PROXMAR: NEURAL PROXIMAL CAUSAL INFERENCE WITH MISSING OUTCOME PROXIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Proximal causal inference enables causal effect estimation under unmeasured confounding by leveraging proxy variables, but existing neural proximal estimators typically assume that these proxies are fully observed. When outcome-inducing proxies are missing at random (MAR), complete-case analysis can distort both bridge learning and the population averaging used to estimate the causal mean. We propose ProxMAR, a unified framework that addresses proxy missingness at both stages. Rather than filling in missing proxy values, ProxMAR applies conditional-mean imputation, inverse probability weighting, and augmented inverse probability weighting to the estimator-specific functions required for bridge learning and causal-mean evaluation. This avoids modeling the full conditional distribution of potentially high-dimensional proxies. We instantiate the framework with Neural Maximum Moment Restriction (NMMR) and Deep Feature Proxy Variables (DFPV), yielding three missingness-aware variants of each strategy. We establish the observed-data recovery identities underlying these corrections and derive an exact product-remainder identity characterizing the double robustness of augmentation with respect to the missingness propensity and the relevant conditional-mean regression. Under method-specific regularity conditions, we establish consistency of the resulting causal-mean estimators. Numerical studies with low- and high-dimensional proxies demonstrate substantial error reductions relative to complete-case estimation across a range of missingness settings. A semi-synthetic study further shows improved recovery of method-specific full-data estimates.

1 INTRODUCTION

Unmeasured confounding remains a central challenge in causal inference from observational data. By using treatment-inducing and outcome-inducing proxies that carry information about these confounders, proximal causal inference extends causal effect identification to settings where adjustment for observed covariates alone is insufficient (Miao et al., 2018; Tchetgen et al., 2020; Cui et al., 2024). Under suitable identification assumptions, these proxies enable the construction of an outcome bridge function that connects proxy information to observed outcomes. Averaging this function over the target population at a specified treatment level identifies the corresponding causal mean. Recent neural estimators make this approach flexible enough to accommodate nonlinear relationships and high-dimensional proxies (Bennett et al., 2019; Xu et al., 2021; Mastouri et al., 2021; Kompa et al., 2022) . These advances, however, have largely been developed under the assumption that the proxy variables are fully observed.

In practice, even proxies that satisfy the identification assumptions may be missing for some individuals. Outcome-inducing proxies may be obtained through additional assessments or specialized measurements, making them susceptible to missingness due to cost, incomplete follow-up, or other data collection constraints (Little & Rubin, 2019; Chen et al., 2024). We focus on missing at random (MAR), where proxy observation is independent of the proxy values conditional on the fully observed variables (Rubin, 1976). Missingness directly complicates bridge estimation because outcome-inducing proxies enter the bridge function as inputs. Restricting estimation to complete cases can alter the conditional relationships that define the bridge, so the learned function need not be valid for the full population causal effect estimation. Averaging the bridge only over complete cases can further introduce bias because these individuals may not represent the target popula-

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.

tion (Ross et al., 2020). This motivates methods that correct the missingness at both stages: bridge learning and population averaging.

We propose ProxMAR, a neural proximal causal inference framework for estimating the average dose-response function with missing at random (MAR) outcome-inducing proxies. Its central idea is to recover both the estimator-specific quantities required for bridge learning and the population average defining the target causal mean directly from observed data, without modeling the full conditional distribution of a potentially high-dimensional proxy. In particular, ProxMAR uses three important observed-data transformations – conditional-mean imputation (IMP, Wang & Paik, 2006; Zhou et al., 2008), inverse probability weighting (IPW, Horvitz & Thompson, 1952), and augmented inverse probability weighting (AIPW, Robins et al., 1994; Bang & Robins, 2005) – to recover the conditional and marginal expectations required for bridge learning and causal-mean evaluation. Here, imputation replaces unavailable contributions to bridge learning and causal-mean evaluation by their conditional expectations given the fully observed variables, rather than filling in the missing proxy values themselves. We apply this framework to two popular proximal estimation strategies: Neural Maximum Moment Restriction (NMMR, Mastouri et al., 2021; Kompa et al., 2022) and Deep Feature Proxy Variables (DFPV, Xu et al., 2021). NMMR learns the bridge directly through conditional moment restrictions, whereas DFPV exploits a feature-based bridge representation to separate learning into a proxy-feature regression and a subsequent outcome regression. Each of the three transformations can be combined with either strategy, yielding corresponding variants of ProxMAR-NMMR and ProxMAR-DFPV. The framework also accounts for proxy missingness in the final population averaging used to estimate the average dose-response function.

Our contributions are threefold. First, we develop a unified framework that integrates three missing-data corrections with two neural proximal estimation strategies, addressing proxy missingness in both bridge learning and causal-mean evaluation. Second, we establish the observed-data recovery identities underlying these corrections and derive an exact product-remainder identity characterizing the double robustness of AIPW with respect to the missingness propensity and the relevant conditional-mean regression. Under method-specific regularity conditions, we further establish consistency of the resulting average dose-response estimators. Third, numerical studies in low- and high-dimensional settings demonstrate empirical improvements over complete-case and mean imputation proximal estimation strategies.

2 Problem Formulation and Causal Identification

In this section, we first define the target causal mean and express it as a population average of an outcome bridge under suitable identification assumptions, following Miao et al. (2018). We then show that, under MAR and observation positivity, the required full-data quantities can be recovered from incomplete records, thereby identifying the causal mean from the observed data.

2.1 Causal target and proximal bridge

Let $Y \in \mathbb{R}$ denote the outcome, $A \in \mathcal{A} \subseteq \mathbb{R}$ a continuous or binary treatment, $X \in \mathbb{R}^{p_X}$ the observed covariates, and $U \in \mathbb{R}^{p_U}$ the unobserved confounders, where $p_X, p_U \geq 1$. Let $Y(a)$ denote the potential outcome under treatment level a. Our target is the **causal mean** $\beta(a) = \mathbb{E}\{Y(a)\}$, for $a \in \mathcal{A}$. Let $\Delta(a_1, a_0) = \beta(a_1) - \beta(a_0)$. For continuous treatment, $\beta(a)$ is the average dose-response function, and for binary treatment with $\mathcal{A} = \{0, 1\}$, $\Delta(1, 0)$ is the average treatment effect.

Proximal causal inference enables identification of causal parameters in the presence of an unobserved confounder U by leveraging two types of proxies: a treatment-inducing proxy $Z \in \mathbb{R}^{p_Z}$ and an outcome-inducing proxy $W \in \mathbb{R}^{p_W}$, with $p_Z, p_W \geq 1$. Under the proximal identification conditions stated below, the causal mean is identified through outcome-bridge functions that relate these proxies to the outcome via a conditional moment equation (Miao et al., 2018; Tchetgen et al., 2020). Figure 1 illustrates a causal structure compatible with these proxy restrictions.

Assumption 1 (Proximal identification conditions)**.**

(i) **Causal conditions.** For every $a \in \mathcal{A}$, $\mathbb{E}|Y(a)| < \infty$, $Y(a) \perp\!\!\!\perp A \mid (U, X)$, and $f_{A|U,X}(a \mid U, X) > 0$ almost surely, where $f_{A|U,X}$ denotes a conditional density or probability mass function, as appropriate. Moreover, $Y = Y(A)$ almost surely (a.s.).

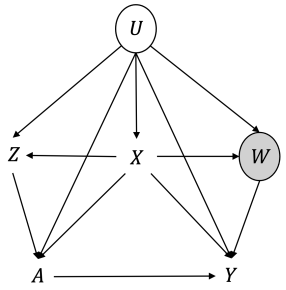


Figure 1: A causal DAG compatible with the proximal proxy restrictions. The unshaded circle denotes the unmeasured confounders U, and the shaded circle denotes the partially observed outcome proxy W. The variables A, X, Z, and Y are fully observed. The missingness mechanism is not shown.

(ii) **Proxy restrictions.** The proxies satisfy $W \perp\!\!\!\perp (A, Z) \mid (U, X)$, and $Y \perp\!\!\!\perp Z \mid (A, U, X)$.

(iii) **Outcome bridge.** There exists a measurable function h_0 such that, for every $a \in \mathcal{A}$ and P_X-almost every x,

$$\mathbb{E}(Y - h_0(W, a, x) \mid A = a, Z, X = x) = 0 \quad \text{a.s.} \tag{1}$$

Moreover, $\mathbb{E}\big[Y^2 + h_0(W, a, x)^2 \mid A = a, X = x\big] < \infty$, and $\mathbb{E}|h_0(W, a, X)| < \infty$ for every $a \in \mathcal{A}$.

(iv) **Completeness.** For every $a \in \mathcal{A}$, P_X-almost every x, and every $v \in L^2(P_{U|A=a,X=x})$, $\mathbb{E}\{v(U) \mid Z, A = a, X = x\} = 0$ a.s. $\Rightarrow v(U) = 0$ a.s.. The two almost-sure statements are with respect to $P_{Z|A=a,X=x}$ and $P_{U|A=a,X=x}$, respectively.

Assumption 1 collects standard conditions for outcome-bridge identification in proximal causal inference (Miao et al., 2018; Tchetgen et al., 2020; Cui et al., 2024). Together with the pointwise regularity conditions in Appendix A.1, these conditions yield the proximal identification

$$\beta(a) = \mathbb{E}\{h_0(W, a, X)\}, \quad a \in \mathcal{A}. \tag{2}$$

Bridge existence is assumed separately from completeness (see Appendix A.3), and uniqueness is not required: every bridge satisfying the stated conditions yields the same causal mean.

The expectation in equation 2 is taken over the target-population distribution of (W, X). When W is missing, both the bridge equation and this population average involve unavailable proxy values. The next subsection derives their observed-data representations under MAR.

2.2 Observed data and MAR recovery

We consider the case where, for each individual, the outcome-inducing proxies are either fully observed or fully missing. Specifically, let $\delta_W \in \{0, 1\}$ indicate whether W is fully observed. We observe n independent and identically distributed records

$$O_i = (Y_i, A_i, X_i, Z_i, \delta_{W,i}, \delta_{W,i} W_i), \qquad i = 1, \ldots, n.$$

The product $\delta_{W,i} W_i$ encodes the observed proxy when $\delta_{W,i} = 1$ and a zero placeholder otherwise; the indicator distinguishes a missing proxy from an observed zero. Define

$$L = (A, Z, X), \qquad L^+ = (A, Z, X, Y).$$

Here, L contains the conditioning variables in the bridge equation, while L^+ contains all fully observed variables. The following assumption specifies the missingness mechanism for W.

Assumption 2 (Missing at random and observation positivity)**.** *The proxy observation indicator satisfies*

$$W \perp\!\!\!\perp \delta_W \mid L^+, \qquad e(L^+) := \Pr(\delta_W = 1 \mid L^+) \geq \varepsilon \quad \textit{a.s.},$$

for some constant $\varepsilon > 0$, where $e(L^+)$ is the observation probability.

The MAR in Assumption 2 allows observation of W to depend on L^+, but not on W after conditioning on these variables (Rubin, 1976). Such a MAR mechanism is non-trivial: when missingness depends on L^+ rather than L, directly applying a standard proximal estimator to complete cases can

yield a bridge that does not satisfy the population bridge equation 1. Consequently, using this bridge in equation 2 need not recover the target causal mean. The observation positivity is a standard MAR condition that rules out systematic exclusion of subgroups defined by these variables from proxy measurement, providing the overlap needed to identify conditional proxy distributions throughout the target population.

We next introduce our framework to recover the causal mean from the MAR data assumed above. For any fixed $\mathbb{R}^p$-valued function $s(W, L^+)$ with integrable components, where $p \geq 1$, define its conditional mean by

$$m_s(L^+) := \mathbb{E}(s \mid L^+).$$

Under Assumption 2,

$$m_s(L^+) = \mathbb{E}\{s(W, L^+) \mid L^+, \delta_W = 1\}, \tag{3}$$

which is identified from fully observed pair (L^+, δ_W).

We consider three observed-data transformations: conditional-mean imputation (IMP), inverse probability weighting (IPW), and augmented inverse probability weighting (AIPW),

$$\begin{aligned} T_{\text{imp}}(s) &= \delta_W s + (1 - \delta_W) m_s(L^+), \quad T_{\text{ipw}}(s) = \frac{\delta_W}{e(L^+)} s, \\ T_{\text{aipw}}(s) &= m_s(L^+) + \frac{\delta_W}{e(L^+)} \{s - m_s(L^+)\}. \end{aligned} \tag{4}$$

Recalling equation 3, the above terms involving W can all be fully recovered from observed data. The IMP transformation replaces the unavailable value of s by its conditional mean; IPW reweights its observed value by the inverse observation probability; AIPW adds a weighted residual correction to the conditional mean. These transformations act on the required function s, which may be a proxy feature, a bridge residual, or a bridge evaluation. For nonlinear s, imputing s is generally different from evaluating s after replacing W by its conditional mean.

Proposition 1 (Observed-data recovery)**.** *Under Assumption 2, for any function $s(W, L^+)$ as defined above and each $j \in \mathcal{J} := \{\text{IMP}, \text{IPW}, \text{AIPW}\}$,*

$$\mathbb{E}\{T_j(s) \mid L^+\} = m_s(L^+) \quad \textit{a.s.}$$

Consequently,

$$\mathbb{E}\{T_j(s) \mid L\} = \mathbb{E}(s \mid L), \qquad \mathbb{E}\{T_j(s)\} = \mathbb{E}(s). \tag{5}$$

For vector-valued s, these identities hold componentwise.

Proposition 1 establishes our ProxMAR framework for bridge learning and causal-mean evaluation with MAR incomplete proxies. For example, take $s(W, L^+) = r_0$ where $r_0 = Y - h_0(W, A, X)$ is the bridge residual. Provided that r_0 and h_0 are integrable, Proposition 1 recovers the target outcome bridge condition equation 1 and causal mean equation 2 from observed data respectively by

$$\mathbb{E}[T_j(r_0) \mid L] = 0, \quad \beta(a) = \mathbb{E}[T_j\{h_0(W, a, X)\}], \tag{6}$$

for each $j \in \mathcal{J}$.

3 Neural Proximal Estimator under Incomplete Outcome Proxies

For any suitable function s, suppose $\widehat{e}$ and $\widehat{m}_s$ are some estimators of $e(L^+)$ and $m_s(L^+)$, respectively, using only observed data (see Section 3.3 for estimators details). Then, for any $i = 1, \ldots, n$, letting $s_i = s(W_i, L_i^+)$, the three observed-data transformations evaluated at the i-th individual can be estimated by

$$\begin{aligned} \widehat{T}_{\text{imp},i}(s_i) &= \delta_{W,i} s_i + (1 - \delta_{W,i}) \widehat{m}_s(L_i^+), \quad \widehat{T}_{\text{ipw},i}(s_i) = \frac{\delta_{W,i}}{\widehat{e}(L_i^+)} s_i, \\ \widehat{T}_{\text{aipw},i}(s_i) &= \widehat{m}_s(L_i^+) + \frac{\delta_{W,i}}{\widehat{e}(L_i^+)} \{s_i - \widehat{m}_s(L_i^+)\}. \end{aligned} \tag{7}$$

Proposition 1 establishes our ProxMAR framework under incomplete outcome proxies. In Section 3.1, we directly model h_0 by a neural network and propose a ProxMAR-NMMR estimator

using equation 6. Then in Section 3.2, we exploit further a basis expansion structure of h_0 and apply Proposition 1 to a different s to propose an alternative estimator, ProxMAR-DFPV.

The identification equation 6 requires to estimate h_0 using observed data and then take the average over (W, X) to estimate $\beta(a)$. To stabilize the computation, we use cross-fitting. For notation simplicity, we present the estimators in Sections 3.1 and 3.2 using a single training-evaluation split. Specifically, we let $\mathcal{T}$ be a bridge-training index set of size $N = |\mathcal{T}|$ and $\mathcal{I}$ be a disjoint index set for averaging estimated bridge over (W, X) to estimate $\beta(a)$. Finally, Section 3.3 discusses multi-folds sample splitting and aggregation.

