{ATE}}(a, X, Z) \mid A = a, X, W] \mid A = a\big].$$

Since $T_\varphi \varphi_0^{\mathrm{ATE}} = r_0$, the inner conditional expectation equals

$$(T_\varphi u)(a, X, W) - r_0(a, X, W).$$

Cauchy–Schwarz gives

$$\left\| \bar{f}_u^{(\varphi)}(a) - f_{\mathrm{ATE}}(a) \right\|_2^2 \le c_h(a)^2 \, \|(T_\varphi u)(a, X, W) - r_0(a, X, W)\|^2_{L^2(\mathbb{P}_{X,W \mid A=a})}.$$

Integrating over $a \sim \mathbb{P}_A$ and using

$$c_h(a)^2 \le C_h^2$$

almost surely proves the result. □

Conditional on the trained bridge $\hat{\varphi}$, let $D_3^{(\varphi)} := \{(y_i, a_i, x_i, z_i)\}_{i \in \mathcal{I}_3^{(\varphi)}}$ be an independent i.i.d. sample from $\mathbb{P}_{Y,A,X,Z}$ with $|\mathcal{I}_3^{(\varphi)}|$. Set

$$\Gamma_\varphi := Y\hat{\varphi}(A, X, Z) \in \mathbb{R}^{d_Y}, \qquad \mathcal{F}_{t_\varphi}^{(\varphi)} \subset \{f : \mathcal{A} \to \mathbb{R}^{d_Y}\}.$$

Define

$$\widehat{\mathcal{R}}_{\varphi,3}(f) := \frac{1}{t_\varphi} \sum_{i \in \mathcal{I}_3^{(\varphi)}} \left\| y_i^{(\varphi)} - f(a_i) \right\|_2^2,$$

$$\hat{f}_{\mathrm{ATE}}^{(\varphi)} \in \mathcal{F}_{t_\varphi}^{(\varphi)} \text{ with } \widehat{\mathcal{R}}_{\varphi,3}(\hat{f}_{\mathrm{ATE}}^{(\varphi)}) \le \inf_{f \in \mathcal{F}_{t_\varphi}^{(\varphi)}} \widehat{\mathcal{R}}_{\varphi,3}(f) + \xi_{\varphi,3,t_\varphi}, \quad \xi_{\varphi,3,t_\varphi} \ge 0,$$

$$\kappa_{\varphi,3,t_\varphi} := \inf_{f \in \mathcal{F}_{t_\varphi}^{(\varphi)}} \|f - \bar{f}_{\hat{\varphi}}^{(\varphi)}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

Here $y_i^{(\varphi)} := y_i\hat{\varphi}(a_i, x_i, z_i)$, as in the algorithmic description, and for the neural class $\hat{f}_{\mathrm{ATE}}^{(\varphi)}(a) = k^\varphi(a; \hat{\theta}_3^\varphi)$. The fitted function is assumed measurable and satisfies the preceding approximate empirical minimization condition. For a continuous treatment, this is the third-stage response regression. For a finite treatment space, the saturated choice of $\mathcal{F}_{t_\varphi}^{(\varphi)}$ yields the within-treatment sample averages on nonempty treatment cells, so no response network is required.

Conditional on $\hat{\varphi}$, define the third-stage loss class

$$\mathcal{L}_{\varphi,3,t_\varphi}(\hat{\varphi}) := \left\{ (y, a, x, z) \mapsto \|y\hat{\varphi}(a, x, z) - f(a)\|_2^2 : f \in \mathcal{F}_{t_\varphi}^{(\varphi)} \right\}.$$

We assume this conditional loss class is measurable and separable.

Assumption H.22 (Third-stage boundedness)**.** *Conditional on the trained treatment bridge $\hat{\varphi}$, assume that:*

- *The third-stage loss class is uniformly bounded:*

$$0 \le \ell(y, a, x, z) \le M_{\varphi,3} \quad \text{for every } \ell \in \mathcal{L}_{\varphi,3,t_\varphi}(\hat{\varphi}) \text{ and } (y, a, x, z),$$

  for a common deterministic constant $M_{\varphi,3} < \infty$, uniform in t_φ.

- *The third-stage sample $D_3^{(\varphi)}$ is independent of the samples used to train $\hat{\varphi}$.*

Define the third-stage empirical process error

$$\Delta_{\varphi,3,t_\varphi}(\delta) := 2\widehat{\mathfrak{R}}_{D_3^{(\varphi)}}(\mathcal{L}_{\varphi,3,t_\varphi}(\hat{\varphi})) + 3M_{\varphi,3}\sqrt{\frac{\log(2/\delta)}{2t_\varphi}}.$$

Lemma H.23 (Third-stage Rademacher bound)**.** *Under Assumption H.22, conditional on the trained treatment bridge $\hat{\varphi}$, with probability at least $1 - 2\delta$,*

$$\|\hat{f}_{\mathrm{ATE}}^{(\varphi)} - \bar{f}_{\hat{\varphi}}^{(\varphi)}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le \kappa_{\varphi,3,t_\varphi} + 2\Delta_{\varphi,3,t_\varphi}(\delta) + \xi_{\varphi,3,t_\varphi}.$$

Proof. Conditional on $\hat{\varphi}$, this is ordinary least-squares regression with response

$$\Gamma_\varphi := Y\hat{\varphi}(A, X, Z) \in \mathbb{R}^{d_Y}$$

and covariate A. Since

$$\bar{f}_{\hat{\varphi}}^{(\varphi)}(A) = \mathbb{E}[\Gamma_\varphi \mid A, \hat{\varphi}],$$

the least-squares projection identity gives

$$\mathbb{E}[\|\Gamma_\varphi - f(A)\|_2^2 \mid \hat{\varphi}] - \mathbb{E}[\|\Gamma_\varphi - \bar{f}_{\hat{\varphi}}^{(\varphi)}(A)\|_2^2 \mid \hat{\varphi}] = \|f - \bar{f}_{\hat{\varphi}}^{(\varphi)}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

By Proposition H.3, each deviation direction is bounded by $\Delta_{\varphi,3,t_\varphi}(\delta)$ with conditional failure probability at most δ. On their intersection, approximate empirical optimality of $\hat{f}^{(\varphi)}_{\mathrm{ATE}}$ and the projection identity give the result, the optimization gap entering additively once.

□

Theorem H.24 (TREATMENTNET dose-response consistency)**.** *Suppose Assumptions H.20 and H.22 hold, and suppose $\hat{\varphi}$ is square-integrable. For every $\delta \in (0, 1/2)$, conditional on the trained bridge, with probability at least $1 - 2\delta$,*

$$\|\hat{f}^{(\varphi)}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le 2\kappa_{\varphi,3,t_\varphi} + 4\Delta_{\varphi,3,t_\varphi}(\delta) + 2\xi_{\varphi,3,t_\varphi} + 2C_h^2\|T_\varphi\hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})}.$$

The multiplier 2 on $\xi_{\varphi,3,t_\varphi}$ is the same $\|u+v\|^2 \le 2\|u\|^2 + 2\|v\|^2$ factor that doubles $\kappa_{\varphi,3,t_\varphi}$ and $2\Delta_{\varphi,3,t_\varphi}(\delta)$. Since $T_\varphi\varphi_0^{\mathrm{ATE}} = r_0$, the projected error is equivalently

$$\|T_\varphi\hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})} = \|\hat{\varphi} - \varphi_0^{\mathrm{ATE}}\|^2_{T_\varphi}.$$

Thus, the preceding bound requires only projected-norm consistency of the treatment bridge. If, in addition, the assumptions of Theorem H.17 hold, then for $\delta \in (0, 1/6)$, intersecting this event with that theorem's event gives joint probability at least $1 - 6\delta$, and

$$\begin{aligned}
&\|\hat{f}^{(\varphi)}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \\
&\quad\le 2\kappa_{\varphi,3,t_\varphi} + 4\Delta_{\varphi,3,t_\varphi}(\delta) + 8C_h^2\kappa_{\varphi,2,\nu_\varphi} \\
&\qquad + 12C_h^2B_\theta^2\kappa_{\varphi,1,\nu_\varphi} + 24C_h^2B_\theta^2\Delta_{\varphi,1,\nu_\varphi}(\delta) + 8C_h^2\Delta^{\mathrm{RF},U}_{\varphi,2,\nu_\varphi}(\delta) \\
&\qquad + 2\xi_{\varphi,3,t_\varphi} + 12C_h^2B_\theta^2\xi_{\varphi,1,\nu_\varphi} + 4C_h^2\xi_{\varphi,2,\nu_\varphi}.
\end{aligned}$$

This follows by substituting the projected-norm bound in Theorem H.17 of Appendix H.4.

Proof. Decompose

$$\hat{f}^{(\varphi)}_{\mathrm{ATE}} - f_{\mathrm{ATE}} = (\hat{f}^{(\varphi)}_{\mathrm{ATE}} - \bar{f}^{(\varphi)}_{\hat{\varphi}}) + (\bar{f}^{(\varphi)}_{\hat{\varphi}} - f_{\mathrm{ATE}}).$$

Using $\|u+v\|_2^2 \le 2\|u\|_2^2 + 2\|v\|_2^2$,

$$\|\hat{f}^{(\varphi)}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le 2\|\hat{f}^{(\varphi)}_{\mathrm{ATE}} - \bar{f}^{(\varphi)}_{\hat{\varphi}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} + 2\|\bar{f}^{(\varphi)}_{\hat{\varphi}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

The first term is controlled by Lemma H.23.

The second is controlled by Lemma H.21. Substituting both bounds gives the result. □

Corollary H.25 (TREATMENTNET dose-response consistency)**.** *Suppose the assumptions of Theorems H.17 and H.24 hold, $n_\varphi, m_\varphi, t_\varphi \to \infty$, and*

$$\kappa_{\varphi,1,\nu_\varphi} \to 0, \qquad \kappa_{\varphi,2,\nu_\varphi} \to 0, \qquad \kappa_{\varphi,3,t_\varphi} \xrightarrow{p} 0.$$

Assume also that

$$\widehat{\mathfrak{R}}_{D_1^{(\varphi)}}(\mathcal{L}_{\varphi,1,\nu_\varphi}) \xrightarrow{p} 0, \qquad \widehat{\mathfrak{R}}^{(2),U}_{D_2^{(\varphi)}}(\mathcal{L}^{\mathrm{RF}}_{\varphi,2,\nu_\varphi}) \xrightarrow{p} 0,$$

and

$$\widehat{\mathfrak{R}}_{D_3^{(\varphi)}}(\mathcal{L}_{\varphi,3,t_\varphi}(\hat{\varphi})) \xrightarrow{p} 0.$$

Assume in addition that the optimization gaps vanish,

$$\xi_{\varphi,1,\nu_\varphi} \xrightarrow{p} 0, \qquad \xi_{\varphi,2,\nu_\varphi} \xrightarrow{p} 0, \qquad \xi_{\varphi,3,t_\varphi} \xrightarrow{p} 0.$$

Then

$$\|\hat{f}^{(\varphi)}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \xrightarrow{p} 0.$$

H.6 DOSE-RESPONSE CONSISTENCY OF DOUBLY ROBUST METHOD

We analyze the DR-PCL-NET estimator for $Y \in \mathbb{R}^{d_Y}$. Condition throughout on the joint training of $\hat{h}$ and $\hat{\varphi}$; the two fitted bridges need not be independent of one another. The final-stage fold, including any data-dependent preprocessing, is independent of these bridge function training procedures. The result uses only projected bridge residuals and does not require strong L^2-convergence of either bridge.

Population plug-in DR functional. Define

$$\mu_{\hat{h}}(a) := \mathbb{E}[\hat{h}(a, X, W)], \qquad K_{\hat{h},\hat{\varphi}}(a) := \mathbb{E}[\hat{\varphi}(a, X, Z)\{Y - \hat{h}(a, X, W)\} \mid A = a].$$

The population plug-in DR-PCL-NET functional is

$$\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) := \mu_{\hat{h}}(a) + K_{\hat{h},\hat{\varphi}}(a).$$

Lemma H.26 (ATE doubly robust identity)**.** *For $\mathbb{P}_A$-almost every a,*

$$f_{\mathrm{ATE}}(a) - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) = \mathbb{E}[\{\varphi_0^{\mathrm{ATE}}(a, X, Z) - \hat{\varphi}(a, X, Z)\}\{h_0(a, X, W) - \hat{h}(a, X, W)\} \mid A = a].$$

Proof. By the outcome bridge equation,

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

Since $\hat{\varphi}(a, X, Z)$ is a function of (X, Z),

$$\mathbb{E}[\hat{\varphi}(a, X, Z)\{Y - h_0(a, X, W)\} \mid A = a] = 0.$$

Therefore,

$$\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) = \mathbb{E}[\hat{h}(a, X, W)] + \mathbb{E}[\hat{\varphi}(a, X, Z)\{h_0(a, X, W) - \hat{h}(a, X, W)\} \mid A = a].$$

By the outcome bridge representation,

$$f_{\mathrm{ATE}}(a) = \mathbb{E}[h_0(a, X, W)].$$

Hence

$$\begin{aligned}\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) - f_{\mathrm{ATE}}(a) = \mathbb{E}[\hat{h}(a, X, W) - h_0(a, X, W)] \\ + \mathbb{E}[\hat{\varphi}(a, X, Z)\{h_0(a, X, W) - \hat{h}(a, X, W)\} \mid A = a].\end{aligned}$$

The treatment bridge reweighting identity gives, for any square-integrable $q(a, X, W)$,

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, X, Z)q(a, X, W) \mid A = a] = \mathbb{E}[q(a, X, W)].$$

Applying this identity with

$$q(a, X, W) = \hat{h}(a, X, W) - h_0(a, X, W)$$

yields

$$\mathbb{E}[\hat{h}(a, X, W) - h_0(a, X, W)] = \mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, X, Z)\{\hat{h}(a, X, W) - h_0(a, X, W)\} \mid A = a].$$

Substitution gives

$$\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) - f_{\mathrm{ATE}}(a) = \mathbb{E}[\{\hat{\varphi}(a, X, Z) - \varphi_0^{\mathrm{ATE}}(a, X, Z)\}\{h_0(a, X, W) - \hat{h}(a, X, W)\} \mid A = a].$$

Multiplying both sides by -1 proves the claim. □

Projected residual bounds for the population DR remainder. Define the conditional bridge-error constants

$$B^2_{e,\varphi} := \operatorname*{ess\,sup}_{a \sim \mathbb{P}_A} \mathbb{E}[\{\hat{\varphi}(a, X, Z) - \varphi_0^{\mathrm{ATE}}(a, X, Z)\}^2 \mid A = a],$$

and

$$B^2_{e,h} := \operatorname*{ess\,sup}_{a \sim \mathbb{P}_A} \mathbb{E}[\|h_0(a, X, W) - \hat{h}(a, X, W)\|_2^2 \mid A = a].$$

Conditional on bridge training, assume $B_{e,\varphi}, B_{e,h} < \infty$. The consistency statement below strengthens this to $B_{e,\varphi}, B_{e,h} = O_p(1)$; this is boundedness, not strong bridge convergence.