3.1 PROXMAR-NMMR

Let h_θ be a neural bridge model for h_0 with parameter $\theta \in \Theta$, where Θ is the parameter space, and the corresponding models for r_0 and m_{r_0} are respectively

$$r_\theta = Y - h_\theta(W, A, X), \qquad m_{r_\theta}(L^+) = \mathbb{E}(r_\theta \mid L^+) = \mathbb{E}(r_\theta \mid L^+, \delta_W = 1).$$

Let κ be a bounded and integrally strictly positive-definite (ISPD) kernel on the support of L, e.g., the Gaussian RBF kernel. The full-data MMR criterion (e.g., Mastouri et al., 2021; Kompa et al., 2022) is

$$\mathcal{R}_\kappa(\theta) := \mathbb{E}\{r_\theta r'_\theta \kappa(L, L')\},$$

where (r'_θ, L') is an independent copy of (r_θ, L). By the definition of ISPD kernel, $\mathcal{R}_\kappa(\theta) = 0 \Leftrightarrow \mathbb{E}(r_\theta|L) = 0$ (see e.g. Zhang et al., 2020), which defines the outcome bridge in equation 1. They thus propose to estimate the outcome bridge parameter θ by minimizing a regularized V-statistic criterion approximating $\mathcal{R}_\kappa$. For MAR incomplete data, taking $s = r_\theta$, by Proposition 1, we can recover

$$\mathcal{R}_\kappa(\theta) = \mathbb{E}\big[T_j(r_\theta)T'_j(r'_\theta)\kappa(L, L')\big], \qquad j \in \mathcal{J}, \tag{8}$$

where T'_j is an independent copy of T_j.

Recalling the definition in equation 7, we then propose the following ProxMAR-NMMR estimator of θ from the MAR data, for some regularization parameter $\lambda \geq 0$,

$$\widehat{\theta}_j \in \arg\min_{\theta\in\Theta} \left\{ \frac{1}{N^2} \sum_{i,\ell\in\mathcal{T}} \widehat{T}_{j,i}(r_{\theta,i})\widehat{T}_{j,\ell}(r_{\theta,\ell})\kappa(L_i, L_\ell) + \lambda\|\theta\|_2^2 \right\}, \tag{9}$$

which minimizes an analogue regularized V-statistic criteria to those in Mastouri et al. (2021); Kompa et al. (2022). Using equation 6, $\beta(a)$ can be estimated by

$$\widehat{\beta}^{(\mathcal{I})}_{\mathrm{NMMR},j}(a) = \frac{1}{|\mathcal{I}|} \sum_{i\in\mathcal{I}} \widehat{T}_{j,i}(h_{\widehat{\theta}_j,i}), \qquad j \in \mathcal{J}. \tag{10}$$

3.2 PROXMAR-DFPV

Consider the following feature representation of the bridge:

$$h_\eta(w, a, x) = u^\top\{\psi_A(a) \otimes \phi_X(x) \otimes \psi_\omega(w)\}, \tag{11}$$

where ψ_A, ψ_ω, and ϕ_X are finite-dimensional neural features for treatment, outcome proxy, and covariates, respectively, ω parameterizes the proxy feature map, u is a coefficient vector, $\otimes$ denotes the tensor product, and η collects all bridge parameters in ψ_A, ϕ_X and u and ω. Letting $\mu_\omega(L) = \mathbb{E}\{\psi_\omega(W) \mid L\}$, the outcome bridge condition equation 1 and the causal mean in equation 2 are then modelled respectively by

$$\mathbb{E}(Y|L) = u^\top\{\psi_A(A) \otimes \phi_X(X) \otimes \mu_\omega(L)\},$$

and

$$\beta(a) = u^\top\mathbb{E}\{\psi_A(a) \otimes \phi_X(X) \otimes \mu_\omega(L)\}.$$

The idea of DFPV is to estimate η through two-stage regression (Xu et al., 2021). Stage 1 approximates $\mu_\omega(L)$ by some regression estimator, $\widehat{\mu}_\omega(L)$, for any fixed ω. Stage 2 regresses Y on

$\{\psi_A(A) \otimes \phi_X(X) \otimes \mu_\omega(L)\}$ to estimate η. Finally, taking the average over L of the estimated $u^\top \mathbb{E}\{\psi_A(a) \otimes \phi_X(X) \otimes \mu_\omega(L)\}$, they obtain an estimator of $\beta(a)$.

When W satisfies the MAR conditions in Assumption 2, taking $s = \psi_\omega$ in Proposition 1, we can recover μ_ω from incomplete data for any ω by $\mathbb{E}\{T_j(\psi_\omega)|L\}$ with $m_s = m_{\psi_\omega}$ for $j \in \mathcal{J}$. Therefore, using equation 7, we propose to approximate μ_ω by a neural network regression of $\widehat{T}_j(\psi_\omega)$ on L in Stage 1. Specifically, we further model $\mathbb{E}\{T_j(\psi_\omega)|L\} = V\Phi_\vartheta(L)$, where $\Phi_\vartheta(L)$ is some neural feature vector with parameter ϑ, and V is the coefficient matrix of compatible dimensions. For any fixed ω, Stage 1 solves

$$
\begin{aligned}
(\widehat{V}_j(\omega), \widehat{\vartheta}_j(\omega)) &\in \arg\min_{V,\vartheta} \left\{ \frac{1}{N} \sum_{i \in \mathcal{T}} \|\widehat{T}_{j,i}(\psi_\omega) - V\Phi_\vartheta(L_i)\|_2^2 + \lambda_1 \|V\|_F^2 \right\}, \\
\widehat{\mu}_{\omega,j}(L) &= \widehat{V}_j(\omega)\Phi_{\widehat{\vartheta}_j(\omega)}(L),
\end{aligned}
\tag{12}
$$

where $\lambda_1 > 0$ is a regularization parameter and $\|\cdot\|_F$ denotes the Frobenius norm. Stage 2 uses the recovered conditional features to fit the bridge:

$$
\widehat{\eta}_j \in \arg\min_{\eta} \left\{ \frac{1}{N} \sum_{i \in \mathcal{T}} \left[Y_i - u^\top \{\psi_A(A_i) \otimes \phi_X(X_i) \otimes \widehat{\mu}_{\omega,j}(L_i)\} \right]^2 + \lambda_2 \|u\|_2^2 \right\}, \tag{13}
$$

where $\lambda_2 > 0$ is a second regularization parameter.

Finally, we can have a ProxMAR-DFPV estimator of $\beta(a)$ as

$$
\widehat{\beta}^{(\mathcal{I})}_{\mathrm{DFPV},j}(a) = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \widehat{u}_j^\top \left\{ \widehat{\psi}_{A,j}(a) \otimes \widehat{\phi}_{X,j}(X_i) \otimes \widehat{\mu}_{\widehat{\omega}_j,j}(L_i) \right\}, \qquad j \in \mathcal{J}. \tag{14}
$$

Here, $\widehat{u}_j$, $\widehat{\psi}_{A,j}$, $\widehat{\omega}_j$, and $\widehat{\phi}_{X,j}$ are the fitted coefficients and feature parameters from equation 13.

3.3 Practical implementation

Both ProxMAR methods use nested cross-fitting. We partition the sample into $K \geq 2$ outer folds $\mathcal{I}_1, \ldots, \mathcal{I}_K$ and define $\mathcal{T}_k = \{1, \ldots, n\} \setminus \mathcal{I}_k$. Within each $\mathcal{T}_k$, we fit the bridge and the required nuisance estimates by inner cross-fitting, reserving $\mathcal{I}_k$ for estimating $\beta(a)$. The final estimates combine the held-out contributions:

$$
\widehat{\beta}_{M,j}(a) = \sum_{k=1}^{K} \frac{|\mathcal{I}_k|}{n} \widehat{\beta}^{(\mathcal{I}_k)}_{M,j}(a), \qquad M \in \{\mathrm{NMMR}, \mathrm{DFPV}\}, \quad j \in \mathcal{J}. \tag{15}
$$

Figure 2 illustrates this procedure. Detailed algorithms for ProxMAR-NMMR and ProxMAR-DFPV are given in Algorithms 1 and 2 in the Appendix, respectively.

In all the cases, the observation probability $e(L^+)$ is estimated by $\widehat{e}(L^+) := \max\{\tilde{e}(L^+), \underline{e}\}$, for some truncation parameter $\underline{e} \in (0, \varepsilon)$, where $\tilde{e}$ is obtained by a neural network classifier (see e.g., Richard & Lippmann, 1991) using observations of (δ_W, L^+). For the estimator $\widehat{m}_s$ with $s = r_\theta, h_{\widehat{\theta}_j}$, and ψ_ω required in equation 9, equation 10, and equation 12, respectively, based on the identification equation 3, we use Nadaraya-Watson estimator with Gaussian kernel (see e.g., Watson, 1964) on the complete cases with cross-validation bandwidths.

4 Statistical Properties

We establish consistency results for ProxMAR-NMMR and ProxMAR-DFPV, linking the IMP, IPW, and AIPW recovery conditions to bridge learning and causal mean estimation. Throughout, the numbers of outer and inner folds are fixed, the partitions are independent of the observations, and all fold sizes tend to infinity. We suppress the outer-fold superscript $(-k)$ on fitted quantities, and their convergence statements apply to every outer fit. $\widehat{\beta}_{M,j}(a)$ denotes the estimator aggregated across outer evaluation folds, where $M \in \{\mathrm{NMMR}, \mathrm{DFPV}\}$ and $j \in \mathcal{J} = \{\mathrm{imp}, \mathrm{ipw}, \mathrm{aipw}\}$. Write P for the population law and $\|f\|_{L_2(P)} = \{\mathbb{E}\|f\|_2^2\}^{1/2}$, with fitted functions held fixed when integrating over an independent observation.

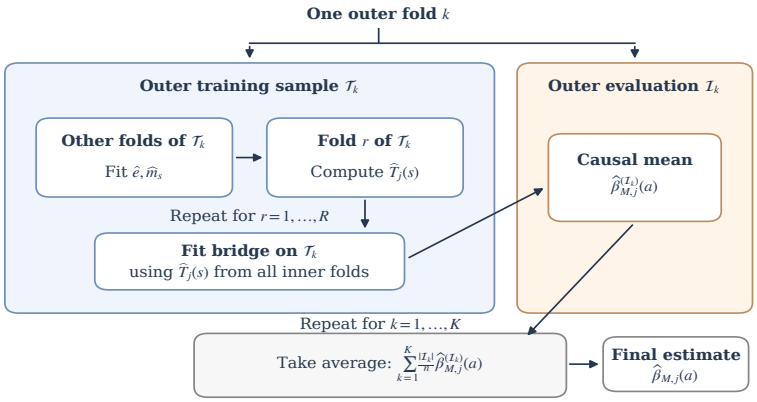


Figure 2: Nested cross-fitting for ProxMAR. The bridge is fitted on $\mathcal{T}_k$ using nuisance estimates from the inner folds, and the causal mean is evaluated on $\mathcal{I}_k$. The estimates from the outer folds are then averaged.

Consistency of ProxMAR-NMMR. Let $\widehat{\mathcal{R}}_{\kappa,j}(\theta)$ denote the unpenalized V-statistic term in Equation 9, computed on an outer training sample with sample size $N = |\mathcal{T}|$, and let $\widehat{\theta}_j$ denote the corresponding fitted bridge parameter, for $j \in \mathcal{J}$. Recall the population kernel criterion $\mathcal{R}_\kappa(\theta)$ and define $b_a(\theta) = \mathbb{E}\{h_\theta(W, a, X)\}$. To connect the population NMMR risk to the target causal mean $\beta(a)$, we impose the target-mean separation condition in Assumption C.5:

$$\inf_{\theta \in \Theta:\, |b_a(\theta) - \beta(a)| \geq \epsilon} \mathcal{R}_\kappa(\theta) > 0 \qquad \text{for every } \epsilon > 0, \tag{16}$$

where the infimum of the empty set is $+\infty$. For each $\epsilon > 0$, this condition bounds the risk away from zero for candidate bridges whose population averages at treatment a differ from $\beta(a)$ by at least ϵ. It permits multiple zero-risk bridges with the same target mean. Appendix C.2 gives sufficient conditions for equation 16.

Theorem 1. *Fix $j \in \mathcal{J}$. Suppose Assumption 2 holds and the bridge class contains a solution h_{θ_0} of the conditional moment restriction (Assumption C.4). Assume the variant-specific recovery conditions: uniform residual-imputation consistency for IMP (Assumption C.6), propensity consistency for IPW (Assumption C.7(i)), or a vanishing product remainder for AIPW (Assumption C.9, Equation 70). Under the remaining technical and approximate-minimization conditions for variant j in Appendix C, with $\lambda \to 0$ as $N \to \infty$, we have*

$$\sup_{\theta \in \Theta} \left| \widehat{\mathcal{R}}_{\kappa,j}(\theta) - \mathcal{R}_\kappa(\theta) \right| \xrightarrow{p} 0, \qquad \mathcal{R}_\kappa(\widehat{\theta}_j) \xrightarrow{p} 0. \tag{17}$$

For a fixed $a \in \mathcal{A}$, additionally suppose that Assumption 1, Equation 16, and the corresponding evaluation conditions hold: Assumption C.10 for IMP, Assumption C.7(ii) and Assumption C.11 for IPW, or Assumption C.12 for AIPW. When the treatment is continuous, impose also the pointwise conditions in Appendix A.1. Then

$$\widehat{\beta}_{\mathrm{NMMR},j}(a) \xrightarrow{p} \beta(a). \tag{18}$$

The theorem establishes that the corrected empirical criterion recovers the full-data population objective uniformly over candidate bridges and yields a consistent causal mean under the stated conditions. The missingness recovery requirements apply to both bridge learning and mean evaluation. Assumption C.5 links the resulting vanishing risk to the causal mean estimation consistency. The variant-specific proofs are given in Sections C.4.1, C.4.2, and C.4.3 of the Appendix.

Consistency of ProxMAR-DFPV. Let $\widehat{h}_j = h_{\widehat{\eta}_j}$ denote the fitted bridge, and $\widehat{\mathbb{E}}[\widehat{T}_j(\psi_\omega) \mid L] = \widehat{\mu}_{\omega,j}(L)$ denote the Stage 1 regression in Equation 12. The proxy-map parameter set Ω_n is specified in Appendix D.2.

Theorem 2. *Fix $j \in \mathcal{J}$. Suppose Assumptions 1 and 2 hold. Assume the variant-specific nuisance recovery condition (Assumption D.2), the moment and fitted-nuisance bounds (Assumptions D.1*

and D.3), and the remaining technical and approximate-minimization conditions for variant j in Appendix D.2. Then, as $n \to \infty$,

$$\sup_{\omega \in \Omega_n} \left\| \widehat{\mathbb{E}}[\widehat{T}_j(\psi_\omega) \mid L] - \mathbb{E}[T_j(\psi_\omega) \mid L] \right\|_{L_2(P_L)} \xrightarrow{p} 0, \quad \left\| \mathbb{E}(Y \mid L) - \mathbb{E}[T_j(\widehat{h}_j) \mid L] \right\|_{L_2(P_L)} \xrightarrow{p} 0. \tag{19}$$

For a fixed $a \in \mathcal{A}$, additionally suppose that stability relative to the valid-bridge set (Assumption D.6) and the fixed-treatment transfer condition (Assumption D.7(ii)) hold. When treatment is continuous, impose also the pointwise identification conditions in Appendix A.1. Then

$$\widehat{\beta}_{\mathrm{DFPV},j}(a) \xrightarrow{p} \beta(a). \tag{20}$$

The first convergence in equation 19 recovers conditional proxy features uniformly over candidate maps, while the second controls the fitted bridge's conditional moment. IMP, IPW, and AIPW enter through their respective imputation, weighted propensity-error, and product-remainder conditions. Inverse stability and target transfer then yield consistent causal means in equation 20 after averaging over held-out observations. Detailed bounds and the proof are given in Theorem D.1 of the Appendix.