Define the projected residuals

$$\mathcal{R}_h^{\mathrm{proj}} := \|T_h \hat{h} - m_0\|^2_{L^2(\mathbb{P}_{Q_h};\mathbb{R}^{d_Y})},$$

and

$$\mathcal{R}_\varphi^{\mathrm{proj}} := \|T_\varphi \hat{\varphi} - r_0\|^2_{L^2(\mathbb{P}_{Q_\varphi})}.$$

Lemma H.27 (Population DR remainder controlled by projected residuals)**.** *Assume $B_{e,\varphi} < \infty$ and $B_{e,h} < \infty$. Then*

$$\|\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le B^2_{e,\varphi} \mathcal{R}_h^{\mathrm{proj}},$$

and

$$\|\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le B^2_{e,h} \mathcal{R}_\varphi^{\mathrm{proj}}.$$

Consequently,

$$\|\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le \min\left\{ B^2_{e,\varphi} \mathcal{R}_h^{\mathrm{proj}},\ B^2_{e,h} \mathcal{R}_\varphi^{\mathrm{proj}} \right\}.$$

Proof. By Lemma H.26,

$$f_{\mathrm{ATE}}(a) - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) = \mathbb{E}[\{\varphi_0^{\mathrm{ATE}}(a,X,Z) - \hat{\varphi}(a,X,Z)\}\{h_0(a,X,W) - \hat{h}(a,X,W)\} \mid A = a].$$

First condition on (X, Z). Since

$$\mathbb{E}[h_0(a,X,W) - \hat{h}(a,X,W) \mid A = a, X, Z] = m_0(a,X,Z) - T_h \hat{h}(a,X,Z),$$

we obtain

$$f_{\mathrm{ATE}}(a) - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) = \mathbb{E}[\{\varphi_0^{\mathrm{ATE}}(a,X,Z) - \hat{\varphi}(a,X,Z)\}\{m_0(a,X,Z) - T_h \hat{h}(a,X,Z)\} \mid A = a].$$

By Cauchy–Schwarz,

$$\begin{aligned}\|f_{\mathrm{ATE}}(a) - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a)\|_2^2 \le\ & \mathbb{E}[\{\varphi_0^{\mathrm{ATE}}(a,X,Z) - \hat{\varphi}(a,X,Z)\}^2 \mid A = a] \\ & \times \mathbb{E}[\|m_0(a,X,Z) - T_h \hat{h}(a,X,Z)\|_2^2 \mid A = a].\end{aligned}$$

Using the definition of $B_{e,\varphi}$ and integrating over $a \sim \mathbb{P}_A$ gives

$$\|\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le B^2_{e,\varphi} \mathcal{R}_h^{\mathrm{proj}}.$$

Secondly, condition on (X, W). Since

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a,X,Z) - \hat{\varphi}(a,X,Z) \mid A = a, X, W] = r_0(a,X,W) - T_\varphi \hat{\varphi}(a,X,W),$$

we similarly obtain

$$f_{\mathrm{ATE}}(a) - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) = \mathbb{E}[\{h_0(a,X,W) - \hat{h}(a,X,W)\}\{r_0(a,X,W) - T_\varphi \hat{\varphi}(a,X,W)\} \mid A = a].$$

Cauchy–Schwarz and integration over $a \sim \mathbb{P}_A$ give

$$\|\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le B^2_{e,h} \mathcal{R}_\varphi^{\mathrm{proj}}.$$

Taking the smaller of the two valid bounds completes the proof. □

Final-stage residual regression. Define the residual pseudo-outcome

$$\Gamma_{\mathrm{DR}} := \hat{\varphi}(A, X, Z)\{Y - \hat{h}(A, X, W)\} \in \mathbb{R}^{d_Y}.$$

Then

$$K_{\hat{h},\hat{\varphi}}(a) = \mathbb{E}[\Gamma_{\mathrm{DR}} \mid A = a].$$

Let $\mathcal{D}_{\mathrm{DR}} = \{(y_i, a_i, x_i, z_i, w_i)\}_{i \in \mathcal{I}_3^{(\mathrm{DR})}}$ be the independent final-stage sample, with $|\mathcal{I}_3^{(\mathrm{DR})}| = t_{\mathrm{DR}}$. Set $y_i^{(\mathrm{DR})} := \hat{\varphi}(a_i, x_i, z_i)\{y_i - \hat{h}(a_i, x_i, w_i)\}$, matching the main-text pseudo-outcome notation. Let ν_{DR} index a class

$$\mathcal{K}_{\mathrm{DR},\nu_{\mathrm{DR}}} \subset \{k : \mathcal{A} \to \mathbb{R}^{d_Y}\}.$$

Writing $\widehat{\mathcal{R}}_{\mathrm{DR},3}(k) := t_{\mathrm{DR}}^{-1} \sum_{i \in \mathcal{I}_3^{(\mathrm{DR})}} \|y_i^{(\mathrm{DR})} - k(a_i)\|_2^2$ for the final-stage empirical risk, let $\hat{K}_{\mathrm{DR}} \in \mathcal{K}_{\mathrm{DR},\nu_{\mathrm{DR}}}$ be a measurable approximate empirical minimizer,

$$\widehat{\mathcal{R}}_{\mathrm{DR},3}(\hat{K}_{\mathrm{DR}}) \le \inf_{k \in \mathcal{K}_{\mathrm{DR},\nu_{\mathrm{DR}}}} \widehat{\mathcal{R}}_{\mathrm{DR},3}(k) + \xi_{\mathrm{DR},3,\nu_{\mathrm{DR}}}, \qquad \xi_{\mathrm{DR},3,\nu_{\mathrm{DR}}} \ge 0.$$

For the neural parameterization in the algorithmic section, $\hat{K}_{\mathrm{DR}}(a) = k^{\mathrm{DR}}(a; \hat{\theta}_3^{\mathrm{DR}})$. Define the approximation error

$$\kappa_{\mathrm{DR},\nu_{\mathrm{DR}}} := \inf_{k \in \mathcal{K}_{\mathrm{DR},\nu_{\mathrm{DR}}}} \|k - K_{\hat{h},\hat{\varphi}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

Define the residual-regression loss class

$$\mathcal{L}_{\mathrm{DR},\nu_{\mathrm{DR}}} := \left\{(\gamma, a) \mapsto \|\gamma - k(a)\|_2^2 : k \in \mathcal{K}_{\mathrm{DR},\nu_{\mathrm{DR}}}\right\}.$$

Assumption H.28 (DR final-stage boundedness)**.** *Conditional on the trained bridge estimators $(\hat{h}, \hat{\varphi})$, assume:*

- *The residual-regression loss class has a common deterministic envelope $M_{\mathrm{DR}} < \infty$:*

$$0 \le \ell \le M_{\mathrm{DR}} \quad \textit{for every } \ell \in \mathcal{L}_{\mathrm{DR},\nu_{\mathrm{DR}}}.$$

- *The final-stage sample $\mathcal{D}_{\mathrm{DR}}$ is independent of all data used to train $\hat{h}$ and $\hat{\varphi}$, including fitted preprocessing transformations.*

Define

$$\Delta_{\mathrm{DR},\nu_{\mathrm{DR}}}(\delta) := 2\widehat{\mathfrak{R}}_{\mathcal{D}_{\mathrm{DR}}}(\mathcal{L}_{\mathrm{DR},\nu_{\mathrm{DR}}}) + 3M_{\mathrm{DR}}\sqrt{\frac{\log(2/\delta)}{2t_{\mathrm{DR}}}}.$$

Lemma H.29 (Final-stage residual-regression bound)**.** *Under Assumption H.28, conditional on the trained bridge estimators, with probability at least $1 - 2\delta$,*

$$\|\hat{K}_{\mathrm{DR}} - K_{\hat{h},\hat{\varphi}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le \kappa_{\mathrm{DR},\nu_{\mathrm{DR}}} + 2\Delta_{\mathrm{DR},\nu_{\mathrm{DR}}}(\delta) + \xi_{\mathrm{DR},3,\nu_{\mathrm{DR}}}.$$

Proof. Conditional on $(\hat{h}, \hat{\varphi})$, this is ordinary least-squares regression with response Γ_{DR} and covariate A. Since

$$K_{\hat{h},\hat{\varphi}}(A) = \mathbb{E}[\Gamma_{\mathrm{DR}} \mid A, \hat{h}, \hat{\varphi}],$$

the least-squares projection identity gives

$$\mathbb{E}[\|\Gamma_{\mathrm{DR}} - k(A)\|_2^2 \mid \hat{h}, \hat{\varphi}] - \mathbb{E}[\|\Gamma_{\mathrm{DR}} - K_{\hat{h},\hat{\varphi}}(A)\|_2^2 \mid \hat{h}, \hat{\varphi}] = \|k - K_{\hat{h},\hat{\varphi}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

Apply both directions of Proposition H.3 and use approximate empirical optimality of $\hat{K}_{\mathrm{DR}}$, which adds $\xi_{\mathrm{DR},3,\nu_{\mathrm{DR}}}$ once. Their intersection has conditional probability at least $1 - 2\delta$, which proves the result. □

On the same final-stage sample, set

$$\hat{\mu}_h(a) := \frac{1}{t_{\mathrm{DR}}} \sum_{i=1}^{t_{\mathrm{DR}}} \hat{h}(a, x_i, w_i), \qquad \mathcal{E}_{\mu,h} := \|\hat{\mu}_h - \mu_{\hat{h}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

Theorem H.14 gives $\mathcal{E}_{\mu,h} = O_p(t_{\mathrm{DR}}^{-1})$. The two final-stage components may use the same fold; their mutual independence is unnecessary.

The DR-PCL-NET estimator is

$$\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}}(a) := \hat{\mu}_h(a) + \hat{K}_{\mathrm{DR}}(a).$$

Theorem H.30 (DR-PCL-NET ATE error bound)**.** *Suppose valid h_0 and φ_0^{ATE} exist, the conditional bridge-error bounds above are finite, and Assumption H.28 holds. Conditional on bridge training, with probability at least $1 - 2\delta$,*

$$\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le 4\mathcal{E}_{\mu,h} + 4\kappa_{\mathrm{DR},\nu_{\mathrm{DR}}} + 8\Delta_{\mathrm{DR},\nu_{\mathrm{DR}}}(\delta) + 4\xi_{\mathrm{DR},3,\nu_{\mathrm{DR}}} + 2\min\left\{B^2_{e,\varphi}\mathcal{R}^{\mathrm{proj}}_h,\ B^2_{e,h}\mathcal{R}^{\mathrm{proj}}_\varphi\right\}.$$

The multiplier 4 is the one already applied to the residual-regression error: the decomposition uses $\|u+v\|^2 \le 2\|u\|^2 + 2\|v\|^2$ twice. The projected bridge errors $\mathcal{R}^{\mathrm{proj}}_h, \mathcal{R}^{\mathrm{proj}}_\varphi$ are those of Theorems H.8 and H.17, and therefore already carry their own optimization gaps; the exact doubly robust product-remainder identity is unchanged.

Proof. Decompose

$$\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}} = (\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}) + (\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}).$$

Using $\|u+v\|_2^2 \le 2\|u\|_2^2 + 2\|v\|_2^2$,

$$\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le 2\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} + 2\|\bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

Moreover,

$$\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}} = (\hat{\mu}_h - \mu_{\hat{h}}) + (\hat{K}_{\mathrm{DR}} - K_{\hat{h},\hat{\varphi}}).$$

Therefore,

$$\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - \bar{f}^{\mathrm{DR}}_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le 2\mathcal{E}_{\mu,h} + 2\|\hat{K}_{\mathrm{DR}} - K_{\hat{h},\hat{\varphi}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

The residual-regression term is bounded by Lemma H.29. The population DR remainder is bounded by Lemma H.27. Combining these bounds gives the theorem. □

Assumption H.31 (Bounded complementary bridge errors)**.** *Along the asymptotic sequence, the action-uniform conditional bridge-error envelopes defined above satisfy*

$$B_{e,\varphi} = O_p(1), \qquad B_{e,h} = O_p(1).$$

This assumption requires stochastic boundedness of the bridge error on the route complementary to the vanishing projected error; it does not require strong consistency of either fitted bridge.

Corollary H.32 (Projected doubly robust consistency)**.** *Let $\delta_t \downarrow 0$ satisfy $\log(1/\delta_t)/t_{\mathrm{DR}} \to 0$, and define*

$$\rho_{\mathrm{DR},3,\nu_{\mathrm{DR}}} := \kappa_{\mathrm{DR},\nu_{\mathrm{DR}}} + 2\Delta_{\mathrm{DR},\nu_{\mathrm{DR}}}(\delta_t).$$

Define the optimization-aware final-stage rate $\tilde{\rho}_{\mathrm{DR},3,\nu_{\mathrm{DR}}} := \rho_{\mathrm{DR},3,\nu_{\mathrm{DR}}} + \xi_{\mathrm{DR},3,\nu_{\mathrm{DR}}}$.

Suppose Assumption H.31 holds, $t_{\mathrm{DR}} \to \infty$,

$$\kappa_{\mathrm{DR},\nu_{\mathrm{DR}}} = o_p(1), \qquad \widehat{\mathfrak{R}}_{\mathcal{D}_{\mathrm{DR}}}(\mathcal{L}_{\mathrm{DR},\nu_{\mathrm{DR}}}) = o_p(1), \qquad \xi_{\mathrm{DR},3,\nu_{\mathrm{DR}}} = o_p(1).$$

Suppose further that the rate conclusions

$$\mathcal{R}^{\mathrm{proj}}_h = O_p(\tilde{\rho}_{h,\nu_h}), \qquad \mathcal{R}^{\mathrm{proj}}_\varphi = O_p(\tilde{\rho}_{\varphi,\nu_\varphi})$$

of Corollaries H.9 and H.18 hold, and that Theorem H.14 applies with $t_h = t_{\mathrm{DR}}$. Then $\tilde{\rho}_{\mathrm{DR},3,\nu_{\mathrm{DR}}} = o_p(1)$, and these results give

$$\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATE}} - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p\big(t_{\mathrm{DR}}^{-1} + \tilde{\rho}_{\mathrm{DR},3,\nu_{\mathrm{DR}}} + \min\{\tilde{\rho}_{h,\nu_h}, \tilde{\rho}_{\varphi,\nu_\varphi}\}\big)\,. \tag{75}$$

Hence DR-PCL-NET *is consistent if the final-stage errors and the final-stage optimization gap vanish and either $\tilde{\rho}_{h,\nu_h} \to 0$ or $\tilde{\rho}_{\varphi,\nu_\varphi} \to 0$, that is, if one route's approximation, empirical-process and optimization terms all vanish. There is no additive density-ratio estimation error.*

Remark H.33 (Projected double robustness)**.** *The theorem gives double robustness in projected-residual form. The estimator is consistent if either the outcome-side projected residual*

$$T_h\hat{h} - m_0$$

vanishes, or the treatment-side projected residual

$$T_\varphi\hat{\varphi} - r_0$$

vanishes. No strong L^2-convergence of either bridge is required. A second-order product rate in strong bridge errors would require additional strong-norm or conditional-profile assumptions and is not used here.

The bridge on the opposite side enters only through the $O_p(1)$ conditional error bound. For a fixed number K_{cf} of cross-fitting folds, the fold-averaged estimator inherits the result because

$$\left\| K_{\mathrm{cf}}^{-1} \sum_{k=1}^{K_{\mathrm{cf}}} (\hat{f}_k - f_{\mathrm{ATE}}) \right\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})} \le K_{\mathrm{cf}}^{-1} \sum_{k=1}^{K_{\mathrm{cf}}} \|\hat{f}_k - f_{\mathrm{ATE}}\|^2_{L^2(\mathbb{P}_A;\mathbb{R}^{d_Y})}.$$

H.7 SELF-NORMALIZED DENSITY-RATIO-FREE HETEROGENEOUS DOSE-RESPONSE

Write $X = (S, V)$ and $\hat{\varphi}^{\mathrm{ATE}}(a, v, s, z) := \hat{\varphi}(a, (s, v), z)$, where $\hat{\varphi}$ is the estimator from Appendix H.4. Condition throughout on the samples used to fit $\hat{h}$ and $\hat{\varphi}^{\mathrm{ATE}}$. Let

$$\begin{aligned} r_A &:= r_0 := r_{0,\mathrm{ATE}}, & r_C &:= r_{0,\mathrm{CATE}}, \\ e_A &:= T_\varphi\hat{\varphi}^{\mathrm{ATE}} - r_A, & \mathcal{E}_{\varphi,A} &:= \|e_A\|^2_{L^2(\mathbb{P}_{A,V,S,W})}. \end{aligned}$$

Also, define

$$\begin{aligned} D_{\hat{\varphi}}(a, v) &:= \mathbb{E}[\hat{\varphi}^{\mathrm{ATE}}(a, v, S, Z) \mid A = a, V = v], \\ \bar{e}_A(a, v) &:= \mathbb{E}[e_A(a, v, S, W) \mid A = a, V = v], \\ \widetilde{\varphi}(a, v, s, z) &:= \frac{\hat{\varphi}^{\mathrm{ATE}}(a, v, s, z)}{D_{\hat{\varphi}}(a, v)}. \end{aligned}$$

Lemma H.34 (ATE-to-CATE projected-error transfer)**.** *Suppose, on the evaluation region,*

$$D_{\hat{\varphi}}(a, v) \ge d_C > 0, \qquad |r_C(a, v, s, w)| \le B_C < \infty.$$

Then

$$\mathcal{E}_{\varphi,C} := \|T^{\mathrm{CATE}}_\varphi\widetilde{\varphi} - r_C\|^2_{L^2(\mathbb{P}_{A,V,S,W})} \le C_C\mathcal{E}_{\varphi,A}, \qquad C_C := \frac{2(1 + B_C^2)}{d_C^2}. \tag{76}$$

Proof. By total law of expectation and Proposition 3.4, we have $D_{\hat{\varphi}} = d_0 + \bar{e}_A$ and $r_C = r_A/d_0$. As result

$$T^{\mathrm{CATE}}_\varphi\widetilde{\varphi} - r_C = \frac{e_A - r_C\bar{e}_A}{D_{\hat{\varphi}}}.$$

Then, the result follows from $(x + y)^2 \le 2x^2 + 2y^2$ and $\|\bar{e}_A\|_{L^2(\mathbb{P}_{A,V})} \le \|e_A\|_{L^2(\mathbb{P}_{A,V,S,W})}$, by Jensen Inequality. □

Joint third-stage regression. For $Y \in \mathbb{R}^{d_Y}$, define

$$\Gamma_N := Y\hat{\varphi}^{\mathrm{ATE}}(A,V,S,Z),$$
$$\Gamma_K := \hat{\varphi}^{\mathrm{ATE}}(A,V,S,Z)\{Y - \hat{h}(A,V,S,W)\},$$
$$\Gamma_D := \hat{\varphi}^{\mathrm{ATE}}(A,V,S,Z),$$

and let $N := N_{\hat{\varphi}} := \mathbb{E}[\Gamma_N \mid A,V]$ and $K := K_{\hat{h},\hat{\varphi}} := \mathbb{E}[\Gamma_K \mid A,V]$; also note that $\mathbb{E}[\Gamma_D \mid A,V] = D_{\hat{\varphi}}$. For $G \in \{N,K\}$, let

$$\mathcal{D}_{3,G} := \{(a_i, v_i, \gamma_{G,i}, \gamma_{D,i})\}_{i=1}^{t_G}$$

be conditionally i.i.d. and independent of the bridge-fitting samples. Let $\mathcal{F}_{G,t_G}$ be a shared-network of (d_Y+1)-output class contained in

$$\{(g,d) : \mathcal{A} \times \mathcal{V} \longrightarrow \mathbb{R}^{d_Y} \times \mathbb{R}\}$$

with target $F_G^0 := (G, D_{\hat{\varphi}})$, and define

$$\kappa_{G,t_G} := \inf_{F \in \mathcal{F}_{G,t_G}} \|F - F_G^0\|^2_{L^2(\mathbb{P}_{A,V};\mathbb{R}^{d_Y+1})}.$$

Let $\mathcal{L}_{G,t_G}$ be the induced real-valued joint squared-loss class,

$$\mathcal{L}_{G,t_G} := \left\{(a,v,\gamma_G,\gamma_D) \mapsto \|\gamma_G - g(a,v)\|_2^2 + \{\gamma_D - d(a,v)\}^2 : (g,d) \in \mathcal{F}_{G,t_G}\right\}.$$

Assume this class is bounded by $M_G < \infty$, and set

$$\Delta_{G,t_G}(\delta) := 2\widehat{\mathfrak{R}}_{\mathcal{D}_{3,G}}(\mathcal{L}_{G,t_G}) + 3M_G\sqrt{\frac{\log(2/\delta)}{2t_G}}, \qquad \rho_{G,3,t_G}(\delta) := \kappa_{G,t_G} + 2\Delta_{G,t_G}(\delta).$$

Let $\hat{F}_G = (\hat{G}, \hat{D}_G) \in \mathcal{F}_{G,t_G}$ be a measurable approximate empirical minimizer of the joint empirical risk $\widehat{\mathcal{R}}_{G,3}$,

$$\widehat{\mathcal{R}}_{G,3}(\hat{F}_G) \le \inf_{F \in \mathcal{F}_{G,t_G}} \widehat{\mathcal{R}}_{G,3}(F) + \xi_{G,3,t_G}, \qquad \xi_{G,3,t_G} \ge 0, \qquad G \in \{N,K\},$$

and define the optimization-aware third-stage rate $\tilde{\rho}_{G,3,t_G}(\delta) := \rho_{G,3,t_G}(\delta) + \xi_{G,3,t_G}$. Because the numerator and the normalizer are fitted by one shared network, there is a single gap per regression and no separate denominator gap. With probability at least $1-2\delta$, Proposition H.3 gives

$$\|\hat{G} - G\|^2_{L^2(\mathbb{P}_{A,V};\mathbb{R}^{d_Y})} + \|\hat{D}_G - D_{\hat{\varphi}}\|^2_{L^2(\mathbb{P}_{A,V})} \le \tilde{\rho}_{G,3,t_G}(\delta). \tag{77}$$

Here $(\hat{G}, \hat{D}_G) = (\hat{N}^{(\varphi)}, \hat{D}^{(\varphi)})$ for TREATMENTNET and $(\hat{K}^{\mathrm{DR}}, \hat{D}^{\mathrm{DR}})$ for the doubly robust estimator. Thus, joint learning changes only the third-stage approximation and Rademacher terms: the shared class must approximate both targets simultaneously.

Choose $0 < \epsilon_D < d_C$ and set $\hat{D}_{G,\epsilon} := \max\{\hat{D}_G, \epsilon_D\}$. If $\|G(a,v)\|_2 \le B_G$, then Equation 77 yields

$$\left\| \frac{\hat{G}}{\hat{D}_{G,\epsilon}} - \frac{G}{D_{\hat{\varphi}}} \right\|^2_{L^2(\mathbb{P}_{A,V};\mathbb{R}^{d_Y})} \le C_G\tilde{\rho}_{G,3,t_G}(\delta), \qquad C_G := \frac{2}{\epsilon_D^2}\left(1 + \frac{B_G^2}{d_C^2}\right). \tag{78}$$

This follows from the 1-Lipschitz property of clipping and a ratio decomposition.

Theorem H.35 (Self-normalized CATE consistency)**.** *Suppose the conditions above hold and that valid bridges h_0 and $\varphi_0^{\mathrm{CATE}}$ exist. Let*

$$(T_h^{\mathrm{CATE}}h)(a,v,s,z) := \mathbb{E}[h(a,v,s,W) \mid A=a, V=v, S=s, Z=z],$$
$$m_0^{\mathrm{CATE}}(a,v,s,z) := \mathbb{E}[Y \mid A=a, V=v, S=s, Z=z].$$

Define

$$M_{\hat{h}}(a,v) := \mathbb{E}[\hat{h}(a,v,S,W) \mid V=v],$$
$$\mathcal{E}_{h,C} := \|T_h^{\mathrm{CATE}}\hat{h} - m_0^{\mathrm{CATE}}\|^2_{L^2(\mathbb{P}_{A,V,S,Z};\mathbb{R}^{d_Y})},$$

$$\mathcal{E}_M := \|\hat{f}^{(h)}_{\mathrm{CATE}} - M_{\hat{h}}\|^2_{L^2(\mathbb{P}_{A,V};\mathbb{R}^{d_Y})}.$$

Assume the relevant conditional second-moment bounds are finite, and write them as

$$C^2_{h,C} := \operatorname*{ess\,sup}_{a,v} \mathbb{E}[\|h_0(a,v,S,W)\|_2^2 \mid A=a, V=v],$$
$$B^2_{\varphi,C} := \operatorname*{ess\,sup}_{a,v} \mathbb{E}[(\widetilde{\varphi}(a,v,S,Z) - \varphi_0^{\mathrm{CATE}}(a,v,S,Z))^2 \mid A=a, V=v],$$
$$B^2_{\hat{h},C} := \operatorname*{ess\,sup}_{a,v} \mathbb{E}[\|\hat{h}(a,v,S,W) - h_0(a,v,S,W)\|_2^2 \mid A=a, V=v].$$

For every $\delta \in (0, 1/4)$, each bound below holds conditionally with probability at least $1-2\delta$; if both third-stage regressions are fitted, the two bounds hold simultaneously with probability at least $1-4\delta$. The TREATMENTNET *estimator in Equation 62 satisfies*

$$\|\hat{f}^{(\varphi)}_{\mathrm{CATE}} - f_{\mathrm{CATE}}\|^2_{L^2(\mathbb{P}_{A,V};\mathbb{R}^{d_Y})} \le 2C_N \tilde{\rho}_{N,3,t_N}(\delta) + 2C^2_{h,C} C_C \mathcal{E}_{\varphi,A}. \tag{79}$$

Moreover, the estimator in Equation 64 satisfies

$$\|\hat{f}^{\mathrm{DR}}_{\mathrm{CATE}} - f_{\mathrm{CATE}}\|^2_{L^2(\mathbb{P}_{A,V};\mathbb{R}^{d_Y})} \le 3\mathcal{E}_M + 3C_K \tilde{\rho}_{K,3,t_K}(\delta) + 3\min\left\{B^2_{\varphi,C}\mathcal{E}_{h,C}, B^2_{\hat{h},C} C_C \mathcal{E}_{\varphi,A}\right\}. \tag{80}$$

Proof. The outcome bridge and a change of measure give

$$\frac{N_{\hat{\varphi}}}{D_{\hat{\varphi}}} - f_{\mathrm{CATE}} = \mathbb{E}\big[h_0(a,v,S,W)\{T^{\mathrm{CATE}}_{\varphi}\widetilde{\varphi} - r_C\} \,\big|\, A=a, V=v\big].$$

Thus, conditional Cauchy–Schwarz gives $\|N_{\hat{\varphi}}/D_{\hat{\varphi}} - f_{\mathrm{CATE}}\|_2^2 \le C^2_{h,C}\mathcal{E}_{\varphi,C}$. Combine this with Equations 76 and 78 to obtain Equation 79. For the DR method, the exact population remainder is

$$\mathbb{E}[(\widetilde{\varphi} - \varphi_0^{\mathrm{CATE}})\{h_0 - \hat{h}\} \mid A, V].$$

Conditioning successively on (A,V,S,Z) and (A,V,S,W), as in Lemma H.27, bounds its squared $L^2(\mathbb{P}_{A,V})$-norm by

$$\min\{B^2_{\varphi,C}\mathcal{E}_{h,C}, B^2_{\hat{h},C}\mathcal{E}_{\varphi,C}\}.$$

Equation 80 now follows from Equation 78 and $\|u+v+w\|_2^2 \le 3(\|u\|_2^2 + \|v\|_2^2 + \|w\|_2^2)$. □

Corollaries H.9 and H.18 give

$$\mathcal{E}_{h,C} = O_p(\tilde{\rho}_{h,\nu_h}), \qquad \mathcal{E}_{\varphi,A} = O_p(\tilde{\rho}_{\varphi,\nu_\varphi}),$$

so the CATE bounds inherit the bridge-stage optimization gaps through the ATE treatment bridge, which the CATE construction reuses; no new bridge-training gap is introduced here. Choose $\delta_G \downarrow 0$ such that $\log(1/\delta_G)/t_G \to 0$. If $\kappa_{G,t_G} = o_p(1)$, $\widehat{\mathfrak{R}}_{\mathcal{D}_{3,G}}(\mathcal{L}_{G,t_G}) = o_p(1)$, $M_G = O_p(1)$, and $\xi_{G,3,t_G} = o_p(1)$, then $\tilde{\rho}_{G,3,t_G} := \tilde{\rho}_{G,3,t_G}(\delta_G) = o_p(1)$. Provided all remaining constants in the bounds are $O_p(1)$, TREATMENTNET is consistent when $\tilde{\rho}_{N,3,t_N} = o_p(1)$ and $\tilde{\rho}_{\varphi,\nu_\varphi} \to 0$. The self-normalized DR estimator is consistent when $\mathcal{E}_M, \tilde{\rho}_{K,3,t_K} = o_p(1)$ and either $\tilde{\rho}_{h,\nu_h} \to 0$ or $\tilde{\rho}_{\varphi,\nu_\varphi} \to 0$. Thus, joint learning adds one joint third-stage regression error and its optimization gap; it neither alters projected double robustness nor requires strong bridge convergence.

H.8 TREATMENT-BRIDGE CONSISTENCY FOR CONDITIONAL DOSE-RESPONSE

We analyze the complete ordered-pair objective in Equation 69. Retain the first-stage objects $\Phi_{\nu_\varphi}, \mathcal{M}_{\nu_\varphi}, \mu^{(\varphi)}_{\phi}, \hat{\mu}^{(\varphi)}_{\phi}, \kappa_{\varphi,1,\nu_\varphi}$, and $\Delta_{\varphi,1,\nu_\varphi}$ from Appendix H.4. The reference level A' does not enter this conditional mean embedding.

Let

$$\Theta^{\mathrm{ATT}}_{\nu_\varphi} \subset \{\theta : \mathcal{A}\times\mathcal{A}\times\mathcal{X} \to \mathbb{R}^{d_{\nu_\varphi}}\},$$

and define

$$\varphi^{\mathrm{ATT}}_{\theta,\phi}(a,a',x,z) := \langle \theta(a,a',x), \phi(z)\rangle,$$

$$g^{\mathrm{ATT}}_{\theta,\phi}(a,a',x,w) := \langle \theta(a,a',x), \mu^{(\varphi)}_{\phi}(a,x,w)\rangle = (T_{\mathrm{ATT}}\varphi^{\mathrm{ATT}}_{\theta,\phi})(a,a',x,w),$$
$$\hat{g}^{\mathrm{ATT}}_{\theta,\phi}(a,a',x,w) := \langle \theta(a,a',x), \hat{\mu}^{(\varphi)}_{\phi}(a,x,w)\rangle.$$

For this operator, write

$$\|u\|_{T_{\mathrm{ATT}}} := \|T_{\mathrm{ATT}}u\|_{L^2(\mathbb{P}_{\mathrm{ATT}})}.$$

Assume the bound $\|\theta(a,a',x)\|_2 \le B_\theta$ holds uniformly over $\Theta^{\mathrm{ATT}}_{\nu_\varphi}$, in addition to the first-stage bounds in Assumption H.15. Then

$$\sup_{\theta,\phi} \|\hat{g}^{\mathrm{ATT}}_{\theta,\phi} - g^{\mathrm{ATT}}_{\theta,\phi}\|^2_{L^2(\mathbb{P}_{\mathrm{ATT}})} \le B^2_\theta \epsilon_{\varphi,1,\nu_\varphi}, \tag{81}$$

because the independent anchor A' integrates out.

For $q = (a,x,w)$ and $q' = (a',x',w')$, define the conditional plug-in kernel

$$\hat{\ell}^{\mathrm{ATT}}_{\theta,\phi}(q,q') := \{\hat{g}^{\mathrm{ATT}}_{\theta,\phi}(a,a',x,w)\}^2 - 2\hat{g}^{\mathrm{ATT}}_{\theta,\phi}(a,a',x',w').$$

Notice that the conditional embedding in the transported term is evaluated at $\hat{\mu}^{(\varphi)}_{\phi}(a,x',w')$: a remains the source treatment, while (a',x',w') comes from the anchor unit. Define the deterministic pairwise loss class

$$\mathcal{L}^{\mathrm{ATT,RF}}_{\varphi,2,\nu_\varphi} := \Big\{(q,q') \mapsto \langle \theta(a,a',x), \eta(a,x,w)\rangle^2 - 2\langle \theta(a,a',x), \eta(a,x',w')\rangle : \theta \in \Theta^{\mathrm{ATT}}_{\nu_\varphi}, \quad \eta \in \mathcal{M}_{\nu_\varphi}\Big\}.$$

Its envelope is

$$M^{\mathrm{ATT,RF}}_{\varphi,2} := B^2_\theta B^2_\mu + 2B_\theta B_\mu.$$

On the independent second-stage sample $D^{(\varphi)}_2 = \{\tilde{q}_{\varphi,i}\}_{i\in\mathcal{I}^{(\varphi)}_2}$, set

$$\Delta^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi}(\delta) := 2\widehat{\mathfrak{R}}^{(2),U}_{D^{(\varphi)}_2}(\mathcal{L}^{\mathrm{ATT,RF}}_{\varphi,2,\nu_\varphi}) + 3M^{\mathrm{ATT,RF}}_{\varphi,2}\sqrt{\frac{8\log(2/\delta)}{m_\varphi}},$$

and define

$$\kappa^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} := \inf_{\theta\in\Theta^{\mathrm{ATT}}_{\nu_\varphi},\, \phi\in\Phi_{\nu_\varphi}} \|g^{\mathrm{ATT}}_{\theta,\phi} - r_{0,\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_{\mathrm{ATT}})}.$$

Write

$$\widehat{\mathcal{R}}^{\mathrm{ATT,RF},U}_{\varphi,2}(\theta,\phi) := \frac{1}{m_\varphi(m_\varphi-1)}\sum_{i\neq j}\hat{\ell}^{\mathrm{ATT}}_{\theta,\phi}(\tilde{q}_{\varphi,i},\tilde{q}_{\varphi,j})$$

and let $(\hat{\theta}^{\mathrm{ATT}},\hat{\phi}^{\mathrm{ATT}}) \in \Theta^{\mathrm{ATT}}_{\nu_\varphi}\times\Phi_{\nu_\varphi}$ be a measurable approximate empirical minimizer, that is, $\widehat{\mathcal{R}}^{\mathrm{ATT,RF},U}_{\varphi,2}(\hat{\theta}^{\mathrm{ATT}},\hat{\phi}^{\mathrm{ATT}}) \le \inf_{\theta,\phi}\widehat{\mathcal{R}}^{\mathrm{ATT,RF},U}_{\varphi,2}(\theta,\phi) + \xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi}$ with $\xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} \ge 0$, and set $\hat{\varphi}^{\mathrm{ATT}} := \varphi^{\mathrm{ATT}}_{\hat{\theta}^{\mathrm{ATT}},\hat{\phi}^{\mathrm{ATT}}}$.