5 Experiments

We evaluate ProxMAR-DFPV and ProxMAR-NMMR on a benchmark experiment, Demand design with continuous treatment and one-dimensional W, used in Xu et al. (2021) and Kompa et al. (2022), a self-designed high-dimensional setting, and semi-synthetic data, Right Heart Catheterization (RHC, Connors et al., 1996).

5.1 Demand benchmark

For the Demand design experiment, the underlying causal structure model is provided in Appendix E.1.1. We generate the missing indicators for W using the observation probability

$$e_r(L^+) = \mathrm{expit}(c_r + 1.6A^\circ + 0.8Z_1^\circ - 0.8Z_2^\circ + 1.2Y^\circ), \tag{21}$$

where the superscript $\circ$ denotes standardization to variables of mean 0 and variance 1, and c_r controls the target missing rate r. In particular, we use $c_{0.3} = 1.33$, $c_{0.5} = 0.25$, and $c_{0.7} = -1$ to obtain data with missing rate around $30\%, 50\%$ and 70%, respectively. The sample sizes used are 1,000, 2,500 and 5,000. More details for missing mechanism, evaluation metrics and implementation details are given in Appendices E.1.2 and E.1.3.

Figure 3 summarizes the performances of our proposed estimators and the naive complete case analysis across target missing rates and sample sizes. More results and the comparison with the full data case (no missing data) can be found in Appendices E.1.4 and E.1.5.

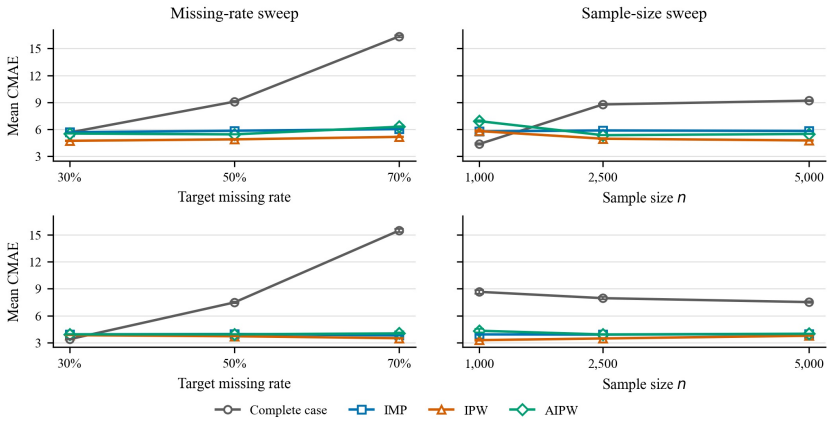


Figure 3: The mean CMAE curves with 1 standard deviation bar of DFPV methods (top) and NMMR methods (bottom) across 100 replications for changing missing rates at $n = 5000$ (left) and changing sample sizes at missing rate 50% (right).

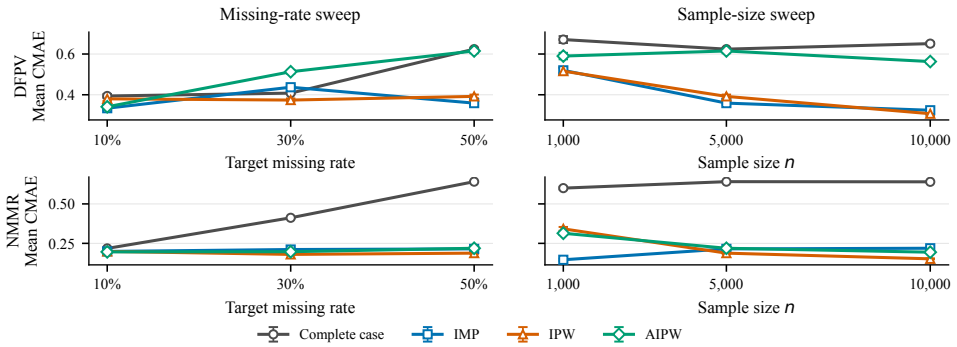


Figure 4: High-dimensional proxy mean CMAE curves with 1 standard deviation bar over 50 replications.

5.2 High-dimensional proxy study

To evaluate the performance of ProxMAR under high-dimensional proxy, we design a causal structural model and missing mechanism described in Appendices E.2.1 and E.2.2. We vary missing rate over 10%, 30%, and 50% at $(n, p_W) = (5000, 200)$ and vary $n \in \{1000, 5000, 10000\}$ at $(p_W, r) = (200, 0.5)$. Figure 4 depicts mean and standard deviation of CMAE across 50 replications for different methods. All ProxMAR-NMMR variants outperform complete-case NMMR; among the ProxMAR-DFPV variants, IPW outperforms complete-case DFPV in all six settings, while IMP and AIPW each do so in five. In particular, at $(n, p_W, r) = (5000, 200, 0.5)$, ProxMAR-NMMR-IPW reduces it from 0.641 to 0.188, while ProxMAR-DFPV-IMP reduces mean CMAE from 0.624 to 0.359 (see more detailed results in Appendices E.2.6 and E.2.5).

5.3 RHC

In the semi-synthetic RHC study with a binary treatment, we impose MAR missingness on W at rates of 30% and 50% using the missing mechanism in Appendix E.3.1. We evaluate recovery of each method's estimates of the full-data causal mean. At both missingness rates, every ProxMAR-NMMR and ProxMAR-DFPV variant achieves a lower recovery error than its corresponding complete-case estimator in each of the 50 replications. Appendix E.3 gives the details of the data, missingness mechanism, evaluation metrics, implementation, and complete results.

6 Discussion

The ProxMAR framework extends bridge estimation to studies in which outcome proxies are missing at random. Beyond NMMR and DFPV, its recovery principle can be adapted to other parametric or nonparametric proximal estimators, such as DeepGMM (Bennett et al., 2019).

A related extension is to allow the treatment-inducing proxy Z to be missing at random. Using the treatment confounding bridge $q(Z, a, X)$ proximal causal inference (Cui et al., 2024), analogous missingness corrections could be developed. When both proxy blocks may be missing, however, the problem is not simply a combination of the two single-proxy settings, as identification and estimation may depend on the joint observation patterns of (Z, W), as in general multivariate or nonmonotone missing-data problems (Robins & Gill, 1997). Future work could therefore extend ProxMAR to settings with jointly missing proxy blocks and to departures from MAR, including missing-not-at-random mechanisms, which would generally require additional identifying restrictions or sensitivity analysis (Robins, 1997; Scharfstein et al., 1999).

REFERENCES

Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

Andrew Bennett, Nathan Kallus, and Tobias Schnabel. Deep generalized method of moments for instrumental variable analysis. *Advances in neural information processing systems*, 32, 2019.

Alain Berlinet and Christine Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer Science+Business Media Dordrecht, New York, NY, 1 edition, 2004. ISBN 978-1-4020-7679-4. doi: 10.1007/978-1-4419-9096-9. Springer Book Archive.

Jacob M Chen, Rohit Bhattacharya, and Katherine A Keith. Proximal causal inference with text data. *Advances in Neural Information Processing Systems*, 37:135983–136017, 2024. doi: 10.52202/079017-4320. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/f53a37f820d5be5930415d964f4a0187-Abstract-Conference.html.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018.

Alfred F. Connors, Theodore Speroff, Neil V. Dawson, Charles Thomas, Frank E. Harrell, Douglas Wagner, Norman Desbiens, Lee Goldman, Albert W. Wu, Robert M. Califf, William J. Fulkerson, Humberto Vidaillet, Sharon Broste, Paul Bellamy, Joanne Lynn, and William A. Knaus. The effectiveness of right heart catheterization in the initial care of critically ill patients. *JAMA*, 276 (11):889–897, 1996. doi: 10.1001/jama.276.11.889.

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.

Max H Farrell, Tengyuan Liang, and Sanjog Misra. Deep neural networks for estimation and inference. *Econometrica*, 89(1):181–213, 2021.

Richard D. Gill and James M. Robins. Causal Inference for Complex Longitudinal Data: The Continuous Case. *The Annals of Statistics*, 29(6):1785 – 1811, 2001. doi: 10.1214/aos/1015345962. URL https://doi.org/10.1214/aos/1015345962.

Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.

Benjamin Kompa, David Bellamy, Tom Kolokotrones, james m robins, and Andrew Beam. Deep learning methods for proximal inference via maximum moment restriction. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 11189–11201. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/487c9d6ef55e73aa9dfd4b48fe3713a6-Paper-Conference.pdf.

Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*. John Wiley & Sons, 2019.

Afsaneh Mastouri, Yuchen Zhu, Limor Gultchin, Anna Korba, Ricardo Silva, Matt Kusner, Arthur Gretton, and Krikamol Muandet. Proximal causal learning with kernels: Two-stage estimation and moment restriction. In *International conference on machine learning*, pp. 7512–7523. PMLR, 2021.

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

Krikamol Muandet, Wittawat Jitkrittum, and Jonas Kübler. Kernel conditional moment test via maximum moment restriction. In *Conference on Uncertainty in Artificial Intelligence*, pp. 41–50. PMLR, 2020.

Whitney K. Newey and Daniel McFadden. Chapter 36 large sample estimation and hypothesis testing. In *Handbook of Econometrics*, volume 4 of *Handbook of Econometrics*, pp. 2111–2245. Elsevier, 1994. doi: https://doi.org/10.1016/S1573-4412(05)80005-4. URL https://www.sciencedirect.com/science/article/pii/S1573441205800054.

Michael D. Richard and Richard P. Lippmann. Neural network classifiers estimate Bayesian a posteriori probabilities. *Neural Computation*, 3(4):461–483, 1991. doi: 10.1162/neco.1991.3.4.461.

James M Robins. Non-response models for the analysis of non-monotone non-ignorable missing data. *Statistics in medicine*, 16(1):21–37, 1997.

James M Robins and Richard D Gill. Non-response models for the analysis of non-monotone ignorable missing data. *Statistics in medicine*, 16(1):39–56, 1997.

James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89 (427):846–866, 1994.

Rachael K Ross, Alexander Breskin, and Daniel Westreich. When is a complete-case approach to missing data valid? the importance of effect-measure modification. *American journal of epidemiology*, 189(12):1583–1589, 2020.

Donald B Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.

Daniel O Scharfstein, Andrea Rotnitzky, and James M Robins. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *Journal of the American Statistical Association*, 94 (448):1096–1120, 1999.

Bharath K. Sriperumbudur, Kenji Fukumizu, and Gert R.G. Lanckriet. Universality, characteristic kernels and rkhs embedding of measures. *Journal of Machine Learning Research*, 12(70):2389–2410, 2011. URL http://jmlr.org/papers/v12/sriperumbudur11a.html.

Eric J Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal learning. *arXiv preprint arXiv:2009.10982*, 2020.

Aad W. van der Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. Springer, New York, NY, 1 edition, 1996. ISBN 978-1-4757-2547-6. doi: 10.1007/978-1-4757-2545-2.

Cuiling Wang and Myunghee Cho Paik. Efficiencies of methods dealing with missing covariates in regression analysis. *Statistica Sinica*, pp. 1169–1192, 2006.

Geoffrey S. Watson. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)*, 26(4):359–372, 1964. ISSN 0581572X. URL http://www.jstor.org/stable/25049340.

Liyuan Xu, Heishiro Kanagawa, and Arthur Gretton. Deep proxy causal learning and its application to confounded bandit policy evaluation. *Advances in Neural Information Processing Systems*, 34: 26264–26275, 2021.

Rui Zhang, Masaaki Imaizumi, Bernhard Schölkopf, and Krikamol Muandet. Maximum moment restriction for instrumental variable regression. *arXiv preprint arXiv:2010.07684*, 3(4), 2020.

Yong Zhou, Alan T K Wan, and Xiaojing Wang. Estimating equations inference with missing data. *Journal of the American Statistical Association*, 103(483):1187–1199, 2008.

AI USE STATEMENT

Generative AI tools were used to assist with polishing and editing the manuscript, preparing code, and visualizing experimental results. The core theoretical ideas and overall proof frameworks were developed by the authors. AI assistance with the theoretical analysis involved checking and discussing assumptions and mathematical arguments, suggesting refinements, and helping fill in intermediate steps and technical details in the proofs. All AI-assisted theoretical and computational outputs were independently verified by the authors. The authors take full responsibility for the accuracy of the analyses, results, claims, and final submitted manuscript.

A IDENTIFICATION AND BRIDGE CONDITIONS

A.1 POINTWISE CONDITIONS FOR CONTINUOUS TREATMENT

For continuous treatment, conditional statements at a fixed target value require versions of the treatment-specific conditional distributions that are continuous on the joint supports of their conditioning variables. We use these versions throughout. Continuity means that the integral of every bounded continuous function is continuous as a function of the conditioning values. The distributions at each target a must arise from the same joint conditional law, so their marginalization and conditioning identities agree.

For P_X-almost every x, let

$$f_{A|X}(a \mid x) = \int f_{A|U,X}(a \mid u, x)\, dP_{U|X=x}(u)$$

be finite and positive, and define $P_{U|A=a,X=x}$ by Bayes' rule using conditional densities. For every target a, impose the pointwise causal-mean restriction

$$\mathbb{E}(Y \mid A = a, U, X) = \mathbb{E}\{Y(a) \mid U, X\} \quad P_{U,X}\text{-a.s.} \tag{22}$$

The negative-control restrictions are required at the same target value:

$$\begin{aligned} \mathbb{E}(Y \mid A = a, Z, U, X) &= \mathbb{E}(Y \mid A = a, U, X), \\ P_{W|A=a,Z,U,X} &= P_{W|U,X}, \end{aligned} \tag{23}$$

for P_X-almost every x and $P_{U,Z|A=a,X=x}$-almost every (u, z), with $X = x$ in the last display. These restrictions are imposed pointwise at every specified target value.

For binary treatment, positivity requires $\Pr(A = 0 \mid U, X) > 0$ and $\Pr(A = 1 \mid U, X) > 0$ almost surely. Equation 22 follows from consistency, conditional exchangeability, and these positive conditional treatment probabilities. Equation 23 follows from the negative-control restrictions at each of the two treatment levels. The discrete treatment support directly defines the relevant conditional distributions at $a = 0$ and $a = 1$, where both equalities hold almost surely in the remaining conditioning variables.

Equations 22 and 23 respectively equate the observed conditional mean at treatment a with the mean outcome under an intervention setting treatment to a, given (U, X), and impose the proxy restrictions at that same treatment value. For continuous treatment, $\Pr(A = a) = 0$, and continuous versions define the conditional distributions at the target value through their values at nearby treatment levels. The pointwise restrictions then specify the required equalities at a (Gill & Robins, 2001, Secs. 3–4).

A.2 PROXY INTERPRETATION

For example, let Z and W be two pretreatment measurements of an underlying health state U. The clinician uses Z when choosing treatment, while W is collected through a research process separate from treatment assignment. When the measurement errors are independent given (U, X), (U, X) contains the common causes of the measurements, treatment, and outcome, and the pathway from Z to Y operates through treatment, this design satisfies the proxy restrictions (Miao et al., 2018; Tchetgen et al., 2020).

A.3 BRIDGE EXISTENCE AND COMPLETENESS

Assumption 1 postulates a measurable outcome bridge h_0 satisfying, for every $a \in \mathcal{A}$, P_X-almost every x, and $P_{Z|A=a,X=x}$-almost every z,

$$\mathbb{E}(Y \mid A = a, Z = z, X = x) = \mathbb{E}\{h_0(W, a, x) \mid A = a, Z = z, X = x\}.$$

For every $a \in \mathcal{A}$ and P_X-almost every x, the bridge also satisfies

$$\mathbb{E}\big\{Y^2 + h_0(W, a, x)^2 \mid A = a, X = x\big\} < \infty, \qquad \mathbb{E}|h_0(W, a, X)| < \infty.$$

Sufficient conditions for bridge existence can be stated through solvability of the associated conditional integral equation. Appendix B.1 of Xu et al. (2021) gives conditions on the conditional expectation operator that ensure the existence of at least one square-integrable solution (Assumptions 4–5

therein), while Appendix B.2 adds a moment condition on Y used for identifiability (Assumption 6 therein).