Theorem H.36 (Projected ATT bridge error)**.** *Suppose the pair-law setup of Appendix G.2 holds, including $\mathbb{Q}_{\mathrm{ATT}} \ll \mathbb{P}_{\mathrm{ATT}}$ and $r_{0,\mathrm{ATT}} \in L^2(\mathbb{P}_{\mathrm{ATT}})$. Assume also that the preceding classes are measurable and separable and that both stages use measurable approximate empirical minimizers with gaps $\xi_{\varphi,1,\nu_\varphi}, \xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} \ge 0$. For every $\delta \in (0,1/4)$, with probability at least $1-4\delta$,*

$$\begin{aligned}\mathcal{R}^{\mathrm{proj}}_{\varphi,\mathrm{ATT}} &:= \|T_{\mathrm{ATT}}\hat{\varphi}^{\mathrm{ATT}} - r_{0,\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_{\mathrm{ATT}})} \\ &\le 4\kappa^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} + 6B^2_\theta\kappa_{\varphi,1,\nu_\varphi} + 12B^2_\theta\Delta_{\varphi,1,\nu_\varphi}(\delta) + 4\Delta^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi}(\delta) \\ &\quad + 6B^2_\theta\xi_{\varphi,1,\nu_\varphi} + 2\xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi}.\end{aligned}$$

The pairwise U-process deviation result itself is unchanged. If a valid bridge satisfies $T_{\mathrm{ATT}}\varphi^{\mathrm{ATT}}_0 = r_{0,\mathrm{ATT}}$, the left-hand side equals $\|\hat{\varphi}^{\mathrm{ATT}} - \varphi^{\mathrm{ATT}}_0\|^2_{T_{\mathrm{ATT}}}$.

Proof. Conditional on the first-stage sample, the plug-in kernels form a subclass of $\mathcal{L}^{\mathrm{ATT,RF}}_{\varphi,2,\nu_\varphi}$. Proposition H.5 therefore yields both uniform deviation directions, each with probability at least $1-\delta$.

Approximate empirical optimality and Equation 68 give

$$\|\hat{g}^{\mathrm{ATT}}_{\hat{\theta}^{\mathrm{ATT}},\hat{\phi}^{\mathrm{ATT}}} - r_{0,\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_{\mathrm{ATT}})} \le 2\kappa^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} + 2B^2_\theta \epsilon_{\varphi,1,\nu_\varphi} + 2\Delta^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi}(\delta) + \xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi}.$$

Since $T_{\mathrm{ATT}}\hat{\varphi}^{\mathrm{ATT}} = g^{\mathrm{ATT}}_{\hat{\theta}^{\mathrm{ATT}},\hat{\phi}^{\mathrm{ATT}}}$, a second use of Equation 81 yields

$$\mathcal{R}^{\mathrm{proj}}_{\varphi,\mathrm{ATT}} \le 4\kappa^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} + 6B^2_\theta \epsilon_{\varphi,1,\nu_\varphi} + 4\Delta^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi}(\delta) + 2\xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi}.$$

Lemma H.16 bounds $\epsilon_{\varphi,1,\nu_\varphi}$ by $\kappa_{\varphi,1,\nu_\varphi} + 2\Delta_{\varphi,1,\nu_\varphi}(\delta) + \xi_{\varphi,1,\nu_\varphi}$, which contributes $6B^2_\theta \xi_{\varphi,1,\nu_\varphi}$. The two first-stage and two conditional U-process events have total failure probability at most 4δ. □

Corollary H.37 (Projected joint-anchor ATT bridge consistency)**.** *Let $\gamma_{\varphi,1,\nu_\varphi}$ and $\gamma^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi}$ be deterministic sequences such that*

$$\widehat{\mathfrak{R}}_{D_1^{(\varphi)}}(\mathcal{L}_{\varphi,1,\nu_\varphi}) = O_p(\gamma_{\varphi,1,\nu_\varphi}), \quad \widehat{\mathfrak{R}}^{(2),U}_{D_2^{(\varphi)}}(\mathcal{L}^{\mathrm{ATT,RF}}_{\varphi,2,\nu_\varphi}) = O_p(\gamma^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi}).$$

Define

$$\rho_{\varphi,\mathrm{ATT},\nu_\varphi} := \kappa^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} + B^2_\theta \kappa_{\varphi,1,\nu_\varphi} + B^2_\theta \gamma_{\varphi,1,\nu_\varphi} + \gamma^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi} + \frac{B^2_\theta M_{\varphi,1}}{\sqrt{n_\varphi}} + \frac{M^{\mathrm{ATT,RF}}_{\varphi,2}}{\sqrt{m_\varphi}}.$$

Following the same convention of absorbing fixed bounded multipliers, define the optimization-aware ATT bridge rate

$$\tilde{\rho}_{\varphi,\mathrm{ATT},\nu_\varphi} := \rho_{\varphi,\mathrm{ATT},\nu_\varphi} + B^2_\theta \xi_{\varphi,1,\nu_\varphi} + \xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi}.$$

Then

$$\|T_{\mathrm{ATT}}\hat{\varphi}^{\mathrm{ATT}} - r_{0,\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_{\mathrm{ATT}})} = O_p(\tilde{\rho}_{\varphi,\mathrm{ATT},\nu_\varphi}).$$

Suppose $\gamma_{\varphi,1,\nu_\varphi} \to 0$, $\gamma^{\mathrm{ATT,RF},U}_{\varphi,2,\nu_\varphi} \to 0$, and $\kappa_{\varphi,1,\nu_\varphi} \to 0$, $\kappa^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} \to 0$ and $\xi_{\varphi,1,\nu_\varphi}, \xi^{\mathrm{ATT}}_{\varphi,2,\nu_\varphi} \to 0$ in probability with $n_\varphi, m_\varphi \to \infty$. Then,

$$\|T_{\mathrm{ATT}}\hat{\varphi}^{\mathrm{ATT}} - r_{0,\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_{\mathrm{ATT}})} \to 0 \ \textit{in probability.}$$

If $T_{\mathrm{ATT}}\varphi^{\mathrm{ATT}}_0 = r_{0,\mathrm{ATT}}$, then $\|\hat{\varphi}^{\mathrm{ATT}} - \varphi^{\mathrm{ATT}}_0\|^2_{T_{\mathrm{ATT}}} \to 0$ in probability.

The result controls the full joint-anchor projection under $\mathbb{P}_{\mathrm{ATT}}$; it does not require strong recovery of the raw bridge.

H.9 Conditional dose-response consistency of TreatmentNet

Let $Y \in \mathbb{R}^{d_Y}$, and condition on the samples used to fit $\hat{\varphi}^{\mathrm{ATT}}$. For a scalar candidate bridge u, define the induced treatment-conditional counterfactual mean

$$\bar{f}^{(\varphi,\mathrm{ATT})}_u(a,a') := \mathbb{E}[Yu(a,a',X,Z) \mid A = a] \in \mathbb{R}^{d_Y}.$$

Lemma H.38 (Projected ATT residual controls the response surface)**.** *Suppose Assumption H.20 holds. Then, for every candidate u for which the displayed products and conditional expectations are well-defined and square-integrable,*

$$\|\bar{f}^{(\varphi,\mathrm{ATT})}_u - f_{\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_A \otimes \mathbb{P}_A;\mathbb{R}^{d_Y})} \le C^2_h \|T_{\mathrm{ATT}}u - r_{0,\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_{\mathrm{ATT}})}. \tag{82}$$

Proof. The outcome-bridge equation and iterated expectation give

$$\bar{f}^{(\varphi,\mathrm{ATT})}_u(a,a') = \mathbb{E}[h_0(a,X,W)(T_{\mathrm{ATT}}u)(a,a',X,W) \mid A = a],$$
$$f_{\mathrm{ATT}}(a,a') = \mathbb{E}[h_0(a,X,W)r_{0,\mathrm{ATT}}(a,a',X,W) \mid A = a].$$

The second identity is the change of measure from $A = a$ to $A = a'$. Conditional Cauchy–Schwarz and integration over $(a,a') \sim \mathbb{P}_A \otimes \mathbb{P}_A$ prove the result. □

Let $D_{3,\mathrm{ATT}}^{(\varphi)} = \{o_i\}_{i=1}^{t_{\mathrm{ATT}}}$ be an i.i.d. sample independent of bridge training. For $i \neq j$, define

$$\gamma_{ij}^{\mathrm{ATT},\varphi} := y_i \hat{\varphi}^{\mathrm{ATT}}(a_i, a_j, x_i, z_i) \in \mathbb{R}^{d_Y}.$$

For two independent random copies O_1 and O_2, define $\Gamma_{12}^{\mathrm{ATT},\varphi} := Y_1 \hat{\varphi}^{\mathrm{ATT}}(A_1, A_2, X_1, Z_1)$. Its conditional mean is

$$\mathbb{E}[\Gamma_{12}^{\mathrm{ATT},\varphi} \mid A_1 = a, A_2 = a'] = \bar{f}_{\hat{\varphi}^{\mathrm{ATT}}}^{(\varphi,\mathrm{ATT})}(a, a').$$

Let $\mathcal{F}_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} \subset \{f : \mathcal{A} \times \mathcal{A} \to \mathbb{R}^{d_Y}\}$, and define

$$\kappa_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} := \inf_{f \in \mathcal{F}_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}} \|f - \bar{f}_{\hat{\varphi}^{\mathrm{ATT}}}^{(\varphi,\mathrm{ATT})}\|_{L^2(\mathbb{P}_A \otimes \mathbb{P}_A; \mathbb{R}^{d_Y})}^2,$$

$$\mathcal{L}_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}(\hat{\varphi}^{\mathrm{ATT}}) := \left\{(o_i, o_j) \mapsto \|y_i \hat{\varphi}^{\mathrm{ATT}}(a_i, a_j, x_i, z_i) - f(a_i, a_j)\|_2^2 : f \in \mathcal{F}_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}\right\}.$$

Conditional on bridge training, assume this real-valued loss class is measurable, separable, and bounded by a deterministic envelope $M_{\varphi,3}^{\mathrm{ATT}} < \infty$. Also, define

$$\Delta_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT},U}(\delta) := 2\widehat{\mathfrak{R}}_{D_{3,\mathrm{ATT}}^{(\varphi)}}^{(2),U}\left(\mathcal{L}_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}(\hat{\varphi}^{\mathrm{ATT}})\right) + 3M_{\varphi,3}^{\mathrm{ATT}}\sqrt{\frac{8\log(2/\delta)}{t_{\mathrm{ATT}}}}.$$

Let $\hat{f}_{\mathrm{ATT}}^{(\varphi)}$ be a measurable approximate minimizer of the complete ordered-pair loss

$$\widehat{\mathcal{R}}_{\varphi,3}^{\mathrm{ATT}}(f) := \frac{1}{t_{\mathrm{ATT}}(t_{\mathrm{ATT}} - 1)} \sum_{i \neq j} \|\gamma_{ij}^{\mathrm{ATT},\varphi} - f(a_i, a_j)\|_2^2,$$

that is, $\widehat{\mathcal{R}}_{\varphi,3}^{\mathrm{ATT}}(\hat{f}_{\mathrm{ATT}}^{(\varphi)}) \leq \inf_f \widehat{\mathcal{R}}_{\varphi,3}^{\mathrm{ATT}}(f) + \xi_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}$ with $\xi_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} \geq 0$.

Lemma H.39 (Pairwise ATT final-stage error)**.** *Conditional on bridge training, with probability at least* $1 - 2\delta$*,*

$$\|\hat{f}_{\mathrm{ATT}}^{(\varphi)} - \bar{f}_{\hat{\varphi}^{\mathrm{ATT}}}^{(\varphi,\mathrm{ATT})}\|_{L^2(\mathbb{P}_A \otimes \mathbb{P}_A; \mathbb{R}^{d_Y})}^2 \leq \kappa_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} + 2\Delta_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT},U}(\delta) + \xi_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}.$$

Proof. Under the law of two independent observations, $\bar{f}_{\hat{\varphi}^{\mathrm{ATT}}}^{(\varphi,\mathrm{ATT})}(A_1, A_2)$ is the conditional mean of $\Gamma_{12}^{\mathrm{ATT},\varphi}$. The vector-valued least-squares projection identity therefore holds, while the induced squared loss remains scalar. Apply both directions of Proposition H.5 and approximate empirical optimality. The latter contributes $\xi_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}$ once to the squared-error bound. □

Theorem H.40 (TREATMENTNET conditional dose-response error)**.** *Under the preceding conditions, for* $\delta \in (0, 1/6)$*, conditional on bridge training, with probability at least* $1 - 2\delta$*,*

$$\begin{aligned}&\|\hat{f}_{\mathrm{ATT}}^{(\varphi)} - f_{\mathrm{ATT}}\|_{L^2(\mathbb{P}_A \otimes \mathbb{P}_A; \mathbb{R}^{d_Y})}^2 \\ &\quad \leq 2\kappa_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} + 4\Delta_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT},U}(\delta) + 2\xi_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} + 2C_h^2\mathcal{R}_{\varphi,\mathrm{ATT}}^{\mathrm{proj}}.\end{aligned}$$

The multiplier 2 *arises from the same inequality* $\|u+v\|^2 \leq 2\|u\|^2 + 2\|v\|^2$ *step applies to* $\kappa_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}$*, and* $\mathcal{R}_{\varphi,\mathrm{ATT}}^{\mathrm{proj}}$ *already carries the bridge-stage gaps of Theorem H.36. Intersecting this event with Theorem H.36 gives an end-to-end event of probability at least* $1 - 6\delta$*.*

Proof. Decompose the fitted surface around $\bar{f}_{\hat{\varphi}^{\mathrm{ATT}}}^{(\varphi,\mathrm{ATT})}$, apply $\|u + v\|_2^2 \leq 2\|u\|_2^2 + 2\|v\|_2^2$, and use Lemma H.39 and Equation 82. □

Let $\delta_t \downarrow 0$ satisfy $\log(1/\delta_t)/t_{\mathrm{ATT}} \to 0$, and define

$$\rho_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} := \kappa_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} + 2\Delta_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT},U}(\delta_t).$$

If $t_{\mathrm{ATT}} \to \infty$, $\kappa_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} = o_p(1)$, and

$$\widehat{\mathfrak{R}}_{D_{3,\mathrm{ATT}}^{(\varphi)}}^{(2),U}\left(\mathcal{L}_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}}(\hat{\varphi}^{\mathrm{ATT}})\right) = o_p(1),$$

then $\rho_{\varphi,3,t_{\mathrm{ATT}}}^{\mathrm{ATT}} = o_p(1)$.

Define the optimization-aware final-stage rate

$$\tilde{\rho}^{\mathrm{ATT}}_{\varphi,3,t_{\mathrm{ATT}}} := \rho^{\mathrm{ATT}}_{\varphi,3,t_{\mathrm{ATT}}} + \xi^{\mathrm{ATT}}_{\varphi,3,t_{\mathrm{ATT}}}.$$

If

$$\tilde{\rho}^{\mathrm{ATT}}_{\varphi,3,t_{\mathrm{ATT}}} = o_p(1), \qquad \tilde{\rho}_{\varphi,\mathrm{ATT},\nu_\varphi} = o_p(1),$$

then Corollary H.37 gives

$$\left\| \hat{f}^{(\varphi)}_{\mathrm{ATT}} - f_{\mathrm{ATT}} \right\|^2_{L^2(\mathbb{P}_A \otimes \mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p\big(\tilde{\rho}^{\mathrm{ATT}}_{\varphi,3,t_{\mathrm{ATT}}} + \tilde{\rho}_{\varphi,\mathrm{ATT},\nu_\varphi}\big) = o_p(1).$$

The bridge-stage gaps are already contained in $\tilde{\rho}_{\varphi,\mathrm{ATT},\nu_\varphi}$ and need not be required separately.

For a finite treatment space, let

$$N_{a,a'} := \sum_{i \neq j} \mathbb{1}\{a_i = a, a_j = a'\}.$$

The saturated final-stage estimator is exactly

$$\hat{f}^{(\varphi)}_{\mathrm{ATT}}(a,a') = \frac{1}{N_{a,a'}} \sum_{i \neq j} \mathbb{1}\{a_i = a, a_j = a'\} \gamma^{\mathrm{ATT},\varphi}_{ij}, \qquad N_{a,a'} > 0.$$

H.10 Conditional dose-response consistency of the doubly robust method

Condition on the data used to fit the vector-valued outcome bridge $\hat{h}$ and the treatment bridge $\hat{\varphi}^{\mathrm{ATT}}$. The two bridge estimates need not be independent of one another, but the final-stage sample below is independent of both, including their fitted preprocessing transformations. Assume that all displayed products and conditional expectations are well-defined and integrable. For two independent random copies O_i and O_j, $i \neq j$, define

$$\Gamma^{\mathrm{ATT,DR}}_{ij} := \hat{\varphi}^{\mathrm{ATT}}(A_i, A_j, X_i, Z_i)\{Y_i - \hat{h}(A_i, X_i, W_i)\} + \hat{h}(A_i, X_j, W_j). \tag{83}$$

$$\bar{f}^{\mathrm{ATT,DR}}_{\hat{h},\hat{\varphi}^{\mathrm{ATT}}}(a,a') := \mathbb{E}[\Gamma^{\mathrm{ATT,DR}}_{12} \mid A_1 = a, A_2 = a']. \tag{84}$$

Thus, one regression of the paired pseudo-outcome on (A_i, A_j) estimates the full conditional dose-response surface.

Lemma H.41 (Doubly robust remainder)**.** *Suppose valid bridges h_0 and φ^{ATT}_0 exist. Then, for $\mathbb{P}_A \otimes \mathbb{P}_A$-almost every (a,a'),*

$$\begin{aligned} &f_{\mathrm{ATT}}(a,a') - \bar{f}^{\mathrm{ATT,DR}}_{\hat{h},\hat{\varphi}^{\mathrm{ATT}}}(a,a') \\ &\quad = \mathbb{E}\Big[\{\varphi^{\mathrm{ATT}}_0(a,a',X,Z) - \hat{\varphi}^{\mathrm{ATT}}(a,a',X,Z)\}\{h_0(a,X,W) - \hat{h}(a,X,W)\} \mid A = a\Big]. \end{aligned} \tag{85}$$

Consequently, the population surface in Equation 84 equals f_{ATT} whenever either fitted bridge equals a valid bridge.

Proof. The outcome-bridge equation implies

$$\mathbb{E}[\hat{\varphi}^{\mathrm{ATT}}(a,a',X,Z)\{Y - h_0(a,X,W)\} \mid A = a] = 0.$$

Moreover, the joint-anchor treatment bridge transports every square-integrable $q(a,X,W)$ according to

$$\mathbb{E}[\varphi^{\mathrm{ATT}}_0(a,a',X,Z)q(a,X,W) \mid A = a] = \mathbb{E}[q(a,X,W) \mid A = a'].$$

Apply the latter identity to $q = h_0 - \hat{h}$, expand Equations 83–84, and cancel terms. □

Retain $B_{e,h}$ and $\mathcal{R}^{\mathrm{proj}}_h$ from Appendix H.6, and $\mathcal{R}^{\mathrm{proj}}_{\varphi,\mathrm{ATT}}$ from Theorem H.36. Define

$$B^2_{e,\varphi,\mathrm{ATT}} := \operatorname*{ess\,sup}_{(a,a') \sim \mathbb{P}_A \otimes \mathbb{P}_A} \mathbb{E}\big[\{\hat{\varphi}^{\mathrm{ATT}}(a,a',X,Z) - \varphi^{\mathrm{ATT}}_0(a,a',X,Z)\}^2 \mid A = a\big].$$

Lemma H.42 (Projected joint-anchor doubly robust bound)**.** *If $B_{e,\varphi,\mathrm{ATT}}, B_{e,h} < \infty$, then*

$$\|\bar{f}^{\mathrm{ATT,DR}}_{\hat{h},\hat{\varphi}^{\mathrm{ATT}}} - f_{\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_A\otimes\mathbb{P}_A;\mathbb{R}^{d_Y})} \le \min\left\{B^2_{e,\varphi,\mathrm{ATT}}\mathcal{R}^{\mathrm{proj}}_h,\ B^2_{e,h}\mathcal{R}^{\mathrm{proj}}_{\varphi,\mathrm{ATT}}\right\}. \tag{86}$$

Proof. Condition the remainder in Equation 85 first on (X, Z). The conditional mean of $h_0 - \hat{h}$ is $m_0 - T_h\hat{h}$. Conditional Cauchy–Schwarz, followed by integration over (a, a'), gives the first bound. Alternatively, condition first on (X, W). The conditional mean of $\varphi_0^{\mathrm{ATT}} - \hat{\varphi}^{\mathrm{ATT}}$ is $r_{0,\mathrm{ATT}} - T_{\mathrm{ATT}}\hat{\varphi}^{\mathrm{ATT}}$, which gives the second bound. Take their minimum. □

Let $D^{\mathrm{DR}}_{3,\mathrm{ATT}} = \{o_i\}_{i=1}^{t_{\mathrm{ATT,DR}}}$ be the independent final-stage sample. Let

$$\mathcal{F}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} \subset \{f : \mathcal{A}\times\mathcal{A} \to \mathbb{R}^{d_Y}\}.$$

For $o = (a, x, w, z, y)$ and $o' = (a', x', w', z', y')$, set

$$\gamma^{\mathrm{ATT}}_{\mathrm{DR}}(o, o') := \hat{\varphi}^{\mathrm{ATT}}(a, a', x, z)\{y - \hat{h}(a, x, w)\} + \hat{h}(a, x', w').$$

Define

$$\kappa^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} := \inf_{f\in\mathcal{F}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}} \|f - \bar{f}^{\mathrm{ATT,DR}}_{\hat{h},\hat{\varphi}^{\mathrm{ATT}}}\|^2_{L^2(\mathbb{P}_A\otimes\mathbb{P}_A;\mathbb{R}^{d_Y})},$$

$$\mathcal{L}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}(\hat{h}, \hat{\varphi}^{\mathrm{ATT}}) := \left\{(o_i, o_j) \mapsto \|\gamma^{\mathrm{ATT}}_{\mathrm{DR}}(o_i, o_j) - f(a_i, a_j)\|_2^2 : f \in \mathcal{F}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}\right\}.$$

Conditional on bridge training, assume this scalar loss class is measurable, separable, and bounded by a deterministic envelope $M^{\mathrm{ATT}}_{\mathrm{DR},3} < \infty$. Define

$$\Delta^{\mathrm{ATT},U}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}(\delta) := 2\widehat{\mathfrak{R}}^{(2),U}_{D^{\mathrm{DR}}_{3,\mathrm{ATT}}}\left(\mathcal{L}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}(\hat{h}, \hat{\varphi}^{\mathrm{ATT}})\right) + 3M^{\mathrm{ATT}}_{\mathrm{DR},3}\sqrt{\frac{8\log(2/\delta)}{t_{\mathrm{ATT,DR}}}}.$$

Let $\hat{f}^{\mathrm{DR}}_{\mathrm{ATT}}$ be a measurable approximate minimizer of

$$\widehat{\mathcal{R}}^{\mathrm{ATT}}_{\mathrm{DR},3}(f) := \frac{1}{t_{\mathrm{ATT,DR}}(t_{\mathrm{ATT,DR}} - 1)}\sum_{i\neq j}\|\gamma^{\mathrm{ATT}}_{\mathrm{DR}}(o_i, o_j) - f(a_i, a_j)\|_2^2$$

over $\mathcal{F}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}$, with gap $\xi^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} \ge 0$.

Theorem H.43 (Doubly robust conditional dose-response error)**.** *Under the preceding conditions, for $\delta \in (0, 1/2)$, conditional on bridge training, with probability at least $1 - 2\delta$,*

$$\begin{aligned}\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATT}} - f_{\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_A\otimes\mathbb{P}_A;\mathbb{R}^{d_Y})} &\le 2\kappa^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} + 4\Delta^{\mathrm{ATT},U}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}(\delta) + 2\xi^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}\\ &\quad + 2\min\left\{B^2_{e,\varphi,\mathrm{ATT}}\mathcal{R}^{\mathrm{proj}}_h,\ B^2_{e,h}\mathcal{R}^{\mathrm{proj}}_{\varphi,\mathrm{ATT}}\right\}.\end{aligned}$$

Proof. Under the law of two independent observations, $\bar{f}^{\mathrm{ATT,DR}}_{\hat{h},\hat{\varphi}^{\mathrm{ATT}}}(A_1, A_2)$ is the conditional mean of $\Gamma^{\mathrm{ATT,DR}}_{12}$. The vector-valued least-squares projection identity, both directions of Proposition H.5, and approximate empirical optimality give

$$\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATT}} - \bar{f}^{\mathrm{ATT,DR}}_{\hat{h},\hat{\varphi}^{\mathrm{ATT}}}\|^2_{L^2(\mathbb{P}_A\otimes\mathbb{P}_A;\mathbb{R}^{d_Y})} \le \kappa^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} + 2\Delta^{\mathrm{ATT},U}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}(\delta) + \xi^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}.$$

Decompose around the population surface, apply $\|u + v\|_2^2 \le 2\|u\|_2^2 + 2\|v\|_2^2$, and use Lemma H.42. □

Choose $\delta_t \downarrow 0$ such that $\log(1/\delta_t)/t_{\mathrm{ATT,DR}} \to 0$, and define

$$\rho^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} := \kappa^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} + 2\Delta^{\mathrm{ATT},U}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}(\delta_t).$$

If $t_{\mathrm{ATT,DR}} \to \infty$, $M^{\mathrm{ATT}}_{\mathrm{DR},3} = O_p(1)$, $\kappa^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} = o_p(1)$, and

$$\widehat{\mathfrak{R}}^{(2),U}_{D^{\mathrm{DR}}_{3,\mathrm{ATT}}}\left(\mathcal{L}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}(\hat{h}, \hat{\varphi}^{\mathrm{ATT}})\right) = o_p(1),$$

then $\rho^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} = o_p(1)$.

Define the optimization-aware final-stage rate

$$\tilde{\rho}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} := \rho^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} + \xi^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}}.$$

If, in addition,

$$B_{e,\varphi,\mathrm{ATT}}, B_{e,h} = O_p(1),$$
$$\mathcal{R}^{\mathrm{proj}}_h = O_p(\tilde{\rho}_{h,\nu_h}), \qquad \mathcal{R}^{\mathrm{proj}}_{\varphi,\mathrm{ATT}} = O_p(\tilde{\rho}_{\varphi,\mathrm{ATT},\nu_\varphi}),$$

then

$$\|\hat{f}^{\mathrm{DR}}_{\mathrm{ATT}} - f_{\mathrm{ATT}}\|^2_{L^2(\mathbb{P}_A\otimes\mathbb{P}_A;\mathbb{R}^{d_Y})} = O_p\Big(\tilde{\rho}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} + \min\{\tilde{\rho}_{h,\nu_h}, \tilde{\rho}_{\varphi,\mathrm{ATT},\nu_\varphi}\}\Big). \tag{87}$$

Hence the estimator is consistent when $\tilde{\rho}^{\mathrm{ATT}}_{\mathrm{DR},3,t_{\mathrm{ATT,DR}}} = o_p(1)$ and either projected bridge error vanishes. There is no additive density-ratio estimation error.

Remark H.44 (Optimization accuracy)**.** *Exact empirical minimization is the special case in which all ξ-terms equal zero, and every bound above then reduces to its exact-minimizer form. Consistency requires only that the optimization gaps vanish together with the approximation and empirical-process terms of their own stage. If a stage has statistical scale $\rho_{b,s}$, then $\xi_{b,s} = O_p(\rho_{b,s})$ preserves the overall order of that stage, while $\xi_{b,s} = o_p(\rho_{b,s})$ makes the optimization error asymptotically negligible; no universal $O_p(n^{-1})$ optimization requirement is imposed, since the familiar n^{-1} criterion condition applies to parameter-distance targets under quadratic curvature, whereas the quantities bounded here are already squared projected errors. These results do not establish a rate for the specific alternating neural optimizer used in the experiments. They make explicit what empirical suboptimality that optimizer would have to achieve for the statistical conclusions to hold.*

I Supplementary on numerical experiments

I.1 Cross-fitting

In our implementation, we utilize cross-fitting that randomly partitions the observation indices into three disjoint sets $\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3$. For rotation $r \in \{1, 2, 3\}$, with subscripts interpreted modulo three, we use $\mathcal{I}_r$ to learn the first-stage conditional-mean features, $\mathcal{I}_{r+1}$ to estimate the second-stage outcome and treatment bridge heads, and $\mathcal{I}_{r+2}$ to compute the corresponding final-stage causal estimate $\hat{f}^{(r)}$. Thus each split serves once at each stage, while the bridge estimators used in a given final stage are trained without its observations. We report the cross-fitted estimate $\hat{f}^{\mathrm{cf}} := 3^{-1}\sum_{r=1}^{3}\hat{f}^{(r)}$.

I.2 Vectorized computation of U-statistics with einsum

Our TreatmentNet algorithm requires computing U-statistics. In this section, we provide a vectorized computation of this loss component with einsum. To this end, write the compact feature maps in the following separable form

$$\phi^{(\varphi)}_{AXW,1} = \phi^{(\varphi)}_{A,1} \otimes \phi^{(\varphi)}_{XW,1}$$
$$\phi^{(\varphi)}_{AX,2} = \phi^{(\varphi)}_{A,2} \otimes \phi^{(\varphi)}_{X,2}.$$

Fix $\check{\boldsymbol{V}}^{(\varphi)}_t$, let $m := |\mathcal{B}^{(\varphi)}_{2,t}| \geq 2$, and suppress (t, φ) on $\boldsymbol{u}_{ij}$, b_{prod}, b_U, and $\bar{\boldsymbol{u}}$ below.

Recall we defined

$$\boldsymbol{u}^{(\varphi)}_{ij,t} = \phi^{(\varphi)}_{A,2}(a_i) \otimes \phi^{(\varphi)}_{X,2}(x_j) \otimes \left\{\check{\boldsymbol{V}}^{(\varphi)}_t\right\}^\top \left\{\phi^{(\varphi)}_{A,1}(a_i) \otimes \phi^{(\varphi)}_{XW,1}(x_j, w_j)\right\}.$$

The U-statistic term in the ratio-free risk is

$$\frac{1}{m(m-1)} \sum_{\substack{i,j \in \mathcal{B}_{2,t}^{(\varphi)} \\ i \neq j}} \boldsymbol{\varphi}^\top \boldsymbol{u}_{ij,t}^{(\varphi)}.$$

Define the tensor-product features $x_{ij} := \phi_{A,2}^{(\varphi)}(a_i) \otimes \phi_{X,2}^{(\varphi)}(x_j)$, $y_{ij} := \phi_{A,1}^{(\varphi)}(a_i) \otimes \phi_{XW,1}^{(\varphi)}(x_j, w_j)$. The operator identity underlying the implementation is

$$\begin{aligned} \boldsymbol{u}_{ij} &= x_{ij} \otimes \{\check{\boldsymbol{V}}_t^{(\varphi)}\}^\top y_{ij} \\ &= \left(I \otimes \{\check{\boldsymbol{V}}_t^{(\varphi)}\}^\top\right)(x_{ij} \otimes y_{ij}). \end{aligned}$$

Here I is the identity on the $A_2 \otimes X_2$ feature space. A fixed commutation permutation groups the i-only and j-only factors. Reshape $\{\check{\boldsymbol{V}}_t^{(\varphi)}\}^\top$ as $\mathsf{V}_t^{(\varphi)}[z, p, r]$, where z, p, r index the Z_2, A_1, XW_1 modes, respectively, and define

$$\begin{aligned} M_A[p, c] &= \frac{1}{m} \sum_i \phi_{A,1}^{(\varphi)}(a_i)[p] \phi_{A,2}^{(\varphi)}(a_i)[c] \\ M_{XW}[r, d] &= \frac{1}{m} \sum_j \phi_{XW,1}^{(\varphi)}(x_j, w_j)[r] \phi_{X,2}^{(\varphi)}(x_j)[d]. \end{aligned}$$

Then

$$b_{\mathrm{prod}}[c, d, z] := \frac{1}{m^2} \sum_{i,j} \boldsymbol{u}_{ij}[c, d, z] = \sum_{p,r} \mathsf{V}_t^{(\varphi)}[z, p, r] M_A[p, c] M_{XW}[r, d] \tag{88}$$

Equation 88 is computed as

```
einsum('zpr,pc,rd->zcd', V, M_A, M_XW).
```

Finally, let $\bar{\boldsymbol{u}} = m^{-1} \sum_i \boldsymbol{u}_{ii}$. Removing the diagonal turns the empirical-product V-statistic into the complete ordered U-statistic:

$$b_U \coloneqq \frac{1}{m(m-1)} \sum_{i \neq j} \boldsymbol{u}_{ij} = \frac{m b_{\mathrm{prod}} - \bar{\boldsymbol{u}}}{m-1} \tag{89}$$

Hence, the linear term in Equation 51 is $-2\boldsymbol{\varphi}^\top b_U$. Equations 88 and 89 are exact and never construct an $m \times m$ pair array. For fixed feature widths, this contraction has linear sample-axis memory and computation; feature-mode contractions and the bridge-head Gram matrix retain their feature-dimension costs.