The completeness condition in Assumption 1, item (iv), establishes causal identification from a given bridge. Analysis of bridge existence commonly also uses the companion condition that, for each fixed $a \in \mathcal{A}$ and P_X-almost every x, any $v \in L_2(P_{Z|A=a,X=x})$ satisfying

$$\mathbb{E}\{v(Z) \mid W, A = a, X = x\} = 0 \quad P_{W|A=a,X=x}\text{-a.s.}$$

must satisfy $v(Z) = 0$, $P_{Z|A=a,X=x}$-almost surely. Once bridge existence is assumed, the identification result in Equation 2 uses the completeness condition stated in Assumption 1, item (iv).

For an explicit bridge construction, suppose U, Z, W are scalar,

$$\mathbb{E}(W \mid U, X) = U, \qquad \mathbb{E}(Y \mid A = a, U, X) = \alpha(a, X) + \gamma(a, X)U,$$

where α and γ are coefficient functions. The first equality allows W to differ from U for an individual while requiring its conditional mean to equal U. The second makes the conditional mean of Y linear in U. Under the proxy restrictions, $h_0(w, a, x) = \alpha(a, x) + \gamma(a, x)w$ satisfies the bridge equation when the required moments are finite. This is the linear construction of Tchetgen et al. (2020, Sec. 1), with the conditional mean of W set to U.

For completeness, suppose additionally that (U, Z, W) is jointly Gaussian with a positive-definite covariance matrix conditional on $A = a, X = x$. Within this model, $\mathrm{Cov}(U, Z \mid A = a, X = x) \neq 0$ ensures the completeness condition used for causal identification, while $\mathrm{Cov}(Z, W \mid A = a, X = x) \neq 0$ ensures its companion condition above; see Miao et al. (2018, Sec. 3). This covariance characterization applies to the stated jointly Gaussian model.

A.4 Proof of Equation 2

Proof. Fix $a \in \mathcal{A}$. For P_X-almost every x, the bridge equation in Assumption 1, item (iii), gives

$$\mathbb{E}\{Y - h_0(W, a, x) \mid A = a, Z, X = x\} = 0 \quad P_{Z|A=a,X=x}\text{-a.s.}$$

For such x, define

$$v_{a,x}(u) := \mathbb{E}(Y \mid A = a, U = u, X = x) - \mathbb{E}\{h_0(W, a, x) \mid U = u, X = x\}.$$

Using the common conditional law at this a, iterated expectations, and the two pointwise proxy restrictions in Equation 23 gives

$$\mathbb{E}\{v_{a,x}(U) \mid A = a, Z, X = x\} = 0 \quad P_{Z|A=a,X=x}\text{-a.s.}$$

Integrating the second restriction in Equation 23 over Z gives

$$\mathbb{E}\{h_0(W, a, x) \mid U, X = x\} = \mathbb{E}\{h_0(W, a, x) \mid A = a, U, X = x\} \quad P_{U|A=a,X=x}\text{-a.s.}$$

Using $(s - t)^2 \leq 2s^2 + 2t^2$ and conditional Jensen's inequality, the moment condition in Assumption 1, item (iii), yields

$$\mathbb{E}\{v_{a,x}(U)^2 \mid A = a, X = x\} \leq 2\mathbb{E}\big\{Y^2 + h_0(W, a, x)^2 \mid A = a, X = x\big\} < \infty.$$

Thus $v_{a,x} \in L_2(P_{U|A=a,X=x})$, and Assumption 1, item (iv), implies $v_{a,x}(U) = 0$, $P_{U|A=a,X=x}$-almost surely. Bayes' rule and treatment positivity give

$$\frac{dP_{U|A=a,X=x}}{dP_{U|X=x}}(u) = \frac{f_{A|U,X}(a \mid u, x)}{f_{A|X}(a \mid x)} > 0 \quad P_{U|X=x}\text{-a.s.}$$

For binary treatment, the ratio is $\Pr(A = a \mid U = u, X = x)/\Pr(A = a \mid X = x)$ and is positive under Assumption 1, item (i). Consequently, a set with zero probability under $P_{U|A=a,X=x}$ also has zero probability under $P_{U|X=x}$. Therefore,

$$\mathbb{E}(Y \mid A = a, U, X) = \mathbb{E}\{h_0(W, a, X) \mid U, X\} \quad P_{U,X}\text{-a.s.}$$

The pointwise causal-mean restriction in Equation 22 gives

$$\mathbb{E}\{Y(a) \mid U, X\} = \mathbb{E}(Y \mid A = a, U, X) \quad P_{U,X}\text{-a.s.}$$

The integrability conditions permit averaging over (U, X). Combining the last two displays and applying iterated expectations yields

$$\begin{aligned}\mathbb{E}\{Y(a)\} &= \mathbb{E}[\mathbb{E}\{Y(a) \mid U, X\}] \\ &= \mathbb{E}[\mathbb{E}\{h_0(W, a, X) \mid U, X\}] \\ &= \mathbb{E}\{h_0(W, a, X)\}.\end{aligned}$$

Applying this identity at a_1 and a_0 proves the contrast formula. Since the argument applies to every bridge satisfying Assumption 1, item (iii), all such bridges yield the same $\beta(a)$ and treatment contrasts. □

B Additional conditional-moment constructions

B.1 Exact transformation remainders

Let $s = s(W, L^+) \in \mathbb{R}^d$ be integrable and define

$$m_s(L^+) := \mathbb{E}(s \mid L^+), \qquad e(L^+) := \Pr(\delta_W = 1 \mid L^+).$$

For candidate functions $\widetilde{m}_s$ and $\widetilde{e}$, define

$$\begin{aligned}\widetilde{T}_{\mathrm{imp}}(s) &:= \delta_W s + (1 - \delta_W)\widetilde{m}_s(L^+), \\ \widetilde{T}_{\mathrm{ipw}}(s) &:= \frac{\delta_W}{\widetilde{e}(L^+)} s, \\ \widetilde{T}_{\mathrm{aipw}}(s) &:= \widetilde{m}_s(L^+) + \frac{\delta_W}{\widetilde{e}(L^+)}\{s - \widetilde{m}_s(L^+)\}.\end{aligned}$$

Lemma B.1 (Observed-data transformation remainders)**.** *Suppose Assumption 2 holds, $\widetilde{m}_s$ and $\widetilde{e}$ are measurable functions of L^+, $\widetilde{e}(L^+) \geq c > 0$ almost surely, and the quantities below are integrable. Then*

$$\mathbb{E}\{\widetilde{T}_{\mathrm{imp}}(s) \mid L^+\} - m_s(L^+) = \{1 - e(L^+)\}\{\widetilde{m}_s(L^+) - m_s(L^+)\}, \tag{24}$$

$$\mathbb{E}\{\widetilde{T}_{\mathrm{ipw}}(s) \mid L^+\} - m_s(L^+) = \frac{e(L^+) - \widetilde{e}(L^+)}{\widetilde{e}(L^+)} m_s(L^+), \tag{25}$$

$$\mathbb{E}\{\widetilde{T}_{\mathrm{aipw}}(s) \mid L^+\} - m_s(L^+) = \frac{\widetilde{e}(L^+) - e(L^+)}{\widetilde{e}(L^+)}\{\widetilde{m}_s(L^+) - m_s(L^+)\}. \tag{26}$$

Writing the three right-hand sides as $B_{\mathrm{imp}}(L^+)$, $B_{\mathrm{ipw}}(L^+)$, and $B_{\mathrm{aipw}}(L^+)$, respectively, gives

$$\mathbb{E}\{\widetilde{T}_j(s) \mid L\} - \mathbb{E}(s \mid L) = \mathbb{E}\{B_j(L^+) \mid L\}, \qquad \mathbb{E}\{\widetilde{T}_j(s)\} - \mathbb{E}(s) = \mathbb{E}\{B_j(L^+)\}, \tag{27}$$

for $j \in \{\mathrm{imp}, \mathrm{ipw}, \mathrm{aipw}\}$. Hence IMP recovers the target when $\widetilde{m}_s = m_s$, IPW recovers it when $\widetilde{e} = e$, and AIPW recovers it when either equality holds. For vector-valued s, all identities hold componentwise. If the candidate functions and any index determining s are fitted on a training sample, the same identities hold conditionally on that sample when the evaluation observation is independent of it.

Proof. Assumption 2 and the definition of m_s give

$$\mathbb{E}(\delta_W s \mid L^+) = e(L^+) m_s(L^+).$$

Because $\widetilde{m}_s$ and $\widetilde{e}$ are functions of L^+,

$$\begin{aligned}\mathbb{E}\{\widetilde{T}_{\mathrm{imp}}(s) \mid L^+\} &= e(L^+) m_s(L^+) + \{1 - e(L^+)\}\widetilde{m}_s(L^+), \\ \mathbb{E}\{\widetilde{T}_{\mathrm{ipw}}(s) \mid L^+\} &= \frac{e(L^+)}{\widetilde{e}(L^+)} m_s(L^+), \\ \mathbb{E}\{\widetilde{T}_{\mathrm{aipw}}(s) \mid L^+\} &= \widetilde{m}_s(L^+) + \frac{e(L^+)}{\widetilde{e}(L^+)}\{m_s(L^+) - \widetilde{m}_s(L^+)\}.\end{aligned}$$

Subtracting $m_s(L^+)$ yields Equations 24–26. The tower property and $\mathbb{E}\{m_s(L^+) \mid L\} = \mathbb{E}(s \mid L)$ yield Equation 27. □

C ProxMAR-NMMR

C.1 Detailed estimation procedure and algorithm

C.1.1 RKHS training objective and observed-data residuals

ProxMAR–NMMR estimates the bridge by minimizing a corrected kernel moment criterion (Mastouri et al., 2021; Kompa et al., 2022). Let $\mathcal{H}_\kappa$ be an RKHS on the support of $L = (A, Z, X)$ with measurable bounded kernel κ. For $r_\theta = Y - h_\theta(W, A, X)$ and every $\theta \in \Theta$ satisfying $\mathbb{E}\{r_\theta^2\} < \infty$, define

$$\nu_\theta := \mathbb{E}\{r_\theta \kappa(L, \cdot)\} \in \mathcal{H}_\kappa.$$

By the reproducing property and the Hilbert-space dual norm identity (Berlinet & Thomas-Agnan, 2004),

$$\mathcal{R}_\kappa(\theta) := \sup_{\|f\|_{\mathcal{H}_\kappa} \le 1} [\mathbb{E}\{r_\theta f(L)\}]^2 = \|\nu_\theta\|_{\mathcal{H}_\kappa}^2 = \mathbb{E}\{r_\theta r'_\theta \kappa(L, L')\}, \tag{28}$$

where the primed variables form an independent copy. If κ is integrally strictly positive definite under the distribution of L, then

$$\mathcal{R}_\kappa(\theta) = 0 \quad \Longleftrightarrow \quad \mathbb{E}\{r_\theta \mid L\} = 0 \quad \text{almost surely.} \tag{29}$$

Thus, the zero set of the kernel risk equals the solution set of the bridge conditional moment restriction (Muandet et al., 2020; Sriperumbudur et al., 2011).

Recall $\mathcal{J} = \{\text{imp}, \text{ipw}, \text{aipw}\}$ and define $m_{r_\theta}(L^+) := \mathbb{E}\{r_\theta \mid L^+\}$. Applying the three transformations in Equation 4 to r_θ gives

$$\mathbb{E}[T_j\{r_\theta\}\kappa(L, \cdot)] = \mathbb{E}\{r_\theta \kappa(L, \cdot)\}, \qquad j \in \mathcal{J}. \tag{30}$$

Consequently, independent copies of the transformed residual recover the full-data risk:

$$\mathbb{E}\big[T_j\{r_\theta\} T'_j\{r'_\theta\}\kappa(L, L')\big] = \mathcal{R}_\kappa(\theta), \qquad j \in \mathcal{J}. \tag{31}$$

Here T'_j applies the same transformation to the independent copy, including its own auxiliary variables.

Write $\mathcal{T}_k = \{1, \dots, n\} \setminus \mathcal{I}_k$, as in Section 3.3. For outer fold $\mathcal{I}_k$, reserve $\mathcal{I}_k$ for causal-mean evaluation and train the bridge on $\mathcal{T}_k$. Partition $\mathcal{T}_k = \bigcup_{r=1}^R \mathcal{B}_{kr}$ into inner folds and define $\mathcal{T}_{kr} = \mathcal{T}_k \setminus \mathcal{B}_{kr}$. For $j \in \{\text{ipw}, \text{aipw}\}$, fit the observation probability using all observations in $\mathcal{T}_{kr}$ and set

$$\bar{e}_{ki} = \max\{\widehat{e}^{(-k,-r)}(L_i^+), \underline{e}\}, \qquad i \in \mathcal{B}_{kr},$$

where $\underline{e} \in (0, \varepsilon)$ is the fixed probability floor in Section 3.3. For a regression training set S, let $C_S = \{\ell \in S : \delta_{W,\ell} = 1\}$. Define the Gaussian kernel and complete-case weights by

$$K_b(u^+, l^+) := \exp\left\{-\frac{\|u^+ - l^+\|_2^2}{2b^2}\right\},$$
$$\omega_{S,\ell}(l^+; b) := \begin{cases} \dfrac{K_b(L_\ell^+, l^+)}{\sum_{q \in C_S} K_b(L_q^+, l^+)}, & C_S \neq \varnothing,\ \ell \in C_S, \\ 0, & \text{otherwise.} \end{cases} \tag{32}$$

For $j \in \{\text{imp}, \text{aipw}\}$, fit the residual regression on the complete cases in $\mathcal{T}_{kr}$:

$$\widehat{m}_{r_\theta}^{(-k,-r)}(L_i^+) := \sum_{\ell \in C_{\mathcal{T}_{kr}}} \omega_{\mathcal{T}_{kr},\ell}(L_i^+; b_{kr})\{Y_\ell - h_\theta(W_\ell, A_\ell, X_\ell)\}, \qquad i \in \mathcal{B}_{kr}, \tag{33}$$

The fitted residual supplied to the kernel loss is

$$\widehat{r}_{j,i}^{(-k)}(\theta) := \begin{cases} r_{\theta,i}, & j = \text{imp},\ \delta_{W,i} = 1, \\ \widehat{m}_{r_\theta}^{(-k,-r)}(L_i^+), & j = \text{imp},\ \delta_{W,i} = 0, \\ r_{\theta,i}/\bar{e}_{ki}, & j = \text{ipw},\ \delta_{W,i} = 1, \\ 0, & j = \text{ipw},\ \delta_{W,i} = 0, \\ \widehat{m}_{r_\theta}^{(-k,-r)}(L_i^+) + \{r_{\theta,i} - \widehat{m}_{r_\theta}^{(-k,-r)}(L_i^+)\}/\bar{e}_{ki}, & j = \text{aipw},\ \delta_{W,i} = 1, \\ \widehat{m}_{r_\theta}^{(-k,-r)}(L_i^+), & j = \text{aipw},\ \delta_{W,i} = 0, \end{cases} \qquad i \in \mathcal{B}_{kr}. \tag{34}$$

Terms involving an individual's W_i are evaluated only when $\delta_{W,i} = 1$.