I.3 Treatment bridge normalization

Proposition I.1 (Treatment-bridge normalization)**.** *Suppose that φ is integrable and satisfies Eq. 2. Then*

$$\mathbb{E}\{\varphi(A, X, Z) \mid A\} = 1 \qquad \mathbb{P}_A\text{-almost surely}. \tag{90}$$

Proof. For $\mathbb{P}_A$-almost every a, the law of total expectation and the bridge equation give

$$\begin{aligned} \mathbb{E}\{\varphi(a, X, Z) \mid A = a\} &= \int r_{0,\mathrm{ATE}}(a, x, w)\, d\mathbb{P}_{X,W|A=a}(x, w) \\ &= \int \frac{d\mathbb{P}_{X,W}}{d\mathbb{P}_{X,W|A=a}}(x, w)\, d\mathbb{P}_{X,W|A=a}(x, w) = 1, \end{aligned}$$

where the second equality uses the same domination condition as the treatment-bridge equation. □

Let $T_\varphi \varphi(A, X, W) := \mathbb{E}\{\varphi(A, X, Z) \mid A, X, W\}$. Whenever the right-hand side below is finite, the same argument and Jensen's inequality yields

$$\|\mathbb{E}\{\varphi(A, X, Z) \mid A\} - 1\|_{L^2(\mathbb{P}_A)} \leq \|T_\varphi \varphi - r_{0,\mathrm{ATE}}\|_{L^2(\mathbb{P}_{A,X,W})}. \tag{91}$$

Thus projected treatment-bridge consistency already implies asymptotic normalization. The constraint below is a finite-sample stabilization.

Finite moment restrictions. If $\varphi \in L^2(\mathbb{P}_{A,X,Z})$, then for every $g \in L^2(\mathbb{P}_A)$,

$$\mathbb{E}[(\varphi - 1)g(A)] = \langle \mathbb{E}(\varphi \mid A) - 1, g \rangle_{L^2(\mathbb{P}_A)}. \tag{92}$$

Consequently, Eq. 90 is equivalent to the vanishing of these moments over any class whose linear span is dense in $L^2(\mathbb{P}_A)$. A finite collection only imposes the corresponding projected restrictions.

Choose $g_1, \ldots, g_L$ and let $\mathcal{B}_{2,t}^{(\varphi)}$ contain m observations. Write the direct bridge feature and its linear head as

$$\boldsymbol{q}_i^{(\varphi)} := \phi_{AX,2}^{(\varphi)}(a_i, x_i) \otimes \phi_{Z,2}^{(\varphi)}(z_i) \in \mathbb{R}^D, \qquad \varphi_{\boldsymbol{\beta}}(a_i, x_i, z_i) := \boldsymbol{\beta}^\top \boldsymbol{q}_i^{(\varphi)}.$$

Define $\boldsymbol{C}_t \in \mathbb{R}^{L \times D}$ and $\boldsymbol{d}_t \in \mathbb{R}^L$ by

$$(\boldsymbol{C}_t)_{\ell,:} := \frac{1}{m} \sum_{i \in \mathcal{B}_{2,t}^{(\varphi)}} g_\ell(a_i) \boldsymbol{q}_i^{(\varphi)\top}, \qquad (\boldsymbol{d}_t)_\ell := \frac{1}{m} \sum_{i \in \mathcal{B}_{2,t}^{(\varphi)}} g_\ell(a_i). \tag{93}$$

The L empirical normalization moments are exactly $\boldsymbol{C}_t \boldsymbol{\beta} = \boldsymbol{d}_t$. To obtain the constrained head update equation, set

$$\boldsymbol{H}_t := \boldsymbol{G}_t^{(\varphi)} + \lambda_2^{(\varphi)} \boldsymbol{I}_D, \qquad \boldsymbol{r}_t := \boldsymbol{b}_{U,t}^{(\varphi)} + \lambda_2^{(\varphi)} \hat{\boldsymbol{\varphi}}_t.$$

For $\lambda_2^{(\varphi)} > 0$, $\boldsymbol{H}_t$ is positive definite. The constrained optimization problem for the normalized head is

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^D} \left\{ \frac{1}{2} \boldsymbol{\beta}^\top \boldsymbol{H}_t \boldsymbol{\beta} - \boldsymbol{r}_t^\top \boldsymbol{\beta} \right\} \quad \text{subject to} \quad \boldsymbol{C}_t \boldsymbol{\beta} = \boldsymbol{d}_t. \tag{94}$$

Its KKT conditions are

$$\begin{bmatrix} \boldsymbol{H}_t & \boldsymbol{C}_t^\top \\ \boldsymbol{C}_t & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\beta}^\star \\ \boldsymbol{\eta}^\star \end{bmatrix} = \begin{bmatrix} \boldsymbol{r}_t \\ \boldsymbol{d}_t \end{bmatrix}. \tag{95}$$

If $\boldsymbol{C}_t$ has full row rank, the block system is nonsingular. Let $\boldsymbol{\beta}_{\text{unc}} := \boldsymbol{H}_t^{-1} \boldsymbol{r}_t$. Its unique primal solution is

$$\boldsymbol{\beta}^\star = \boldsymbol{\beta}_{\text{unc}} + \boldsymbol{H}_t^{-1} \boldsymbol{C}_t^\top \left(\boldsymbol{C}_t \boldsymbol{H}_t^{-1} \boldsymbol{C}_t^\top \right)^{-1} (\boldsymbol{d}_t - \boldsymbol{C}_t \boldsymbol{\beta}_{\text{unc}}), \tag{96}$$

and $\hat{\boldsymbol{\varphi}}_{t+1} := \boldsymbol{\beta}^\star$. Hence the KKT update is the minimum-$\boldsymbol{H}_t$-norm correction of the unconstrained head that satisfies the selected moments. The implementation solves Eq. 95 directly; it does not form either inverse in Eq. 96. In implementation, we evaluate the bridge features without dropout, orthonormalize the empirical treatment basis, and discard numerically redundant basis directions.

Table 3 isolates the effect of the KKT update on TREATMENTNET. Here, None uses the unconstrained head update, whereas KKT enforces the empirical normalization moments in Eq. 93 during training. For ATT, we use the corresponding tensor-product moments in (a, a').

Table 3: TREATMENTNET normalization ablation. Entries are causal MSE (lower is better), reported as mean ± seed-level standard error over ten data-generating seeds after averaging three repetitions within each seed.

Benchmark	Normalization	$n = 2$k	$n = 5$k	$n = 10$k	$n = 15$k	$n = 20$k
Demand	None	367.42 ± 55.39	374.43 ± 28.81	371.06 ± 37.03	383.51 ± 29.58	350.96 ± 31.13
	KKT	$\mathbf{199.67 \pm 15.77}$	$\mathbf{174.11 \pm 16.50}$	$\mathbf{153.84 \pm 5.57}$	$\mathbf{144.97 \pm 4.24}$	$\mathbf{145.76 \pm 4.40}$
Synthetic ATT ($10^3\times$)	None	$\mathbf{142.83 \pm 7.88}$	128.02 ± 13.47	142.15 ± 8.00	130.36 ± 14.43	131.01 ± 8.86
	KKT	179.49 ± 8.90	$\mathbf{103.68 \pm 9.55}$	$\mathbf{95.08 \pm 6.96}$	$\mathbf{71.87 \pm 8.37}$	$\mathbf{72.49 \pm 7.61}$

KKT lowers the TREATMENTNET error throughout the Demand sweep and for $n \geq 5{,}000$ in the ATT sweep. At $n = 2{,}000$ in the ATT design, the unconstrained update performs better, consistent with viewing the restriction as a finite-sample stabilization rather than a guarantee of uniform improvement.

I.4 Benchmark descriptions and data-generating processes

In this section, we give the variables, proxy partitions, structural equations, and oracle targets used in the numerical experiments.

I.4.1 Bridge-route preference benchmarks

We use two complementary stress tests to compare the outcome- and treatment-bridge routes. In both designs, we consider binary treatment setting with $A \in \{-1, +1\}$, and we set $d_Y = 3$. For a Rademacher variable R, write $B_\lambda(R) := R\xi_\lambda$, where $\xi_\lambda \in \{-1, +1\}$ with $\mathbb{P}(\xi_\lambda = 1) = (1 + \lambda)/2$. Hence $\mathbb{E}\{B_\lambda(R) \mid R\} = \lambda R$.

Scenario A: treatment-route stress. Let $U = (U_1, U_2)$, where U_1 and U_2 are independent Rademacher variables. Define

$$\eta(u_1, u_2) := \kappa_1 u_1 + \kappa_2 u_2 + \kappa_{12} u_1 u_2, \qquad \mathbb{P}(A = a \mid U_1 = u_1, U_2 = u_2) = \frac{1 + a\eta(u_1, u_2)}{2}.$$

with the restriction $|\kappa_1| + |\kappa_2| + |\kappa_{12}| < 1$. We define the proxy variables and potential outcomes as

$$\begin{aligned} Z &= \big(B_{\lambda_{Z_1}}(U_1), B_{\lambda_{Z_2}}(U_2)\big), & W &= \big(B_{\lambda_{W_1}}(U_1), B_{\lambda_{W_2}}(U_2)\big), \\ Y^{(a)} &= a\mathbf{1}_3 + (U_1, U_2, U_1U_2)^\top + \varepsilon, & \varepsilon &\sim \mathcal{N}(0, 0.25^2 I_3). \end{aligned}$$

By Bayes' rule, $\mathbb{P}(U_1 = u_1, U_2 = u_2 \mid A = a) = \frac{1+a\eta(u_1,u_2)}{4}$. Consequently, we have $\mathbb{E}[Y \mid A = a] = a\mathbf{1}_3 + a(\kappa_1,\ \kappa_2,\ \kappa_{12})^\top$ and $f_{\mathrm{ATE}}(a) = a\mathbf{1}_3$. Then, an outcome bridge is given by

$$h_0(a, w_1, w_2) = a\mathbf{1}_3 + \left(\frac{w_1}{\lambda_{W_1}}, \frac{w_2}{\lambda_{W_2}}, \frac{w_1 w_2}{\lambda_{W_1}\lambda_{W_2}}\right)^\top. \tag{97}$$

Indeed, since $\mathbb{E}\left[\frac{W_1}{\lambda_{W_1}} \mid U_1, U_2\right] = U_1$, $\mathbb{E}\left[\frac{W_2}{\lambda_{W_2}} \mid U_1, U_2\right] = U_2$, and $\mathbb{E}\left[\frac{W_1W_2}{\lambda_{W_1}\lambda_{W_2}} \mid U_1, U_2\right] = U_1U_2$, we have

$$\mathbb{E}\left[h_0(a, W) \mid U_1, U_2\right] = a\mathbf{1}_3 + (U_1, U_2, U_1U_2)^\top := m_a(U_1, U_2).$$

Therefore, by the law of total expectations,

$$\begin{aligned} \mathbb{E}[h_0(a, W) \mid A = a, Z] &= \mathbb{E}[\mathbb{E}[h_0(a, W) \mid A = a, Z, U_1, U_2] \mid A = a, Z] \\ &= \mathbb{E}[\mathbb{E}[h_0(a, W) \mid U_1, U_2] \mid A = a, Z] \ \text{ by Assumption 2.1} \\ &= \mathbb{E}[m_a(U_1, U_2) \mid A = a, Z] = \mathbb{E}[\mathbb{E}[Y \mid A = a, Z, U_1, U_2] \mid A = a, Z] \\ &= \mathbb{E}[Y \mid A = a, Z]. \end{aligned}$$

To derive a treatment bridge, define the latent reweighting

$$r_a(u_1, u_2) := \frac{\mathbb{P}(A = a)}{\mathbb{P}(A = a \mid U_1 = u_1, U_2 = u_2)} = \frac{1}{1 + a\eta(u_1, u_2)}, \qquad u_1, u_2 \in \{-1, +1\},$$

Since (U_1, U_2) takes only four possible values, every function of (U_1, U_2) can be written uniquely as

$$r_a(u_1, u_2) = d_0(a) + d_1(a)u_1 + d_2(a)u_2 + d_{12}(a)u_1u_2.$$

The four coefficients are

$$\begin{aligned} d_0(a) &:= \frac{1}{4}\sum_{u_1,u_2\in\{-1,+1\}} r_a(u_1, u_2), \quad d_1(a) := \frac{1}{4}\sum_{u_1,u_2\in\{-1,+1\}} u_1\, r_a(u_1, u_2), \\ d_2(a) &:= \frac{1}{4}\sum_{u_1,u_2\in\{-1,+1\}} u_2\, r_a(u_1, u_2), \quad d_{12}(a) := \frac{1}{4}\sum_{u_1,u_2\in\{-1,+1\}} u_1u_2\, r_a(u_1, u_2). \end{aligned}$$

Then, a treatment bridge is given by

$$\varphi_0^{\mathrm{ATE}}(a, z_1, z_2) = d_0(a) + \frac{d_1(a)}{\lambda_{Z_1}} z_1 + \frac{d_2(a)}{\lambda_{Z_2}} z_2 + \frac{d_{12}(a)}{\lambda_{Z_1}\lambda_{Z_2}} z_1 z_2. \tag{98}$$

Indeed, conditional independence of the two proxy-noise variables gives

$$\mathbb{E}[Z_1 \mid U_1, U_2] = \lambda_{Z_1} U_1, \qquad \mathbb{E}[Z_2 \mid U_1, U_2] = \lambda_{Z_2} U_2, \qquad \mathbb{E}[Z_1 Z_2 \mid U_1, U_2] = \lambda_{Z_1}\lambda_{Z_2} U_1 U_2.$$

Therefore,

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, Z_1, Z_2) \mid U_1, U_2] = d_0(a) + d_1(a)U_1 + d_2(a)U_2 + d_{12}(a)U_1U_2 = r_a(U_1, U_2).$$

Consequently, by iterated expectation,

$$\begin{aligned}\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, Z_1, Z_2) \mid A = a, W] &= \mathbb{E}[\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, Z_1, Z_2) \mid A = a, W, U_1, U_2] \mid A = a, W] \\ &= \mathbb{E}[r_a(U_1, U_2) \mid A = a, W] \\ &= \frac{\mathbb{P}(A = a)\mathbb{P}(W)}{\mathbb{P}(A = a, W)}.\end{aligned}$$

The second equality uses the conditional independence $Z \perp (A, W) \mid (U_1, U_2)$ from Assumption 2.1. The third equality uses Bayes' rule and $A \perp W \mid (U_1, U_2)$ from Assumption 2.1.

In our simulations, we use

$$\kappa_1 = \kappa_{12} = 0.1, \quad \kappa_2 \in \{0, 0.2, 0.4, 0.6\}, \quad (\lambda_{Z_1}, \lambda_{Z_2}, \lambda_{W_1}, \lambda_{W_2}) = (0.8, 0.6, 0.8, 0.8).$$

In this data generating process, the proxy channels and the exact outcome bridge are fixed, while increasing κ_2 strengthens latent treatment selection through U_2. In particular, the smallest latent-stratum treatment probability decreases from 0.4 at $\kappa_2 = 0$ to 0.1 at $\kappa_2 = 0.6$. Consequently, the latent reweighting function

$$r_a(u_1, u_2) = \frac{1}{1 + a\eta(u_1, u_2)}$$

becomes more dispersed, which increases the magnitude and second moment of the exact treatment bridge. Scenario A therefore stresses the weighting route while leaving the exact outcome bridge unchanged. The conditional operator also becomes moderately less well conditioned over the sweep, but remains injective.