Within the shared kernel objective, the three members differ in the observed-data transformation applied to the residual and to the target-treatment bridge value:

Variant	Training quantity	Evaluation quantity	Fitted nuisances
NMMR-IMP	$\widehat{r}_{\mathrm{imp}}$	$\widehat{v}_{\mathrm{imp}}(a,\theta)$	$m_{r_\theta}, q_{a,\theta}$
NMMR-IPW	$\widehat{r}_{\mathrm{ipw}}$	$\widehat{v}_{\mathrm{ipw}}(a,\theta)$	e
NMMR-AIPW	$\widehat{r}_{\mathrm{aipw}}$	$\widehat{v}_{\mathrm{aipw}}(a,\theta)$	$e, m_{r_\theta}, q_{a,\theta}$

Let $N_k = |\mathcal{T}_k|$, let $\mathbf{r}_{-k,j}(\theta)$ collect the residuals in Equation 34, and define $\mathbf{K}_{-k} = [\kappa(L_i, L_\ell)]_{i,\ell\in\mathcal{T}_k}$. Write $\mathbf{K}^{\circ}_{-k}$ for the same matrix with zero diagonal. The normalized penalized U- and V-statistic losses are

$$\widehat{\mathcal{L}}^{(-k)}_{s,n,j}(\theta) := \begin{cases} \dfrac{\mathbf{r}_{-k,j}(\theta)^\top \mathbf{K}^{\circ}_{-k} \mathbf{r}_{-k,j}(\theta)}{N_k(N_k-1)} + \lambda_n \|\theta\|_2^2, & s = U, \\ \dfrac{\mathbf{r}_{-k,j}(\theta)^\top \mathbf{K}_{-k} \mathbf{r}_{-k,j}(\theta)}{N_k^2} + \lambda_n \|\theta\|_2^2, & s = V. \end{cases} \tag{35}$$

The empirical RKHS embedding and its unpenalized V-statistic risk are

$$\widehat{\nu}_{k,j}(\theta) := \frac{1}{N_k} \sum_{i\in\mathcal{T}_k} \widehat{r}^{(-k)}_{j,i}(\theta) \kappa(L_i, \cdot), \tag{36}$$

$$\widehat{\mathcal{R}}^{(-k)}_{V,0,j}(\theta) := \|\widehat{\nu}_{k,j}(\theta)\|^2_{\mathcal{H}_\kappa}. \tag{37}$$

For the selected loss form $s \in \{U, V\}$, write

$$\widehat{\mathcal{Q}}_{k,j}(\theta) := \widehat{\mathcal{R}}^{(-k)}_{s,0,j}(\theta) := \widehat{\mathcal{L}}^{(-k)}_{s,n,j}(\theta) - \lambda_n \|\theta\|_2^2. \tag{38}$$

The estimator $\widehat{\theta}^{(-k)}_j$ is the numerical fit to this penalized objective. The consistency results require the approximate-minimization condition in Equation 54; they do not assert that a fixed number of optimizer updates guarantees that condition. The V-statistic is the objective in Section 3.1; the U-statistic is an alternative obtained by removing the diagonal. An unnormalized quadratic criterion with penalty coefficient λ has normalized coefficient λ/N_k^2 for V and $\lambda/\{N_k(N_k-1)\}$ for U.

C.1.2 Causal-mean evaluation

For a target treatment a, define

$$q_{a,\theta}(L^+) := \mathbb{E}\{h_\theta(W, a, X) \mid L^+\}.$$

Under Assumption 2, $q_{a,\theta}(L^+) = \mathbb{E}\{h_\theta(W, a, X) \mid L^+, \delta_W = 1\}$. For $j \in \{\mathrm{imp}, \mathrm{aipw}\}$, estimate this regression on the complete cases in $\mathcal{T}_k$ by

$$\widehat{q}^{(-k)}_{a,\theta}(L^+_i) := \sum_{\ell\in C_{\mathcal{T}_k}} \omega_{\mathcal{T}_k,\ell}(L^+_i; b_{q,k}(a)) h_\theta(W_\ell, a, X_\ell), \qquad i \in \mathcal{I}_k. \tag{39}$$

For $j \in \{\mathrm{ipw}, \mathrm{aipw}\}$, fit the observation-probability learner on $\mathcal{T}_k$ and set

$$\bar{e}^{\mathrm{eval}}_k(L^+_i) = \max\{\widehat{e}^{(-k)}(L^+_i), \underline{e}\}, \qquad i \in \mathcal{I}_k.$$

The observation probability and the conditional regression are evaluated at the factual $L^+_i = (A_i, Z_i, X_i, Y_i)$; the treatment argument of the bridge is set to a. For $i \in \mathcal{I}_k$, define

$$\widehat{v}^{(-k)}_{j,i}(a,\theta) := \begin{cases} h_\theta(W_i, a, X_i), & j = \mathrm{imp},\ \delta_{W,i} = 1, \\ \widehat{q}^{(-k)}_{a,\theta}(L^+_i), & j = \mathrm{imp},\ \delta_{W,i} = 0, \\ h_\theta(W_i, a, X_i)/\bar{e}^{\mathrm{eval}}_k(L^+_i), & j = \mathrm{ipw},\ \delta_{W,i} = 1, \\ 0, & j = \mathrm{ipw},\ \delta_{W,i} = 0, \\ \widehat{q}^{(-k)}_{a,\theta}(L^+_i) + \{h_\theta(W_i, a, X_i) - \widehat{q}^{(-k)}_{a,\theta}(L^+_i)\}/\bar{e}^{\mathrm{eval}}_k(L^+_i), & j = \mathrm{aipw},\ \delta_{W,i} = 1, \\ \widehat{q}^{(-k)}_{a,\theta}(L^+_i), & j = \mathrm{aipw},\ \delta_{W,i} = 0. \end{cases} \tag{40}$$

The held-out fold contribution and cross-fitted causal-mean estimator are

$$\widehat{\beta}_{k,j}(a) := \frac{1}{|\mathcal{I}_k|} \sum_{i \in \mathcal{I}_k} \widehat{v}_{j,i}^{(-k)}(a, \widehat{\theta}_j^{(-k)}), \tag{41}$$

$$\widehat{\beta}_{\mathrm{NMMR},j}(a) := \sum_{k=1}^{K} \frac{|\mathcal{I}_k|}{n} \widehat{\beta}_{k,j}(a). \tag{42}$$

C.1.3 Kernel implementation and gradients

The MMR kernel κ on L and the Gaussian smoothing kernel K_b on L^+ serve separate roles. Within each regression training sample, the coordinates of L^+ are standardized using that sample's means and standard deviations. Let $\mathcal{B}_n \subset (0, \infty)$ be a finite candidate set fixed before training. For a training sample with at least two complete cases, each candidate predicts every complete-case response from all other complete cases, and the selected bandwidth minimizes the mean leave-one-out squared prediction error, with the smallest candidate selected for ties. A sample with fewer than two complete cases uses $\min \mathcal{B}_n$; a single donor yields its response, and an empty donor set yields the zero function and is recorded.

For residual regression on $\mathcal{B}_{kr}$, cross-validation uses the initial residuals $r_{\theta_k^{(0)},\ell}$ on the complete cases in $\mathcal{T}_{kr}$ and selects b_{kr}. For causal-mean evaluation at treatment a, cross-validation uses $h_{\widehat{\theta}_j^{(-k)}}(W_\ell, a, X_\ell)$ on the complete cases in $\mathcal{T}_k$ and selects $b_{q,k}(a)$. These bandwidths and the resulting normalized kernel weights are held fixed when differentiating with respect to the bridge parameters. Both regressions use the Gaussian weights defined in Equation 32.

At each bridge update for NMMR-IMP and NMMR-AIPW, Equation 33 is evaluated at the current parameter θ. The normalized smoothing weights are fixed, while the donor residuals retain their dependence on θ. In particular,

$$\nabla_\theta \widehat{m}_{r_\theta}^{(-k,-r)}(L_i^+) = -\sum_{\ell \in C_{\mathcal{T}_{kr}}} \omega_{\mathcal{T}_{kr},\ell}(L_i^+; b_{kr}) \nabla_\theta h_\theta(W_\ell, A_\ell, X_\ell). \tag{43}$$

The gradient of Equation 34 follows through this donor regression for IMP and AIPW and through the observed residual for all three variants. Fitted observation probabilities and smoothing weights remain fixed. An empty donor set produces a zero regression and zero derivative.

The quadratic terms in Equation 35 can be computed in row blocks as in Kompa et al. (2022, Appendix D). Each row block is paired with every column in $\mathcal{T}_k$, so the accumulated contributions include pairs drawn from different inner folds. Summing the block contributions with the denominator in Equation 35, accumulating their gradients, and adding the penalty once reproduces the full-matrix loss and gradient.

For a nonempty closed parameter set Θ, let Π_Θ denote Euclidean projection onto Θ, with a fixed measurable rule for ties. For unconstrained network parameters, $\Theta = \mathbb{R}^p$ and Π_Θ is the identity. The optimizer is an implementation choice subject to the approximate-minimization condition used below. Each outer fold initializes and trains its own bridge parameters.

C.1.4 Algorithm

Algorithm 1 summarizes the NMMR variants. Let E denote the observation-probability learner and let $\mathcal{A}_{\mathrm{eval}}$ denote the target treatment values.

C.2 Shared ProxMAR-NMMR conditions

Throughout this appendix, $\Xi_i = (W_i, A_i, X_i, Y_i, Z_i, \delta_{W,i})$ are i.i.d.; $O_i = (L_i^+, \delta_{W,i}, \delta_{W,i} W_i)$ denotes the observed record. All suprema are measurable, the RKHS is separable, and the conditional expectations and fitted functions have jointly measurable versions. We use a common regular conditional distribution of W given L^+ to define m_{r_θ} and $q_{a,\theta}$ for all θ. The numbers of outer and inner folds, $K, R \geq 2$, are fixed. The partitions are independent of the observations, and $\min_k |\mathcal{I}_k| \to \infty$ and $\min_{k,r} |\mathcal{B}_{kr}| \to \infty$. All fitting randomization is independent of the observations. Fitting on $\mathcal{T}_{kr}$

Algorithm 1 NMMR-IMP, NMMR-IPW, and NMMR-AIPW

Require: Observed data; transformation type $j \in \mathcal{J}$; outer and inner fold counts K, R; MMR kernel κ; observation-probability learner E; probability floor $\underline{e}$; bandwidth candidates $\mathcal{B}_n$; loss form $s \in \{U, V\}$; penalty λ_n; optimizer $\mathcal{O}_j$; update count T; target treatments $\mathcal{A}_{\text{eval}}$

Partition $\{1, \dots, n\}$ into outer folds $\mathcal{I}_1, \dots, \mathcal{I}_K$ independently of the observations and keep the partition fixed
for $k = 1, \dots, K$ **do**
  Reserve $\mathcal{I}_k$ for causal-mean evaluation; set $\mathcal{T}_k = \bigcup_{\ell \neq k} \mathcal{I}_\ell$ and $N_k = |\mathcal{T}_k|$
  Partition $\mathcal{T}_k$ into inner folds $\mathcal{B}_{k1}, \dots, \mathcal{B}_{kR}$ independently of the observations and set $\mathcal{T}_{kr} = \mathcal{T}_k \setminus \mathcal{B}_{kr}$
  Initialize $\theta_k^{(0)} \in \Theta$ independently of the observations
  for $r = 1, \dots, R$ **do**
    For $j \in \{\text{ipw}, \text{aipw}\}$, fit E on $\mathcal{T}_{kr}$ and compute $\bar{e}_{ki}$ on $\mathcal{B}_{kr}$
    For $j \in \{\text{imp}, \text{aipw}\}$, select b_{kr} by leave-one-out cross-validation using $r_{\theta_k^{(0)}, \ell}$ and compute the normalized donor weights from $\mathcal{T}_{kr}$ to $\mathcal{B}_{kr}$
  end for
  for $t = 0, \dots, T-1$ **do**
    Compute $\widehat{r}_{j,i}^{(-k)}(\theta_k^{(t)})$ for all $i \in \mathcal{T}_k$ using Equation 34
    Evaluate $\widehat{\mathcal{L}}_{s,n,j}^{(-k)}(\theta_k^{(t)})$ on all pairs in $\mathcal{T}_k$ and compute its gradient using the rules above
    Apply $\mathcal{O}_j$ to the full loss and gradient to obtain $\theta_k^{(t+1)}$
  end for
  Set $\widehat{\theta}_j^{(-k)} = \theta_k^{(T)}$
  For $j \in \{\text{ipw}, \text{aipw}\}$, fit E on $\mathcal{T}_k$ and compute $\bar{e}_k^{\text{eval}}$ on $\mathcal{I}_k$
  for $a \in \mathcal{A}_{\text{eval}}$ **do**
    For $j \in \{\text{imp}, \text{aipw}\}$, select $b_{q,k}(a)$ and fit $\widehat{q}_{a, \widehat{\theta}_j^{(-k)}}^{(-k)}$ using Equation 39
    Compute $\widehat{\beta}_{k,j}(a)$ on $\mathcal{I}_k$ using Equations 40–41
  end for
end for
Combine the outer-fold contributions using Equation 42
return $\{\widehat{\beta}_{\text{NMMR},j}(a) : a \in \mathcal{A}_{\text{eval}}\}$ and empty-donor-set records for $j \in \{\text{imp}, \text{aipw}\}$

uses only that sample and independent randomization; causal-mean evaluation on $\mathcal{I}_k$ uses fits constructed solely from $\mathcal{T}_k$. Statements below may be conditioned on the partitions. No independence between different fitted folds is required.

Assumption C.1 (Kernel conditions)**.** *The kernel* k *is bounded, continuous, symmetric, and integrally strictly positive definite under the distribution of* $L = (A, Z, X)$.

Assumption C.2 (Second moments of the outcome and bridge class)**.** *There is a nonnegative measurable envelope* $H_*(W, A, X)$*, independent of* n*, such that*

$$\sup_{\theta \in \Theta} |h_\theta(W, A, X)| \leq H_*(W, A, X) \quad a.s., \qquad \mathbb{E}\{Y^2 + H_*(W, A, X)^2\} < \infty.$$

Write $F_r = |Y| + H_*(W, A, X)$*, so that* $\sup_\theta |r_\theta| \leq F_r$ *and* $\mathbb{E}(F_r^2) < \infty$.

To state the complexity conditions precisely, let $\mathcal{Z} = \{z_\theta : \theta \in \Theta\}$ be a deterministic class of real-valued functions of Ξ. For $M > 0$, put $z_{\theta,M} = \max\{-M, \min(z_\theta, M)\}$ and define

$$\mathcal{F}_M(\mathcal{Z}) = \{(\xi, \xi') \mapsto z_{\theta,M}(\xi) z_{\theta,M}(\xi') \kappa(l, l') : \theta \in \Theta\}.$$

For a class $\mathcal{F}$ on pairs, its section classes are

$$\mathcal{F}_1 = \{\xi \mapsto f(\xi', \xi) : f \in \mathcal{F},\ \xi' \in \mathcal{X}_\Xi\}, \qquad \mathcal{F}_2 = \{\xi \mapsto f(\xi, \xi') : f \in \mathcal{F},\ \xi' \in \mathcal{X}_\Xi\}.$$

Here $\mathcal{X}_\Xi$ is the support of Ξ. Define the expected Rademacher complexity by

$$\mathfrak{R}_m(\mathcal{G}) = \mathbb{E}_{\Xi_{1:m}, \sigma_{1:m}} \left[\sup_{g \in \mathcal{G}} \frac{1}{m} \sum_{i=1}^m \sigma_i g(\Xi_i) \right],$$

where the σ_i are independent symmetric signs, independent of the observations. We shall use the following condition on $\mathcal{Z}$:

$$\mathfrak{R}_{m-1}\{\mathcal{F}_M(\mathcal{Z})_1\} + \mathfrak{R}_m\{\mathcal{F}_M(\mathcal{Z})_2\} \longrightarrow 0 \quad \text{for every fixed } M < \infty. \tag{44}$$

Assumption C.3 (Full-data complexity)**.** *The residual class $\mathcal{Z}_{\text{full}} = \{r_\theta : \theta \in \Theta\}$ satisfies Equation 44.*

This condition concerns the unpenalized full-data criterion. For uniformly bounded residuals, truncation can be omitted by taking M above their common bound. The truncation argument below allows unbounded outcomes under Assumption C.2.

Assumption C.4 (Bridge specification)**.** *There exists $\theta_0 \in \Theta$ such that $\mathbb{E}(r_{\theta_0} \mid L) = 0$ a.s. The bridge and its parameterization may be nonunique.*

Assumption C.5 (Target-mean separation)**.** *At the fixed treatment value a under consideration, h_{θ_0} satisfies the outcome-bridge condition in Assumption 1. The quantities $b_a(\theta) = \mathbb{E}\{h_\theta(W, a, X)\}$ are finite, and for every $\epsilon > 0$,*

$$c_a(\epsilon) := \inf_{\theta\in\Theta:\, |b_a(\theta)-b_a(\theta_0)|\geq\epsilon} \mathcal{R}_\kappa(\theta) > 0, \tag{45}$$

with the convention $\inf \varnothing = +\infty$.