Scenario B: outcome-route stress. Again let U_1 and U_2 be independent Rademacher variables, but now set

$$\mathbb{P}(A = a \mid U_1 = u_1) = \frac{1 + 0.1au_1}{2}.$$

The proxy channels retain the same form as in Scenario A, with

$$(\lambda_{Z_1}, \lambda_{W_1}) = (0.9, 0.9), \qquad \lambda_{Z_2}\lambda_{W_2} = 0.36,$$

and potential outcomes are defined by

$$Y^{(a)} = a\mathbf{1}_3 + (0.25U_1, U_2, 0.8U_1U_2)^\top + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, 0.25^2 I_3).$$

This construction also yields $f_{\mathrm{ATE}}(a) = a\mathbf{1}_3$. The exact bridges are given by

$$h_0(a, w_1, w_2) = a\mathbf{1}_3 + \left(\frac{0.25w_1}{0.9}, \frac{w_2}{\lambda_{W_2}}, \frac{0.8w_1w_2}{0.9\lambda_{W_2}}\right)^\top, \tag{99}$$

$$\varphi_0^{\mathrm{ATE}}(a, z_1, z_2) = \frac{1 - 0.1az_1/0.9}{1 - 0.1^2}. \tag{100}$$

The outcome-bridge derivation is the same as in Scenario A, up to the scaling coefficients. Indeed,

$$\mathbb{E}[W_1 \mid U_1, U_2] = 0.9U_1, \qquad \mathbb{E}[W_2 \mid U_1, U_2] = \lambda_{W_2} U_2, \qquad \mathbb{E}[W_1W_2 \mid U_1, U_2] = 0.9\lambda_{W_2} U_1U_2.$$

Consequently,

$$\mathbb{E}[h_0(a, W_1, W_2) \mid U_1, U_2] = a\mathbf{1}_3 + (0.25U_1, U_2, 0.8U_1U_2)^\top = \mathbb{E}[Y \mid A = a, U_1, U_2].$$

Using $W \perp (A, Z) \mid (U_1, U_2)$ and iterated expectation therefore yields

$$\mathbb{E}[h_0(a, W) \mid A = a, Z] = \mathbb{E}[Y \mid A = a, Z].$$

For treatment bridge function, observe that

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, Z_1, Z_2) \mid U_1, U_2] = \frac{1 - 0.1a\mathbb{E}[Z_1 \mid U_1, U_2]/0.9}{1 - 0.1^2} = \frac{1 - 0.1aU_1}{1 - 0.1^2} = \frac{\mathbb{P}(A = a)}{\mathbb{P}(A = a \mid U_1)}.$$

Let $r_a(u_1)$ be $\mathbb{P}(A = a)/\mathbb{P}(A = a \mid U_1 = u_1)$. Then, law of total expectations and Assumption 2.1 give

$$\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, Z) \mid A = a, W] = \mathbb{E}[\mathbb{E}[\varphi_0^{\mathrm{ATE}}(a, Z) \mid A = a, W, U_1, U_2] \mid A = a, W] = \mathbb{E}[r_a(U_1) \mid A = a, W] = \frac{\mathbb{P}(A = a)}{\mathbb{P}(A = a \mid W)}.$$

We sweep $\lambda_{W_2} \in \{0.9, 0.7, 0.6, 0.5, 0.4\}$, $\lambda_{Z_2} = 0.36/\lambda_{W_2}$. In this data generating process, treatment depends only on U_1. Hence U_1 is the latent confounder, whereas U_2 is prognostic but does not affect treatment. The exact treatment bridge depends only on Z_1 and is therefore independent of λ_{Z_2} and λ_{W_2}. Moreover,

$$\mathbb{P}(Z_2 = z_2, W_2 = w_2) = \frac{1 + \lambda_{Z_2}\lambda_{W_2} z_2 w_2}{4}.$$

Keeping $\lambda_{Z_2}\lambda_{W_2} = 0.36$ therefore fixes the joint law of (A, Z, W), and hence fixes the entire treatment-bridge nuisance-learning problem and both conditional-expectation operators. Because $\mathbb{P}\{B_{\lambda_{W_2}}(U_2) \neq U_2\} = (1 - \lambda_{W_2})/2$, decreasing λ_{W_2} makes W_2 a noisier measurement of U_2. The exact outcome bridge then requires coefficients proportional to $1/\lambda_{W_2}$, increasing its variability and sensitivity to estimation error. Scenario B therefore isolates deterioration of the outcome route while leaving the treatment bridge and its nuisance law fixed.

I.4.2 Legalized Abortion and Crime

We use the application-motivated data of Donohue & Levitt (2001), as preprocessed and released by Mastouri et al. (2021). The continuous treatment A is the effective abortion rate, and the scalar outcome Y is the murder rate. The treatment-inducing proxy $Z \in \mathbb{R}$ is the generosity of Aid to Families with Dependent Children. The outcome-inducing proxy $W \in \mathbb{R}^3$ contains beer consumption per capita, log prisoners per capita, and the concealed-weapons-law indicator. The released benchmark also supplies the intervention grid and oracle response used to compute causal MSE.

I.4.3 Grade retention (Deaner IM and IR)

We use the calibrated semi-synthetic ECLS-K grade-retention designs based on Fruehwirth et al. (2016) and Deaner (2023), as released by Mastouri et al. (2021). In the IM and IR variants, Y is the age-11 mathematics or reading score, respectively. The scalar treatment A is a continuously dequantized retention measure. The proxy vectors contain earlier elementary-school and kindergarten cognitive measurements, with $Z \in \mathbb{R}^4$ and $W \in \mathbb{R}^3$ under the released partition. We retain the observational and intervention-grid treatment values in $(-0.1, 2)$, and use the accompanying oracle responses to compute causal MSE.

I.4.4 Demand

We use the nonlinear demand design of Xu et al. (2021b). Draw latent demand $D \sim \mathrm{Unif}[0, 10]$ and define

$$\psi(d) = 2\left\{\frac{(d-5)^4}{600} + \exp[-4(d-5)^2] + \frac{d}{10} - 2\right\}.$$

For mutually independent standard-normal errors, generate fuel costs (C_1, C_2), page views R, ticket price P, and sales Y as

$$C_1 = 2\sin(2\pi D/10) + \varepsilon_1, \qquad C_2 = 2\cos(2\pi D/10) + \varepsilon_2,$$
$$R = 7\psi(D) + 45 + \varepsilon_R, \qquad P = 35 + (C_1 + 3)\psi(D) + C_2 + \varepsilon_P,$$
$$Y = P\min\{\exp[(R - P)/10], 5\} - 5\psi(D) + \varepsilon_Y.$$

The treatment is $A = P$, the treatment-inducing proxy is $Z = (C_1, C_2)$, and the outcome-inducing proxy is $W = R$. The oracle population dose-response is

$$f_{\mathrm{ATE}}(a) = \mathbb{E}[a \min\{\exp[(R-a)/10], 5\} - 5\psi(D)]\,.$$

I.4.5 IMAGE-VALUED TREATMENT BENCHMARK: DSPRITE

We construct a modified dSprite benchmark from (Xu et al., 2021b) that uses images of size 64×64 described by the latent parameters: shape, scale, $\mathrm{rotation}$, posX, and posY (Matthey et al., 2017). We draw a confounder $U \sim \mathrm{Unif}[0,1]$. The shape is fixed to the heart, while the scale, orientation, and vertical-position indices of the treatment image are drawn independently and uniformly from $\{0,\dots,5\}$, $\{0,\dots,39\}$, and $\{0,\dots,31\}$, respectively. Its horizontal-position index is

$$\mathrm{posX}_A = \mathrm{round}[31\ \mathrm{clip}\{U + 0.25\varepsilon_A, 0, 1\}]\,,$$

and the observed treatment is the resulting flattened 64×64 image plus $\eta_A \sim \mathcal{N}(0, 0.1^2 I_{4096})$. The treatment-inducing proxy is

$$Z = (U + 0.40\varepsilon_{Z,1},\ U + \varepsilon_{Z,2}).$$

For the outcome-inducing proxy, we fix the scale, orientation, and horizontal- position indices at $3, 0, 16$, set

$$\mathrm{posY}_W = \mathrm{round}[31\ \mathrm{clip}\{U + 0.30\varepsilon_W, 0, 1\}]\,,$$

and let W be the corresponding flattened image plus $\eta_W \sim \mathcal{N}(0, 0.1^2 I_{4096})$. Here $\varepsilon_A, \varepsilon_{Z,1}, \varepsilon_W \sim \mathcal{N}(0,1)$, $\varepsilon_{Z,2} \sim \mathrm{Unif}[-0.60, 0.60]$, and all proxy, rendering, and pixel noises are mutually independent.

With $B_{ij} = |32 - j|/32$, define

$$f_{\mathrm{ATE}}(a) = \frac{\{\mathrm{vec}(B)^\top a\}^2 - 3000}{500}, \qquad q(u) = \frac{\exp\{3(u - 1/2)\}}{(2/3)\sinh(3/2)}.$$

We generate

$$Y = q(U) f_{\mathrm{ATE}}(A) + \varepsilon_Y, \qquad \varepsilon_Y \sim \mathcal{N}(0, 0.5^2).$$

The outcome noise is independent of all preceding variables. Since $\mathbb{E}\{q(U)\} = 1$, the oracle population dose-response is exactly $f_{\mathrm{ATE}}(a)$.

Original benchmark of Xu et al. (2021b). For comparison, Xu et al. (2021b) draw the latent rendering vector $U = (\mathrm{scale}, \mathrm{rotation}, \mathrm{posX}, \mathrm{posY})$ uniformly over the dSprite factors, with shape fixed to the heart. They set

$$\begin{aligned} A &= \mathrm{vec}\{\mathrm{Fig}(U)\} + \eta_A, & Z &= (\mathrm{scale}, \mathrm{rotation}, \mathrm{posX}), \\ W &= \mathrm{vec}\{\mathrm{Fig}(0.8, 0, 0.5, \mathrm{posY})\} + \eta_W, \end{aligned}$$

where $\eta_A, \eta_W \sim \mathcal{N}(0, 0.1^2 I_{4096})$. Thus, Z directly copies three coordinates of U, while W is a noisy image sharing the same vertical position. With $B_{ij} = |32 - j|/32$, the structural response is

$$f_{\mathrm{ATE}}(a) = \frac{\{\mathrm{vec}(B)^\top a\}^2 - 3000}{500}, \qquad Y = 12(\mathrm{posY} - 1/2)^2 f_{\mathrm{ATE}}(A) + \varepsilon_Y, \quad \varepsilon_Y \sim \mathcal{N}(0, 0.5^2).$$

Figure 3 presents results for this original benchmark.

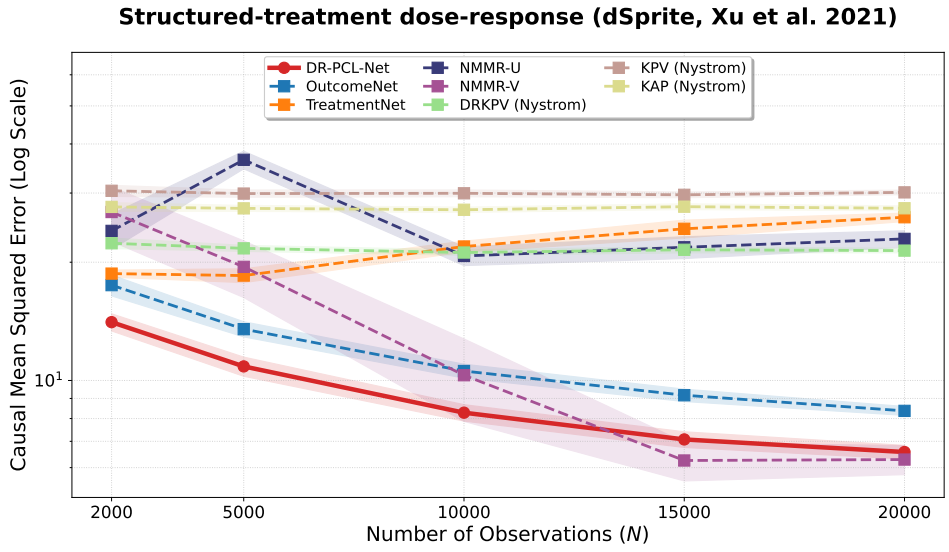

Figure 3: Results on the original dSprite benchmark of Xu et al. (2021b). We report causal MSE over ten data-generating seeds after averaging three repetitions within each seed; shaded regions are 95% confidence intervals across the seed-level means. The vertical axis is logarithmic and lower is better.

I.4.6 Treatment-conditional dose-response benchmark

We use the data-generating process as in the synthetic low-dimensional dose-response benchmark from Wu et al. (2024), but change the estimand to the conditional dose-response curve

$$f_{\mathrm{ATT}}(a, a') := \mathbb{E}[Y^{(a)} \mid A = a']$$

where a' indexes the observed-treatment reference population.

Observational data-generating process: The latent variables are generated as

$$U_1 \sim \mathrm{Unif}[-1, 2], \quad U_2 := R - \mathbf{1}\{0 \le U_1 \le 1\},$$

where $R \sim \mathrm{Unif}[0, 1]$. The observed proxy, treatment, and outcome variables are generated as

$$W_1 = U_2 + \varepsilon_{W,1}, \qquad W_2 = U_1 + \varepsilon_{W,2},$$

$$Z_1 = U_2 + \varepsilon_{Z,1}, \qquad Z_2 = U_1 + \varepsilon_{Z,2},$$

$$A = U_1 + \varepsilon_A,$$

$$Y = 3\cos\big(2(0.3U_2 + 0.3U_1 + 0.2) + 1.5A\big) + \varepsilon_Y,$$

where $\varepsilon_{W,1}, \varepsilon_{Z,2} \sim \mathrm{Unif}[-1, 1]$, $\varepsilon_{W,2}, \varepsilon_{Z,1}, \varepsilon_A, \varepsilon_Y \sim \mathcal{N}(0, 1)$, and all exogenous variables are mutually independent. The treatment proxy is $Z = (Z_1, Z_2)$, the outcome proxy is $W = (W_1, W_2)$, and the latent confounding is driven by (U_1, U_2).

Potential outcomes and the oracle ATT curve: Under the intervention $A = a$, the potential outcome is

$$Y^{(a)} = 3\cos\big(0.6U_2 + 0.6U_1 + 0.4 + 1.5a\big) + \varepsilon_Y.$$

Since $\mathbb{E}[\varepsilon_Y] = 0$, the target ATT curve is

$$f_{\mathrm{ATT}}(a, a') = \mathbb{E}\big[3\cos\big(0.6U_2 + 0.6U_1 + 0.4 + 1.5a\big) \,\big|\, A = a'\big].$$

We first integrate out U_2 conditional on $U_1 = u$. Because $U_2 = R - \mathbf{1}\{0 \le u \le 1\}$ with $R \sim \mathrm{Unif}[0, 1]$, we obtain

$$m(a, u) := \mathbb{E}[Y^{(a)} \mid U_1 = u] = 3\int_0^1 \cos\big(0.6r + c(u, a)\big)\, dr,$$

where $c(u,a) := 0.6u - 0.6\mathbf{1}\{0 \le u \le 1\} + 0.4 + 1.5a$. Evaluating the integral yields

$$m(a,u) = 5\Big[\sin\big(c(u,a)+0.6\big) - \sin\big(c(u,a)\big)\Big].$$

Next, since $A = U_1 + \varepsilon_A$ with $\varepsilon_A \sim \mathcal{N}(0,1)$ and $U_1 \sim \mathrm{Unif}[-1,2]$, Bayes' rule gives the posterior density

$$p(u \mid A = a') = \frac{\phi_{\mathrm{N}}(a'-u)\,\mathbf{1}\{-1 \le u \le 2\}}{\int_{-1}^{2} \phi_{\mathrm{N}}(a'-t)\,dt} = \frac{\phi_{\mathrm{N}}(a'-u)\,\mathbf{1}\{-1 \le u \le 2\}}{\Phi_{\mathrm{N}}(2-a') - \Phi_{\mathrm{N}}(-1-a')},$$

where ϕ_{N} and Φ_{N} denote the standard normal density and distribution functions, respectively.

Therefore, the oracle conditional dose-response curve is

$$f_{\mathrm{ATT}}(a,a') = \int_{-1}^{2} m(a,u)\,p(u \mid A = a')\,du,$$

that is,

$$f_{\mathrm{ATT}}(a,a') = \frac{5}{\Phi_{\mathrm{N}}(2-a') - \Phi_{\mathrm{N}}(-1-a')} \int_{-1}^{2} \Big[\sin\big(c(u,a)+0.6\big) - \sin\big(c(u,a)\big)\Big]\phi_{\mathrm{N}}(a'-u)\,du.$$

This is the ground-truth ATT curve used in the experiments. In practice, we evaluate the one-dimensional integral above numerically on a grid of intervention values a, for each chosen a'.