This condition separates incorrect target means from the zero-risk set without requiring a unique bridge or a unique network parameter. For example, it follows if Θ is compact, both $\mathcal{R}_\kappa$ and b_a are continuous, the risk has a zero, and every zero-risk candidate has target mean $\beta(a) = b_a(\theta_0)$: the closed set in Equation 45 is compact, and a zero minimum on a nonempty such set contradicts the last condition. For continuous treatment, identification at the specified a additionally uses the pointwise conditions in Appendix A.1.

Lemma C.1 (Bounded U-process deviation; Kompa et al., 2022, Lemma 3 and Corollary 4)**.** *Let $\Xi_1, \ldots, \Xi_m$ be i.i.d., $m \geq 2$, and let $\mathcal{F}$ be a measurable class of functions on pairs with $|f| \leq B$. Write*

$$U_m f = \frac{1}{m(m-1)} \sum_{i\neq\ell} f(\Xi_i, \Xi_\ell), \qquad P^2 f = \mathbb{E}\{f(\Xi, \Xi')\},$$

where Ξ' is an independent copy. For $0 < \delta < 1$, with probability at least $1 - \delta$,

$$\sup_{f\in\mathcal{F}} |U_m f - P^2 f| \leq 2\{\mathfrak{R}_{m-1}(\mathcal{F}_1) + \mathfrak{R}_m(\mathcal{F}_2)\} + 2B\sqrt{\frac{2\log(2/\delta)}{m}}. \tag{46}$$

Proof. Set the regularization parameter to zero in Lemma 3 of Kompa et al. (2022). Apply its two one-sided bounds with error probability $\delta/2$ and take a union bound. The resulting two-sided inequality is Equation 46. □

Lemma C.2 (Uniform convergence with a square-integrable envelope)**.** *Suppose Assumption C.1 holds and $\mathcal{Z} = \{z_\theta\}$ satisfies Equation 44. If $\sup_\theta |z_\theta| \leq F$ with $\mathbb{E}(F^2) < \infty$, define*

$$f_\theta(\xi, \xi') = z_\theta(\xi) z_\theta(\xi') \kappa(l, l'), \qquad V_m f_\theta = \frac{1}{m^2} \sum_{i,\ell} f_\theta(\Xi_i, \Xi_\ell).$$

Then

$$\mathbb{E}\left[\sup_\theta |U_m f_\theta - P^2 f_\theta|\right] \to 0, \qquad \sup_\theta |V_m f_\theta - P^2 f_\theta| = o_p(1). \tag{47}$$

Proof. For fixed M, Lemma C.1, with $B = C_\kappa M^2$, gives uniform convergence of the clipped U-statistic. Its supremum deviation is bounded by $2C_\kappa M^2$, so convergence also holds in L_1. Moreover,

$$\sup_\theta |f_\theta(\Xi, \Xi') - f_{\theta,M}(\Xi, \Xi')| \le C_\kappa F F' \{\mathbf{1}(F > M) + \mathbf{1}(F' > M)\}.$$

Independence of distinct observations bounds the expected uniform empirical truncation error, and also the population truncation error, by $2C_\kappa \mathbb{E}(F)\mathbb{E}\{F\mathbf{1}(F > M)\}$ each. Consequently,

$$\limsup_{m\to\infty} \mathbb{E} \sup_\theta |U_m f_\theta - P^2 f_\theta| \le 4C_\kappa \mathbb{E}(F)\mathbb{E}\{F\mathbf{1}(F > M)\} \to 0$$

as $M \to \infty$. Finally,

$$\sup_\theta |V_m f_\theta - U_m f_\theta| \le \frac{2C_\kappa}{m} \left(\frac{1}{m} \sum_{i=1}^m F_i^2 \right) = O_p(m^{-1}). \tag{48}$$

Indeed, $V_m = (1 - m^{-1})U_m + m^{-2} \sum_i f_\theta(\Xi_i, \Xi_i)$ and $|U_m f_\theta| \le C_\kappa m^{-1} \sum_i F_i^2$. This proves the V-statistic assertion. □

On the same outer training sample used by the estimator, define

$$\widehat{\mathcal{R}}_U^{\text{full},(-k)}(\theta) = \frac{1}{N_k(N_k - 1)} \sum_{\substack{i,\ell \in \mathcal{T}_k \\ i \ne \ell}} r_{\theta,i} r_{\theta,\ell} \kappa(L_i, L_\ell), \tag{49}$$

$$\widehat{\mathcal{R}}_V^{\text{full},(-k)}(\theta) = \frac{1}{N_k^2} \sum_{i,\ell \in \mathcal{T}_k} r_{\theta,i} r_{\theta,\ell} \kappa(L_i, L_\ell). \tag{50}$$

Lemma C.3 (Full-data NMMR benchmark)**.** *Under Assumptions C.1–C.3, for $s \in \{U, V\}$,*

$$B_{n,s} := \max_{1 \le k \le K} \sup_{\theta \in \Theta} \left| \widehat{\mathcal{R}}_s^{\text{full},(-k)}(\theta) - \mathcal{R}_\kappa(\theta) \right| = o_p(1). \tag{51}$$

For each fold the supremum U-statistic deviation also tends to zero in L_1.

Proof. Apply Lemma C.2 with $z_\theta = r_\theta$ and envelope F_r. Since K is fixed and all N_k diverge, a union bound establishes Equation 51. □

C.3 A COMMON CONSISTENCY THEOREM AND THE CRITERION DECOMPOSITION

For a selected loss form s, compare the unpenalized corrected criterion with its full-data counterpart:

$$\begin{aligned} \Gamma_{n,k,j} &:= \sup_{\theta \in \Theta} |\widehat{\mathcal{Q}}_{k,j}(\theta) - \mathcal{R}_\kappa(\theta)| \\ &\le \underbrace{\sup_{\theta \in \Theta} |\widehat{\mathcal{Q}}_{k,j}(\theta) - \widehat{\mathcal{R}}_s^{\text{full},(-k)}(\theta)|}_{A_{n,k,j}\text{: observed-data correction}} + \underbrace{\sup_{\theta \in \Theta} |\widehat{\mathcal{R}}_s^{\text{full},(-k)}(\theta) - \mathcal{R}_\kappa(\theta)|}_{B_{n,k,s}\text{: full-data deviation}}. \end{aligned} \tag{52}$$

Lemma C.3 controls the second term. The variant-specific results below control the first. Both criteria in this comparison are unpenalized; regularization is handled in the following theorem.

Theorem C.1 (ProxMAR–NMMR family consistency)**.** *Fix $j \in \mathcal{J}$. For each outer fold k, let $\widehat{\mathcal{Q}}_{k,j}$ be the unpenalized U- or V-statistic criterion used to fit $\widehat{\theta}_j^{(-k)}$. Suppose Assumptions C.1, C.2, and C.4 hold, $\|\theta_0\|_2 < \infty$, and*

$$\sup_{\theta \in \Theta} |\widehat{\mathcal{Q}}_{k,j}(\theta) - \mathcal{R}_\kappa(\theta)| = o_p(1) \tag{53}$$

for every k. Let $\lambda_n \ge 0$ tend to zero and suppose

$$\begin{aligned} &\widehat{\mathcal{Q}}_{k,j}(\widehat{\theta}_j^{(-k)}) + \lambda_n \|\widehat{\theta}_j^{(-k)}\|_2^2 \\ &\qquad \le \inf_{\theta \in \Theta} \{\widehat{\mathcal{Q}}_{k,j}(\theta) + \lambda_n \|\theta\|_2^2\} + o_p(1). \end{aligned} \tag{54}$$

Then $\max_k \mathcal{R}_\kappa(\widehat{\theta}_j^{(-k)}) = o_p(1)$. *At a fixed a, additionally assume Assumptions 1 and C.5 with $h_0 = h_{\theta_0}$, the pointwise regularity conditions in Appendix A.1 when A is continuous, and*

$$\max_k |\widehat{\beta}_{k,j}(a) - b_a(\widehat{\theta}_j^{(-k)})| = o_p(1). \tag{55}$$

Then $\widehat{\beta}_{\mathrm{NMMR},j}(a) \xrightarrow{p} \beta(a)$.

Proof. Comparing Equation 54 with θ_0, using $\mathcal{R}_\kappa(\theta_0) = 0$, and dropping the nonnegative fitted-parameter penalty gives

$$0 \le \mathcal{R}_\kappa(\widehat{\theta}_j^{(-k)}) \le 2\Gamma_{n,k,j} + \lambda_n \|\theta_0\|_2^2 + o_p(1) = o_p(1). \tag{56}$$

The fixed number of folds permits taking the maximum over k. The proximal identification assumptions give $b_a(\theta_0) = \beta(a)$. If the set defining $c_a(\epsilon)$ is empty, all candidates have target error less than ϵ. Otherwise its infimum is finite and positive, and

$$\Pr\{\max_k |b_a(\widehat{\theta}_j^{(-k)}) - \beta(a)| \ge \epsilon\} \le \sum_{k=1}^{K} \Pr\{\mathcal{R}_\kappa(\widehat{\theta}_j^{(-k)}) \ge c_a(\epsilon)\} \longrightarrow 0.$$

Combine this with Equation 55 and the nonnegative fold weights, which sum to one, to obtain the claim. □

We next isolate a perturbation argument used for IPW and AIPW. Define the oracle residual $r^o_{j,\theta} = T_j(r_\theta)$ using the true nuisance functions, and let

$$\widehat{\nu}^o_{k,j}(\theta) = N_k^{-1} \sum_{i \in \mathcal{T}_k} r^o_{j,\theta,i} \kappa(L_i, \cdot), \tag{57}$$

$$\widehat{\nu}^{\mathrm{full}}_k(\theta) = N_k^{-1} \sum_{i \in \mathcal{T}_k} r_{\theta,i} \kappa(L_i, \cdot). \tag{58}$$

Put $\widehat{\mathcal{R}}^{o,(-k)}_{V,0,j} = \|\widehat{\nu}^o_{k,j}\|^2_{\mathcal{H}_\kappa}$.

Lemma C.4 (RKHS embedding perturbation)**.** *Let*

$$D_{n,k,j} = \sup_\theta \|\widehat{\nu}_{k,j}(\theta) - \widehat{\nu}^o_{k,j}(\theta)\|_{\mathcal{H}_\kappa}, \qquad M_{n,k,j} = \sup_\theta \|\widehat{\nu}^o_{k,j}(\theta)\|_{\mathcal{H}_\kappa}.$$

Then

$$\sup_\theta |\widehat{\mathcal{R}}^{(-k)}_{V,0,j}(\theta) - \widehat{\mathcal{R}}^{o,(-k)}_{V,0,j}(\theta)| \le D_{n,k,j}(2M_{n,k,j} + D_{n,k,j}). \tag{59}$$

In particular, $D_{n,k,j} = o_p(1)$ and $M_{n,k,j} = O_p(1)$ imply that the left-hand side is $o_p(1)$.

Proof. Apply $|\|x\|^2 - \|y\|^2| \le \|x - y\|(2\|y\| + \|x - y\|)$ to each pair of embeddings and take the supremum. □

C.4 Correction of the observation transformation

C.4.1 IMP: correction of the training criterion

Assumption C.6 (Uniform residual-imputation consistency)**.** *The inner cross-fitted regression, including its empty-donor convention, satisfies*

$$\max_{k,r} \sup_{\theta \in \Theta} \frac{1}{|\mathcal{B}_{kr}|} \sum_{i \in \mathcal{B}_{kr}} (1 - \delta_{W,i}) \{\widehat{m}^{(-k,-r)}_{r_\theta}(L_i^+) - m_{r_\theta}(L_i^+)\}^2 = o_p(1). \tag{60}$$

This is a uniform estimation condition on the residual regression; bandwidth selection alone does not establish it. For $s \in \{U, V\}$, write $\widehat{\mathcal{R}}^{(-k)}_{s,0,\mathrm{imp}}$ for the corresponding unpenalized loss from Equation 35.

Lemma C.5 (IMP approximation to the full-data criterion)**.** *Under Assumption 2, Assumptions C.1–C.3, and Assumption C.6, for $s \in \{U, V\}$,*

$$\max_k \sup_\theta |\widehat{\mathcal{R}}^{(-k)}_{s,0,\mathrm{imp}}(\theta) - \widehat{\mathcal{R}}^{\mathrm{full},(-k)}_s(\theta)| = o_p(1). \tag{61}$$

Thus Equation 53 holds for IMP.

Proof. Fix k and put $r^o_{\mathrm{imp},\theta,i} = \delta_{W,i} r_{\theta,i} + (1 - \delta_{W,i}) m_{r_\theta}(L_i^+)$. By MAR this equals $\mathbb{E}(r_{\theta,i} \mid O_i)$. Conditional on $\mathcal{O}_k = \sigma(O_i : i \in \mathcal{T}_k)$, the missing components of distinct observations are independent, while $\kappa(L_i, L_\ell)$ is measurable. Hence the oracle U-statistic satisfies

$$\widehat{\mathcal{R}}^{o,(-k)}_{U,0,\mathrm{imp}}(\theta) = \mathbb{E}\{\widehat{\mathcal{R}}^{\mathrm{full},(-k)}_U(\theta) \mid \mathcal{O}_k\}. \tag{62}$$

Let $B_{n,k,U}$ be the full-data deviation in Equation 52. Conditional Jensen's inequality and Lemma C.3 imply

$$\mathbb{E} \sup_\theta |\widehat{\mathcal{R}}^{o,(-k)}_{U,0,\mathrm{imp}}(\theta) - \mathcal{R}_\kappa(\theta)| \le \mathbb{E} B_{n,k,U} \to 0.$$

Now let $\eta_{k,i}(\theta) = \widehat{r}^{(-k)}_{\mathrm{imp},i}(\theta) - r^o_{\mathrm{imp},\theta,i}$ and $d_{n,k} = \sup_\theta N_k^{-1} \sum_{i \in \mathcal{T}_k} \eta_{k,i}(\theta)^2$. Inner-fold averaging and Equation 60 give $d_{n,k} = o_p(1)$. With $F_{\mathrm{imp}} = \mathbb{E}(F_r \mid O)$, conditional Jensen's inequality gives $\sup_\theta |r^o_{\mathrm{imp},\theta}| \le F_{\mathrm{imp}}$ and $\mathbb{E}(F^2_{\mathrm{imp}}) \le \mathbb{E}(F_r^2)$. Thus $Q_{n,k} = N_k^{-1} \sum_{i \in \mathcal{T}_k} F^2_{\mathrm{imp},i} = O_p(1)$. Expansion of the pairwise products and Cauchy–Schwarz give

$$\sup_\theta |\widehat{\mathcal{R}}^{(-k)}_{U,0,\mathrm{imp}}(\theta) - \widehat{\mathcal{R}}^{o,(-k)}_{U,0,\mathrm{imp}}(\theta)| \le 2C_\kappa \sqrt{Q_{n,k} d_{n,k}} + C_\kappa d_{n,k} = o_p(1).$$

The fitted–oracle, oracle–population, and population–full-data differences therefore prove Equation 61 for U. Moreover,

$$\sup_\theta \frac{1}{N_k} \sum_{i \in \mathcal{T}_k} \{\widehat{r}^{(-k)}_{\mathrm{imp},i}(\theta)\}^2 \le 2Q_{n,k} + 2d_{n,k} = O_p(1).$$

The deterministic U/V bound in Equation 48, applied to both fitted and full-data residuals, proves the V case. This also proves the V-statistic extension without treating fitted residuals as independent observations. Finally use Equation 52 and fixed K. □

C.4.2 IPW: CORRECTION OF THE TRAINING CRITERION

Assumption C.7 (Cross-fitted observation-probability estimation)**.** *The lower-truncated probabilities satisfy, for every outer fold,*

$$\frac{1}{N_k} \sum_{i \in \mathcal{T}_k} \{\bar{e}_{ki} - e(L_i^+)\}^2 = o_p(1), \tag{63}$$

$$\frac{1}{|\mathcal{I}_k|} \sum_{i \in \mathcal{I}_k} \{\bar{e}^{\mathrm{eval}}_k(L_i^+) - e(L_i^+)\}^2 = o_p(1). \tag{64}$$

The training assertion uses only the first condition; the second is needed for causal-mean evaluation. The common lower bound $\underline{e} > 0$ is enforced by construction.