I.4.7 Heterogeneous dose-response benchmark

We adapt the binary-treatment design of Abrevaya et al. (2015) to the proxy causal learning setting. Draw $\epsilon_0, \epsilon_1, \epsilon_2, \epsilon_3 \overset{\text{iid}}{\sim} \mathrm{Unif}[-0.5, 0.5]$, set $V = \epsilon_0$, and define

$$U_1 = 1 + 2V + \epsilon_1, \qquad U_2 = 1 + 2V + \epsilon_2, \qquad U_3 = (V-1)^2 + \epsilon_3.$$

Treatment is assigned according to

$$\mathbb{P}(A = 1 \mid V, U_1, U_2, U_3) = \mathrm{expit}\big\{\tfrac{1}{2}(V + U_1 + U_2 + U_3)\big\}.$$

With $\nu \sim \mathcal{N}(0, 0.25^2)$, the observed outcome is

$$Y = A\{VU_1U_2U_3 + \nu\}.$$

The proxies are independent noisy measurements of the latent variables:

$$\begin{aligned}
W_1 &= U_1 + \omega_1, & \omega_1 &\sim \mathrm{Unif}[-b,b], & Z_1 &= U_1 + \zeta_1, & \zeta_1 &\sim \mathcal{N}(0,\sigma^2),\\
W_2 &= U_2 + \omega_2, & \omega_2 &\sim \mathcal{N}(0,\sigma^2), & Z_2 &= U_2 + \zeta_2, & \zeta_2 &\sim \mathrm{Unif}[-b,b],\\
W_3 &= U_3 + \omega_3, & \omega_3 &\sim \mathcal{N}(0,(\sigma/2)^2), & Z_3 &= U_3 + \zeta_3, & \zeta_3 &\sim \mathrm{Unif}[-b,b],
\end{aligned}$$

where $b = \sigma = 0.1$, and all measurement noises are mutually independent and independent of the preceding variables. The oracle response at $A = 1$ is

$$f_{\mathrm{CATE}}(1,v) = \mathbb{E}\{Y^{(1)} \mid V = v\} = v(1+2v)^2(v-1)^2.$$

I.5 Sensitivity to proxy informativeness

We assess sensitivity to proxy informativeness in the Demand design. For $P \in \{W, Z\}$, we perturb one proxy at a time according to

$$\widetilde{P}_{ij} = P_{ij} + s\widehat{\sigma}_{P_j}\varepsilon_{ij}, \qquad \varepsilon_{ij} \sim \mathcal{N}(0,1),$$

where $\widehat{\sigma}_{P_j}$ is the empirical standard deviation of proxy coordinate j. Thus, s measures the added noise relative to the scale of that coordinate. Within each seed, the underlying generated sample, sample splits, initialization, and training randomness are held fixed across s; the $s = 0$ baseline is shared by the two perturbation arms.

Table 4: Sensitivity to proxy informativeness in the Demand design at $n = 4{,}000$. Entries are causal MSE, reported as mean $\pm$ seed-level standard error over ten data-generating seeds after averaging three repetitions within each seed. The R^2 rows measure the informativeness of the perturbed proxy for U using quadratic regression on a separate $n = 20{,}000$ diagnostic sample and therefore have no standard errors. Boldface marks the smallest MSE at each noise level.

	Noise scale s					
Method	0	0.25	0.5	1	2	4
Noise added to the outcome proxy W						
OUTCOMENET	27.5 ± 2.5	27.4 ± 2.4	30.9 ± 2.5	55.5 ± 4.0	118.9 ± 7.1	178.0 ± 6.0
TREATMENTNET	163.4 ± 11.1	175.0 ± 8.6	186.9 ± 10.5	200.9 ± 6.2	203.6 ± 3.5	200.2 ± 2.4
DR-PCL-NET	$\mathbf{22.3 \pm 2.2}$	$\mathbf{22.4 \pm 2.1}$	$\mathbf{26.3 \pm 2.5}$	$\mathbf{50.3 \pm 3.7}$	$\mathbf{113.9 \pm 6.3}$	$\mathbf{170.0 \pm 5.5}$
$R^2(U \mid \widetilde{W})$	0.460	0.432	0.369	0.237	0.097	0.028
Noise added to the treatment proxy Z						
OUTCOMENET	27.5 ± 2.5	38.5 ± 5.0	59.5 ± 7.7	101.7 ± 7.8	157.0 ± 7.7	196.3 ± 7.9
TREATMENTNET	163.4 ± 11.1	175.9 ± 11.2	162.8 ± 14.9	169.9 ± 14.7	183.3 ± 17.8	194.1 ± 18.4
DR-PCL-NET	$\mathbf{22.3 \pm 2.2}$	$\mathbf{31.4 \pm 4.5}$	$\mathbf{50.2 \pm 7.6}$	$\mathbf{91.7 \pm 7.7}$	$\mathbf{148.7 \pm 7.3}$	$\mathbf{190.5 \pm 7.8}$
$R^2(U \mid \widetilde{Z})$	0.402	0.378	0.322	0.201	0.080	0.023

The effect of corrupting one proxy is not confined to a single bridge route. Indeed, for square-integrable $u(a, x, w)$ and $v(a, x, z)$,

$$\langle T_h u, v \rangle_{L^2(\mathbb{P}_{A,X,Z})} = \mathbb{E}\{u(A, X, W)v(A, X, Z)\} = \langle u, T_\varphi v \rangle_{L^2(\mathbb{P}_{A,X,W})}.$$

Thus, $T_\varphi = T_h^*$; when compact, the two operators have the same nonzero singular values. Independent measurement noise in either proxy composes this adjoint pair with an L^2-contraction and can shrink the shared informative spectrum. Hence, when one bridge problem becomes less stable, deterioration of the other route is also expected. Their finite-sample responses need not coincide because their targets, function classes, and regularization differ. Consistent with this interpretation, causal MSE generally rises as $R^2(U \mid \widetilde{P})$ falls. DR-PCL-NET has the smallest MSE in every reported cell, although it also deteriorates under severe proxy corruption.

I.6 NEURAL ARCHITECTURES

We report the component featurizers used in the six benchmarks below. The OUTCOMENET tables use the compact feature notation of Appendix E.1. For TREATMENTNET, the compact maps are implemented as tensor products of the listed component networks; for example, $\phi^{(\varphi)}_{AW,1} = \phi^{(\varphi)}_{A,1} \otimes \phi^{(\varphi)}_{W,1}$. The doubly robust estimator uses the same two sets of bridge featurizers. In the tables, "hidden" gives the successive hidden widths. We write FC, LN, BN, SN, and DO for fully connected affine layers, layer normalization, batch normalization, spectral normalization, and dropout, respectively.

I.6.1 IMAGE-VALUED TREATMENT SETTING - DSPRITE

The reported OUTCOMENET uses a separate network for each image-valued input. Each image network has the form $4096 \to 1024 \to 512 \to 128 \to 32$, with SN on every affine layer, ReLU activations, and BN after the 512-dimensional layer. The two-dimensional treatment-proxy network uses ordinary affine layers and ReLU. The first-stage feature is $[\phi^{(h)}_{A,1}(A), 1] \otimes [\phi^{(h)}_{Z,1}(Z), 1]$, with the duplicated constant removed; it therefore has dimension $33^2 - 1 = 1088$ before the model appends its intercept. Each TREATMENTNET component uses two hidden blocks of FC–LN–GELU–BN–DO(0.05), followed by FC–LN–GELU.

Table 5: Component featurizers for the dSprite V3 benchmark.

Route	Feature map	Input (dimension)	Hidden	Output
OUTCOMENET	$\phi^{(h)}_{A,1}$	A (4096)	1024, 512, 128	32
	$\phi^{(h)}_{Z,1}$	Z (2)	128, 64	32
	$\phi^{(h)}_{A,2}$	A (4096)	1024, 512, 128	32
	$\phi^{(h)}_{W,2}$	W (4096)	1024, 512, 128	32
TREATMENTNET	$\phi^{(\varphi)}_{A,1}$	A (4096)	512, 256	16
	$\phi^{(\varphi)}_{W,1}$	W (4096)	512, 256	8
	$\phi^{(\varphi)}_{Z,2}$	Z (2)	512, 256	8
	$\phi^{(\varphi)}_{A,2}$	A (4096)	512, 256	8

The TREATMENTNET and doubly robust final-stage regressors take the flattened treatment image as input and use the architecture $4096 \to 1024 \to 512 \to 1$, with GELU, DO$(0.05)$, and a linear output.

I.6.2 DEMAND

Each OUTCOMENET featurizer has two FC–LN–GELU–DO(0.01) hidden blocks and an FC–LN–GELU output block. Each TREATMENTNET component has two FC–LN–GELU–BN–DO(0.01) hidden blocks and an FC–LN–GELU output block. The treatment-bridge first- and second-stage compact dimensions are $16 \times 8 = 128$ and $8 \times 8 = 64$, respectively.

Table 6: Component featurizers for the Demand benchmark.

Route	Feature map	Input (dimension)	Hidden	Output
OUTCOMENET	$\phi^{(h)}_{AZ,1}$	(A, Z) (3)	32, 64	32
	$\phi^{(h)}_{A,2}$	A (1)	32, 64	16
	$\phi^{(h)}_{W,2}$	W (1)	32, 64	16
TREATMENTNET	$\phi^{(\varphi)}_{A,1}$	A (1)	32, 16	16
	$\phi^{(\varphi)}_{W,1}$	W (1)	32, 16	8
	$\phi^{(\varphi)}_{Z,2}$	Z (2)	32, 16	8
	$\phi^{(\varphi)}_{A,2}$	A (1)	32, 16	8

For the final stage, the fold-specific OUTCOMENET treatment featurizer supplies the 16-dimensional input. TREATMENTNET and the doubly robust correction use $16 \to 32 \to 64 \to 1$ networks with GELU, DO(0.05), linear outputs, and five-member averaging.

I.6.3 LEGALIZED ABORTION AND CRIME

The OUTCOMENET featurizers apply two FC–LN–GELU–DO(0.01) hidden blocks and an FC–GELU output block. The TREATMENTNET block is the same FC–LN–GELU–BN–DO(0.01) architecture used for Demand.

Table 7: Component featurizers for the Legalized Abortion and Crime benchmark.

Route	Feature map	Input (dimension)	Hidden	Output
OUTCOMENET	$\phi^{(h)}_{AZ,1}$	(A, Z) (2)	256, 128	64
	$\phi^{(h)}_{A,2}$	A (1)	256, 128	8
	$\phi^{(h)}_{W,2}$	W (3)	256, 128	16
TREATMENTNET	$\phi^{(\varphi)}_{A,1}$	A (1)	32, 16	16
	$\phi^{(\varphi)}_{W,1}$	W (3)	32, 16	8
	$\phi^{(\varphi)}_{Z,2}$	Z (1)	32, 16	8
	$\phi^{(\varphi)}_{A,2}$	A (1)	32, 16	8

The TREATMENTNET and doubly robust final-stage regressors use $1 \to 32 \to 64 \to 1$, GELU, DO(0.01).

I.6.4 DEANER IM AND IR

The IM and IR designs use the same networks. The OUTCOMENET featurizers use two FC–LN–GELU–DO(0.01) hidden blocks and an FC–GELU output block. The TREATMENTNET components use two FC–LN–GELU–BN–DO(0.01) hidden blocks and an FC–LN–GELU output block. Their first- and second-stage compact dimensions are $16 \times 8 = 128$ and $8 \times 8 = 64$, respectively.

Table 8: Component featurizers for the Deaner IM and IR benchmarks.

Route	Feature map	Input (dimension)	Hidden	Output
OUTCOMENET	$\phi^{(h)}_{AZ,1}$	(A, Z) (5)	256, 128	64
	$\phi^{(h)}_{A,2}$	A (1)	256, 128	8
	$\phi^{(h)}_{W,2}$	W (3)	256, 128	16
TREATMENTNET	$\phi^{(\varphi)}_{A,1}$	A (1)	32, 16	16
	$\phi^{(\varphi)}_{W,1}$	W (3)	32, 16	8
	$\phi^{(\varphi)}_{Z,2}$	Z (4)	32, 16	8
	$\phi^{(\varphi)}_{A,2}$	A (1)	32, 16	8

The TREATMENTNET and doubly robust final stages first map the standardized scalar treatment to 128 frozen random Fourier features with unit bandwidth, followed by $128 \to 32 \to 32 \to 1$ with GELU, DO(0.05), and a linear output.

I.6.5 HETEROGENEOUS DOSE-RESPONSE

The OUTCOMENET components use two FC–LN–GELU–DO(0.01) hidden blocks and an FC–LN–GELU output block. TREATMENTNET uses the corresponding FC–LN–GELU–BN–DO(0.01) blocks. Here V is the scalar effect modifier and also enters the bridge as the observed covariate. Consequently, the treatment-bridge compact maps are $\phi^{(\varphi)}_{A,1} \otimes \phi^{(\varphi)}_{V,1} \otimes \phi^{(\varphi)}_{W,1}$ and $\phi^{(\varphi)}_{A,2} \otimes \phi^{(\varphi)}_{V,2} \otimes \phi^{(\varphi)}_{Z,2}$, both of dimension 256.

Table 9: Component featurizers for the synthetic CATE benchmark.

Route	Feature map	Input (dimension)	Hidden	Output
OUTCOMENET	$\phi^{(h)}_{AVZ,1}$	(A, V, Z) (5)	64, 128	128
	$\phi^{(h)}_{A,2}$	A (1)	64, 128	4
	$\phi^{(h)}_{V,2}$	V (1)	64, 128	8
	$\phi^{(h)}_{W,2}$	W (3)	64, 128	16
TREATMENTNET	$\phi^{(\varphi)}_{A,1}$	A (1)	32, 16	8
	$\phi^{(\varphi)}_{V,1}$	V (1)	32, 16	4
	$\phi^{(\varphi)}_{W,1}$	W (3)	32, 16	8
	$\phi^{(\varphi)}_{A,2}$	A (1)	32, 16	8
	$\phi^{(\varphi)}_{V,2}$	V (1)	32, 16	4
	$\phi^{(\varphi)}_{Z,2}$	Z (3)	32, 16	8

The OUTCOMENET final stage regresses the 16-dimensional outcome-proxy feature on V with a $1 \to 32 \to 64 \to 16$ network. For TREATMENTNET and the doubly robust estimator, a $2 \to 32 \to 64 \to 2$ network takes (A, V) and jointly estimates the numerator and denominator of the self-normalized response. All three use GELU, DO(0.01), and linear outputs.

I.6.6 TREATMENT-CONDITIONAL DOSE-RESPONSE

The OUTCOMENET components use two FC–LN–GELU–DO(0.05) hidden blocks and an FC–LN–GELU output block. TREATMENTNET again uses two FC–LN–GELU–BN–DO(0.01) hidden blocks and an FC–LN–GELU output block. Its bridge surface is represented by the separable feature $\psi^{(\varphi)}_{A}(a) \otimes \psi^{(\varphi)}_{A'}(a')$. Thus, the first-stage compact map has dimension $16 \times 8 = 128$, and the second-stage map has dimension $8 \times 8 \times 8 = 512$.

Table 10: Component featurizers for the synthetic ATT benchmark.

Route	Feature map	Input (dimension)	Hidden	Output
OUTCOMENET	$\phi^{(h)}_{AZ,1}$	(A, Z) (3)	128, 256	128
	$\phi^{(h)}_{A,2}$	A (1)	128, 256	8
	$\phi^{(h)}_{W,2}$	W (2)	128, 256	16
TREATMENTNET	$\phi^{(\varphi)}_{A,1}$	A (1)	32, 16	16
	$\phi^{(\varphi)}_{W,1}$	W (2)	32, 16	8
	$\psi^{(\varphi)}_{A}$	A (1)	32, 16	8
	$\psi^{(\varphi)}_{A'}$	A' (1)	32, 16	8
	$\phi^{(\varphi)}_{Z,2}$	Z (2)	32, 16	8

The OUTCOMENET final stage uses $1 \to 32 \to 64 \to 16$ to regress the outcome-proxy feature on the anchor treatment. TREATMENTNET and the doubly robust estimator regress scalar pseudo-outcomes on (A, A') with $2 \to 32 \to 64 \to 1$. These heads use GELU, DO(0.01), and linear outputs.

I.7 COMPUTE RESOURCES.

All experiments were run on a Linux x86_64 compute node with Python 3.12.12. GPU experiments used single-node SLURM jobs with one task, eight CPU cores, 10GB system memory, and a 24-hour wall-time limit. Unless otherwise stated, each job used a single NVIDIA RTX A4500 GPU with driver version 580.95.05 and CUDA 13.0; the GPU provided approximately 20GB of memory.