Assumption C.8 (Oracle IPW complexity)**.** *The class $\mathcal{Z}_{\mathrm{ipw}} = \{r^o_{\mathrm{ipw},\theta} = \delta_W r_\theta / e(L^+) : \theta \in \Theta\}$ satisfies Equation 44.*

Proposition C.1 (IPW approximation to the full-data criterion)**.** *Suppose Assumption 2, Assumptions C.1–C.3, Equation 63, and Assumption C.8 hold. For $s \in \{U, V\}$,*

$$\max_k \sup_\theta |\widehat{\mathcal{R}}^{(-k)}_{s,0,\mathrm{ipw}}(\theta) - \widehat{\mathcal{R}}^{\mathrm{full},(-k)}_s(\theta)| = o_p(1), \tag{65}$$

$$\max_k \sup_\theta |\widehat{\mathcal{R}}^{(-k)}_{s,0,\mathrm{ipw}}(\theta) - \mathcal{R}_\kappa(\theta)| = o_p(1). \tag{66}$$

Proof. The oracle residual has envelope F_r/ε. By population recovery and Lemma C.2,

$$\sup_\theta |\widehat{\mathcal{R}}^{o,(-k)}_{V,0,\mathrm{ipw}}(\theta) - \mathcal{R}_\kappa(\theta)| = o_p(1). \tag{67}$$

The empirical embeddings obey

$$\begin{aligned} D_{n,k,\mathrm{ipw}} &\le \frac{\sqrt{C_\kappa}}{\underline{e}\varepsilon}\left(\frac{1}{N_k}\sum_{i\in\mathcal{T}_k} F_{r,i}^2\right)^{1/2}\left(\frac{1}{N_k}\sum_{i\in\mathcal{T}_k}\{\bar{e}_{ki} - e(L_i^+)\}^2\right)^{1/2} = o_p(1), \\ M_{n,k,\mathrm{ipw}} &\le \frac{\sqrt{C_\kappa}}{\varepsilon}\frac{1}{N_k}\sum_{i\in\mathcal{T}_k} F_{r,i} = O_p(1). \end{aligned} \tag{68}$$

Lemma C.4 controls the fitted–oracle criterion difference. Combining it with Equation 67 and the full-data benchmark proves Equation 65 for V. Since the fitted residuals have envelope $F_r/\underline{e}$, the U/V difference is $O_p(N_k^{-1})$ by Equation 48; the same holds for the full-data comparator. This proves the U case. Equation 52 then gives Equation 66. □

C.4.3 AIPW: CORRECTION OF THE TRAINING CRITERION AND DOUBLE ROBUSTNESS

For $i \in \mathcal{B}_{kr}$, write $\bar{e}_{kr}(l^+) = \max\{\widehat{e}^{(-k,-r)}(l^+), \underline{e}\}$, and define

$$r^o_{\mathrm{aipw},\theta} = m_{r_\theta}(L^+) + \frac{\delta_W}{e(L^+)}\{r_\theta - m_{r_\theta}(L^+)\}, \qquad \Delta_{kr,\theta} = \widehat{r}^{(-k,-r)}_{\mathrm{aipw}}(\theta) - r^o_{\mathrm{aipw},\theta}.$$

Here $\widehat{r}^{(-k,-r)}_{\mathrm{aipw}}$ is the fitted transformation evaluated on an independent observation with the inner-training fits held fixed. Let $\mathcal{F}_{kr}$ contain these fits, their training observations, the partitions, and independent fitting randomization. For such an independent observation, MAR gives

$$\mathbb{E}(\Delta_{kr,\theta} \mid L^+, \mathcal{F}_{kr}) = \frac{\bar{e}_{kr} - e}{\bar{e}_{kr}}\{\widehat{m}^{(-k,-r)}_{r_\theta} - m_{r_\theta}\}(L^+). \tag{69}$$

Indeed, the conditional mean of the fitted AIPW residual is $\widehat{m} + (e/\bar{e}_{kr})(m - \widehat{m})$, while the oracle mean is m.

Assumption C.9 (AIPW nuisance remainder and complexity)**.** *The oracle class $\mathcal{Z}_{\mathrm{aipw}} = \{r^o_{\mathrm{aipw},\theta}\}$ satisfies Equation 44, and*

$$\max_{k,r}\sup_\theta \left\| \frac{\bar{e}_{kr} - e}{\bar{e}_{kr}}(\widehat{m}^{(-k,-r)}_{r_\theta} - m_{r_\theta}) \right\|_{L_1(P_{L^+})} = o_p(1), \tag{70}$$

$$\max_{k,r}\sup_\theta \|(\mathbb{P}_{\mathcal{B}_{kr}} - P)[\Delta_{kr,\theta}\kappa(L,\cdot)]\|_{\mathcal{H}_\kappa} = o_p(1). \tag{71}$$

Here P integrates an independent observation with the fits held fixed, and $\mathbb{P}_{\mathcal{B}_{kr}}$ averages the corresponding inner evaluation observations. The displayed functions are integrable.

Under the common probability floor, the product condition follows from either of the following uniform $L_2(P_{L^+})$ conditions:

$$\begin{aligned} &\max_{k,r}\|\bar{e}_{kr} - e\|_2 = o_p(1), && \max_{k,r}\sup_\theta \|\widehat{m}^{(-k,-r)}_{r_\theta} - m_{r_\theta}\|_2 = O_p(1), \\ &\max_{k,r}\|\bar{e}_{kr} - e\|_2 = O_p(1), && \max_{k,r}\sup_\theta \|\widehat{m}^{(-k,-r)}_{r_\theta} - m_{r_\theta}\|_2 = o_p(1). \end{aligned}$$

The centered empirical-process condition is additionally required in both branches; cross-fitting alone does not imply uniform convergence over Θ. The oracle residual is bounded by the square-integrable envelope

$$F_{\mathrm{aipw}} = \varepsilon^{-1}F_r + (1+\varepsilon^{-1})\mathbb{E}(F_r \mid L^+).$$

Thus Lemma C.2 and population recovery imply

$$\max_k \sup_\theta |\widehat{\mathcal{R}}^{o,(-k)}_{V,0,\mathrm{aipw}}(\theta) - \mathcal{R}_\kappa(\theta)| = o_p(1). \tag{72}$$

Proposition C.2 (AIPW approximation to the full-data criterion)**.** *Under Assumption 2, Assumptions C.1–C.3, and Assumption C.9,*

$$\max_k \sup_\theta |\widehat{\mathcal{R}}^{(-k)}_{V,0,\mathrm{aipw}}(\theta) - \widehat{\mathcal{R}}^{\mathrm{full},(-k)}_{V}(\theta)| = o_p(1), \tag{73}$$

$$\max_k \sup_\theta |\widehat{\mathcal{R}}^{(-k)}_{V,0,\mathrm{aipw}}(\theta) - \mathcal{R}_\kappa(\theta)| = o_p(1). \tag{74}$$

The same conclusions hold for U if additionally

$$\max_k \sup_\theta \frac{1}{N_k} \sum_{i \in \mathcal{T}_k} \{\widehat{r}^{(-k)}_{\mathrm{aipw},i}(\theta)\}^2 = O_p(1). \tag{75}$$

Proof. Decompose each inner-fold empirical embedding into its conditional mean and centered part. Equations 69 and 70 give

$$\sup_\theta \|P[\Delta_{kr,\theta}\kappa(L,\cdot)]\|_{\mathcal{H}_\kappa} \le \sqrt{C_\kappa} \sup_\theta \left\| \frac{\bar{e}_{kr} - e}{\bar{e}_{kr}} (\widehat{m}^{(-k,-r)}_{r_\theta} - m_{r_\theta}) \right\|_1 = o_p(1).$$

Consequently, using the nonnegative inner-fold weights,

$$D_{n,k,\mathrm{aipw}} \le \sum_{r=1}^{R} \frac{|\mathcal{B}_{kr}|}{N_k} \sup_\theta \|\mathbb{P}_{\mathcal{B}_{kr}}[\Delta_{kr,\theta}\kappa(L,\cdot)]\|_{\mathcal{H}_\kappa} = o_p(1).$$

The envelope F_{aipw} gives $M_{n,k,\mathrm{aipw}} \le \sqrt{C_\kappa}\mathbb{P}_{\mathcal{T}_k} F_{\mathrm{aipw}} = O_p(1)$. Combine Lemma C.4, Equation 72, and Lemma C.3 to obtain the first conclusion. Equation 52 gives the second. For the U extension, Equation 75 and the deterministic bound 48 control the fitted diagonal; the full-data diagonal is already controlled. None of these steps requires independence between corrected residuals in different inner folds. □

C.5 Recovery of the fixed-treatment causal mean

The training results establish the first requirement of Theorem C.1. We now verify its held-out evaluation requirement, retaining the distinction between the residual regression m_{r_θ} and the target-treatment regression $q_{a,\theta}$.

Assumption C.10 (IMP target-treatment regression)**.** *At the fixed a, there is a nonnegative envelope $F_a(W,X)$ such that*

$$\sup_\theta |h_\theta(W,a,X)| \le F_a(W,X) \quad a.s., \qquad \mathbb{E}(F_a^2) < \infty.$$

The outer-training regression satisfies

$$\max_k \|\widehat{q}^{(-k)}_{a,\widehat{\theta}^{(-k)}_{\mathrm{imp}}} - q_{a,\widehat{\theta}^{(-k)}_{\mathrm{imp}}}\|_{L_2(P_{L^+})} = o_p(1).$$

The norm uses an independent population observation, with the fitted bridge and nuisance functions held fixed.

Assumption C.11 (IPW target-treatment envelope)**.** *At the fixed a, the envelope condition in Assumption C.10 holds.*

An additional uniform target law is unnecessary here: the bridge is fitted outside $\mathcal{I}_k$, so a conditional law of large numbers applies directly to that fitted bridge.

Assumption C.12 (AIPW fixed-treatment recovery)**.** *At the fixed a, the envelope condition in Assumption C.10 holds. For $\widehat{\theta}_k = \widehat{\theta}^{(-k)}_{\mathrm{aipw}}$,*

$$\max_k \left\| \frac{\bar{e}^{\mathrm{eval}}_k - e}{\bar{e}^{\mathrm{eval}}_k} (\widehat{q}^{(-k)}_{a,\widehat{\theta}_k} - q_{a,\widehat{\theta}_k}) \right\|_{L_1(P_{L^+})} = o_p(1), \tag{76}$$

$$\max_k \|\widehat{q}^{(-k)}_{a,\widehat{\theta}_k}\|_{L_2(P_{L^+})} = O_p(1). \tag{77}$$

The product condition follows if either the observation probability or the target regression is L_2-consistent and the other estimation error is L_2-bounded. This is evaluation-stage double robustness; the training-stage conditions above must also hold.

Proposition C.3 (Held-out target recovery)**.** *Suppose Assumption 2 and the stated sample-splitting conditions hold. Equation 55 holds*

$$\max_k |\widehat{\beta}_{k,j}(a) - b_a(\widehat{\theta}_j^{(-k)})| = o_p(1).$$

for: (i) IMP under Assumption C.10; (ii) IPW under Equation 64 and Assumption C.11; and (iii) AIPW under Assumption C.12.

Proof. Let $\mathcal{F}_k$ contain the observed training sample on $\mathcal{T}_k$, the partitions, and independent fitting randomization. Conditional on $\mathcal{F}_k$, all fitted functions for fold k are fixed and the evaluation observations remain i.i.d. with their population law. Write $\widehat{h}_k(W,X) = h_{\widehat{\theta}_j^{(-k)}}(W,a,X)$, and write q_k, $\widehat{q}_k$ for its true and fitted conditional means given the factual L^+.

For IMP, let $G_k = \delta_W \widehat{h}_k(W,X) + (1-\delta_W)\widehat{q}_k(L^+)$. If $\zeta_n = \max_k \|\widehat{q}_k - q_k\|_2 = o_p(1)$, Jensen's inequality implies $\max_k \|\widehat{q}_k\|_2 \le \|F_a\|_2 + \zeta_n$. Hence $\max_k \mathbb{E}(G_k^2 \mid \mathcal{F}_k) = O_p(1)$, and conditional Chebyshev's inequality gives

$$\max_k |\mathbb{P}_{\mathcal{I}_k} G_k - PG_k| = o_p(1). \tag{78}$$

Here and below P holds the fits fixed. MAR yields

$$PG_k - b_a(\widehat{\theta}_{\mathrm{imp}}^{(-k)}) = P[(1-e)(\widehat{q}_k - q_k)],$$

whose absolute value is at most ζ_n.

For IPW, first use the oracle evaluation summand $G_k^o = \delta_W \widehat{h}_k(W,X)/e(L^+)$. Its conditional mean is $b_a(\widehat{\theta}_{\mathrm{ipw}}^{(-k)})$ and its conditional second moment is at most $\varepsilon^{-2}P(F_a^2)$. Conditional Chebyshev therefore gives

$$\max_k |\mathbb{P}_{\mathcal{I}_k} G_k^o - b_a(\widehat{\theta}_{\mathrm{ipw}}^{(-k)})| = o_p(1). \tag{79}$$

Replacing e by $\bar{e}_k^{\mathrm{eval}}$ changes this average by at most

$$\frac{1}{\underline{e}\varepsilon}\{\mathbb{P}_{\mathcal{I}_k} F_a^2\}^{1/2}\left[\mathbb{P}_{\mathcal{I}_k}\{\bar{e}_k^{\mathrm{eval}}(L^+) - e(L^+)\}^2\right]^{1/2} = o_p(1).$$

For AIPW, set $G_k = \widehat{q}_k(L^+) + \delta_W\{\widehat{h}_k(W,X) - \widehat{q}_k(L^+)\}/\bar{e}_k^{\mathrm{eval}}(L^+)$. The probability floor and Equation 77 give

$$P(G_k^2) \le 2\underline{e}^{-2}P(F_a^2) + 2(1+\underline{e}^{-1})^2 P(\widehat{q}_k^2) = O_p(1).$$

Consequently,

$$\max_k |\mathbb{P}_{\mathcal{I}_k} G_k - PG_k| = o_p(1). \tag{80}$$

The same conditional-mean calculation as in Equation 69 yields

$$PG_k - b_a(\widehat{\theta}_{\mathrm{aipw}}^{(-k)}) = P\left[\frac{\bar{e}_k^{\mathrm{eval}} - e}{\bar{e}_k^{\mathrm{eval}}}(\widehat{q}_k - q_k)\right] = o_p(1).$$

In all three cases, finite K and diverging evaluation-fold sizes justify the maximum over folds. More explicitly, an $O_p(1)$ conditional second moment gives an unconditional $o_p(1)$ fluctuation by first restricting to an event where that moment is bounded and then applying conditional Chebyshev. This completes the proof. □

C.6 Consistency of the three ProxMAR–NMMR estimators

Theorem C.2 (Consistency of ProxMAR–NMMR–IMP)**.** *Assume the conditions of Lemma C.5, Assumption C.4 with $\|\theta_0\|_2 < \infty$, and Equation 54 for IMP with $\lambda_n \ge 0$, $\lambda_n \to 0$. Then for either U or V, $\max_k \mathcal{R}_\kappa(\widehat{\theta}_{\mathrm{imp}}^{(-k)}) = o_p(1)$. At a fixed a, if Assumptions 1, C.5, and C.10 also hold, with $h_0 = h_{\theta_0}$ and the pointwise conditions in Appendix A.1 when A is continuous, then $\widehat{\beta}_{\mathrm{NMMR,imp}}(a) \xrightarrow{p} \beta(a)$.*

Proof. Lemma C.5 and the decomposition in Equation 52 give Equation 53. Proposition C.3(i) gives Equation 55. Apply Theorem C.1. □

Theorem C.3 (Consistency of ProxMAR–NMMR–IPW)**.** *Assume the conditions of Proposition C.1, Assumption C.4 with $\|\theta_0\|_2 < \infty$, and Equation 54 for IPW with $\lambda_n \geq 0$, $\lambda_n \to 0$. Then for either U or V, $\max_k \mathcal{R}_\kappa(\widehat{\theta}^{(-k)}_{\mathrm{ipw}}) = o_p(1)$. At a fixed a, additionally assume Assumptions 1, C.5, and C.11, Equation 64, and the same identification and pointwise conditions as in Theorem C.2. Then $\widehat{\beta}_{\mathrm{NMMR,ipw}}(a) \xrightarrow{p} \beta(a)$.*

Proof. Use Proposition C.1, Proposition C.3(ii), and Theorem C.1. □

Theorem C.4 (Consistency of ProxMAR–NMMR–AIPW)**.** *Assume the conditions of Proposition C.2, Assumption C.4 with $\|\theta_0\|_2 < \infty$, and Equation 54 for AIPW with $\lambda_n \geq 0$, $\lambda_n \to 0$. Then the V-statistic estimator satisfies $\max_k \mathcal{R}_\kappa(\widehat{\theta}^{(-k)}_{\mathrm{aipw}}) = o_p(1)$. The U-statistic conclusion additionally requires Equation 75. At a fixed a, if Assumptions 1, C.5, and C.12 hold, with $h_0 = h_{\theta_0}$ and the pointwise conditions in Appendix A.1 when A is continuous, then $\widehat{\beta}_{\mathrm{NMMR,aipw}}(a) \xrightarrow{p} \beta(a)$.*

Proof. Use Proposition C.2, Proposition C.3(iii), and Theorem C.1. □

These conclusions concern the population kernel risk and the identified causal mean. They do not require convergence of network parameters or L_2 convergence of the entire bridge. An L_2 bridge conclusion would require a corresponding bridge-separation or inverse-stability condition.

D ProxMAR-DFPV

D.1 Detailed estimation procedure and algorithm

D.1.1 Cross-fitted two-stage estimation objectives

The DFPV proxy feature map links the two estimation stages. Write $\psi_\omega(W)$ for this map and define the Stage 1 product feature and regression by

$$\Phi_\vartheta(L) := \phi_{A(1)}(A) \otimes \phi_Z(Z) \otimes \phi_{X(1)}(X), \qquad \mu_{V,\vartheta}(L) := V\Phi_\vartheta(L),$$

where ϑ collects the parameters of the three Stage 1 feature maps. The matrix V maps the product feature into the dimension of $\psi_\omega(W)$. Stage 2 uses separate treatment and covariate maps ψ_A and ϕ_X in the bridge parameterization in Equation 11.

Fix a transformation type $j \in \mathcal{J} := \{\mathrm{imp}, \mathrm{ipw}, \mathrm{aipw}\}$. At the current proxy-map parameter ω, define $m_\omega(L^+) := \mathbb{E}\{\psi_\omega(W) \mid L^+\}$. On each inner training sample $I_{-k,-r}$, fit $\widehat{m}^{(-k,-r)}_\omega$ when $j \in \{\mathrm{imp}, \mathrm{aipw}\}$ and fit $\widehat{e}^{(-k,-r)}$ when $j \in \{\mathrm{ipw}, \mathrm{aipw}\}$. The notation m_ω abbreviates m_{ψ_ω} in the main text. For $i \in J_{kr}$, evaluate the estimated transformations in Equation 7 at the proxy feature ψ_ω:

$$\begin{aligned}
\widehat{T}^{(-k,-r)}_{\mathrm{imp},i}(\psi_\omega) &:= \delta_{W,i}\psi_\omega(W_i) + (1-\delta_{W,i})\widehat{m}^{(-k,-r)}_\omega(L^+_i), \\
\widehat{T}^{(-k,-r)}_{\mathrm{ipw},i}(\psi_\omega) &:= \frac{\delta_{W,i}}{\widehat{e}^{(-k,-r)}(L^+_i)}\psi_\omega(W_i), \\
\widehat{T}^{(-k,-r)}_{\mathrm{aipw},i}(\psi_\omega) &:= \widehat{m}^{(-k,-r)}_\omega(L^+_i) + \frac{\delta_{W,i}}{\widehat{e}^{(-k,-r)}(L^+_i)}\left\{\psi_\omega(W_i) - \widehat{m}^{(-k,-r)}_\omega(L^+_i)\right\}.
\end{aligned} \tag{81}$$

Every observation in I_{-k} receives one response from nuisance estimators trained on its complementary inner sample. The three variants differ only in this response and its required nuisance functions:

Variant	Stage 1 response	Fitted nuisance functions
DFPV-IMP	$\widehat{T}_{\mathrm{imp}}(\psi_\omega)$	$\widehat{m}_\omega$
DFPV-IPW	$\widehat{T}_{\mathrm{ipw}}(\psi_\omega)$	$\widehat{e}$
DFPV-AIPW	$\widehat{T}_{\mathrm{aipw}}(\psi_\omega)$	$\widehat{e}, \widehat{m}_\omega$

Terms containing W_i are evaluated only when $\delta_{W,i} = 1$.

Holding ω fixed, combine the inner-fold responses and fit Stage 1 through

$$\begin{aligned}&\left(\widehat{V}_j^{(-k)}(\omega), \widehat{\vartheta}_j^{(-k)}(\omega)\right)\\ &\quad\in \arg\min_{V,\vartheta}\left\{\frac{1}{|I_{-k}|}\sum_{r=1}^{R}\sum_{i\in J_{kr}}\left\|\widehat{T}_{j,i}^{(-k,-r)}(\psi_\omega) - \mu_{V,\vartheta}(L_i)\right\|_2^2 + \lambda_1\|V\|_F^2\right\}.\end{aligned} \tag{82}$$

The resulting conditional proxy-feature estimate is

$$\widehat{\mu}_{\omega,j}^{(-k)}(L) := \widehat{V}_j^{(-k)}(\omega)\Phi_{\widehat{\vartheta}_j^{(-k)}(\omega)}(L). \tag{83}$$

The fitted nuisance functions and $\widehat{\mu}_{\omega,j}^{(-k)}$ are indexed by the current proxy map and are refitted after each update of ω.

For Stage 2, define

$$\widehat{s}_{\psi_A,\phi_X,\omega,j}^{(-k)}(L) := \psi_A(A)\otimes\phi_X(X)\otimes\widehat{\mu}_{\omega,j}^{(-k)}(L).$$

The fold-specific Stage 2 parameters solve

$$\left(\widehat{u}_j^{(-k)}, \widehat{\psi}_{A,j}^{(-k)}, \widehat{\phi}_{X,j}^{(-k)}, \widehat{\omega}_j^{(-k)}\right) \in \arg\min_{u,\psi_A,\phi_X,\omega}\left\{\frac{1}{|I_{-k}|}\sum_{i\in I_{-k}}\left[Y_i - u^\top \widehat{s}_{\psi_A,\phi_X,\omega,j}^{(-k)}(L_i)\right]^2 + \lambda_2\|u\|_2^2\right\}. \tag{84}$$

The proxy parameter ω enters this criterion through the refitted Stage 1 map in Equation 83.

D.1.2 Parameter optimization and causal-mean estimation

The proxy map is initialized through an auxiliary complete-case regression of Y on $\{\psi_\omega(W), A, X\}$. Its fitted proxy-network parameters initialize the alternating Stage 1 and Stage 2 updates. At every epoch, each conditional-mean nuisance regression uses $\psi_\omega(W)$ at the current ω as its response.

Each epoch uses the selected feature response in Equation 81 twice. During the Stage 1 feature update, ω is fixed and the computed feature responses are fixed regression responses. During the proxy-map update, the feature responses are recomputed at the current ω. The fitted nuisance parameters, including any observation probabilities and Gaussian smoothing weights, and the Stage 1 feature maps are held fixed. Differentiation propagates through $\psi_\omega(W_i)$ for all three variants and through $\widehat{m}_\omega^{(-k,-r)}(L_i^+)$ for DFPV-IMP and DFPV-AIPW. The ridge coefficient V is recomputed from these differentiable responses, and the Stage 2 gradient propagates through the resulting $\widehat{\mu}_{\omega,j}^{(-k)}$. This update follows the dependence of the fitted Stage 1 coefficients on the proxy map in the original DFPV construction (Xu et al., 2021).

For fixed neural feature parameters, V and u are recomputed as the ridge minimizers in Equations 82 and 84. The fitted value of u is then held fixed while the Stage 2 feature parameters and ω are updated. Each outer fold initializes and trains its own feature maps and ridge coefficients.

After the final epoch, every required inner-fold nuisance regression is refitted at $\widehat{\omega}_j^{(-k)}$, the feature responses are reconstructed, and Equation 82 is solved again. Equation 84 is then solved for $\widehat{u}_j^{(-k)}$ with the final feature maps and refitted Stage 1 map held fixed. These final refits align all nuisance, Stage 1, and Stage 2 quantities at the same proxy map.

Let $\mathcal{A}_{\text{eval}} \subseteq \mathcal{A}$ denote the treatment values at which the causal mean is estimated. For each outer fold, define

$$\widehat{M}_{k,j} := \frac{1}{|I_k|}\sum_{i\in I_k}\left\{\widehat{\phi}_{X,j}^{(-k)}(X_i)\otimes\widehat{\mu}_{\widehat{\omega}_j^{(-k)},j}^{(-k)}(L_i)\right\}. \tag{85}$$

For each fixed treatment value a, define the cross-fitted causal-mean estimator by

$$\widehat{\beta}_{\text{DFPV},j}(a) = \sum_{k=1}^{K}\frac{|I_k|}{n}\,\widehat{u}_j^{(-k)\top}\left\{\widehat{\psi}_{A,j}^{(-k)}(a)\otimes\widehat{M}_{k,j}\right\}, \qquad a\in\mathcal{A}_{\text{eval}}. \tag{86}$$

D.1.3 ALGORITHM

Algorithm 2 DFPV-IMP, DFPV-IPW, and DFPV-AIPW

Require: Observed data; transformation type $j \in \mathcal{J}$; outer and inner fold counts K, R; penalties λ_1, λ_2; epoch budget $T_{\max}$; target treatments $\mathcal{A}_{\text{eval}}$

Partition $\{1, \dots, n\}$ into outer folds $I_1, \dots, I_K$ independently of the observations and keep the partition fixed
for $k = 1, \dots, K$ **do**
  Reserve I_k for causal-mean evaluation and set $I_{-k} = \{1, \dots, n\} \setminus I_k$
  Partition I_{-k} into inner folds $J_{k1}, \dots, J_{kR}$ independently of the observations; set $I_{-k,-r} = I_{-k} \setminus J_{kr}$ and keep the partition fixed
  Initialize the fold-specific feature maps and pretrain ω on complete cases in I_{-k}
  for $t = 1, \dots, T_{\max}$ **do**
    for $r = 1, \dots, R$ **do**
      Fit $\widehat{m}_{\omega}^{(-k,-r)}$ for $j \in \{\text{imp}, \text{aipw}\}$ and $\widehat{e}^{(-k,-r)}$ for $j \in \{\text{ipw}, \text{aipw}\}$ on $I_{-k,-r}$
      Compute $\widehat{T}_{j,i}^{(-k,-r)}(\psi_\omega)$ on J_{kr} using Equation 81
    end for
    Form the Stage 1 pairs over all inner folds; every $i \in I_{-k}$ appears once
    Hold ω and the feature responses fixed; update the Stage 1 feature maps and compute V using Equation 82
    Recompute the feature responses at the current ω; hold the fitted nuisance parameters, smoothing weights, and Stage 1 feature maps fixed; retain the response derivatives through ω; recompute V and $\widehat{\mu}_{\omega,j}^{(-k)}$
    Hold the current Stage 2 feature maps, ω, and $\widehat{\mu}_{\omega,j}^{(-k)}$ fixed; compute u using Equation 84
    Hold u at its fitted value; differentiate the Stage 2 loss through $\widehat{\mu}_{\omega,j}^{(-k)}$ and update the Stage 2 feature maps and ω
  end for
  for $r = 1, \dots, R$ **do**
    At $\widehat{\omega}_j^{(-k)}$, refit the nuisance functions required by j on $I_{-k,-r}$
    Recompute $\widehat{T}_{j,i}^{(-k,-r)}(\psi_{\widehat{\omega}_j^{(-k)}})$ for every $i \in J_{kr}$
  end for
  Refit $\widehat{\mu}_{\widehat{\omega}_j^{(-k)},j}^{(-k)}$ using all inner-fold pairs and Equation 82
  Recompute $\widehat{u}_j^{(-k)}$ using Equation 84 with the final feature maps and Stage 1 map fixed
  Compute $\widehat{M}_{k,j}$ on I_k using Equation 85
end for
For each $a \in \mathcal{A}_{\text{eval}}$, combine the outer-fold contributions using Equation 86
return $\{\widehat{\beta}_{\text{DFPV},j}(a) : a \in \mathcal{A}_{\text{eval}}\}$

Experiment-specific epoch budgets, validation schedules, and stopping rules are reported with the corresponding designs in Appendices E.1 and E.2.

D.2 SHARED PROXMAR-DFPV ASSUMPTIONS

The original proxy satisfies $W \in \mathcal{W} \subseteq \mathbb{R}^{p_W}$, whereas $\psi_\omega(W) \in \mathbb{R}^{p_\omega}$. Thus p_W is the dimension of W and p_ω is the proxy-feature dimension. The Stage 1 feature dimension is $p_{\omega,1}$, so $V \in \mathbb{R}^{p_\omega \times p_{\omega,1}}$. The Stage 2 coefficient has dimension $p_A p_X p_\omega$.

For each sample size n, let Ω_n be the deterministic set of admissible parameter vectors ω for the proxy map ψ_ω. It specifies the allowed weights and biases of the chosen network architecture, including any constraints. The architecture and constraints may vary with n. The associated hypothesis class is

$$\mathcal{H}_{\psi,n} := \{\psi_\omega : \mathcal{W}_n \to \mathbb{R}^{p_\omega} : \omega \in \Omega_n\}.$$

Following the function-class description in Xu et al. (2021, Appendix C.1), define the Stage 1 regression class and bridge class by

$$\mathcal{G}_n := \{\mu_{V,\vartheta} : L \mapsto V\Phi_\vartheta(L) : (V, \vartheta) \in \mathcal{V}_n \times \Theta_{1,n}\},$$
$$\mathcal{H}_{h,n} := \{h_\eta : \eta \in \Xi_n\}, \qquad h_\eta(w, a, x) = u^\top\{\psi_A(a) \otimes \phi_X(x) \otimes \psi_\omega(w)\}.$$

Here $\Theta_{1,n}$ is the deterministic set of admissible Stage 1 feature parameters, $\mathcal{V}_n \subseteq \mathbb{R}^{p_\omega \times p_{\omega,1}}$ is the set of admissible Stage 1 coefficient matrices, and Ξ_n contains all admissible bridge parameters: u, the parameters of ψ_A and ϕ_X, and $\omega \in \Omega_n$. These sets are specified before fitting, and all folds use the sets for the total sample size n. The fitted parameters belong to these sets. All functions and the suprema below are assumed measurable. Elements of $\mathcal{G}_n$ are vector-valued regression functions denoted by $\mu_{V,\vartheta}$.

Let P denote the complete-data law and its induced marginal laws, and write $Pf = \mathbb{E}\{f(O)\}$ for an integrable function of an observed record. Expectations involving W are taken under the corresponding complete-data law. For fitted functions, P integrates over a fresh observation independent of the training data, with the fitted functions held fixed. Define

$$P_{n,-k}f := \frac{1}{|I_{-k}|}\sum_{i\in I_{-k}} f(O_i), \quad P_{n,kr}f := \frac{1}{|J_{kr}|}\sum_{i\in J_{kr}} f(O_i), \quad P_{n,k}f := \frac{1}{|I_k|}\sum_{i\in I_k} f(O_i).$$

For a vector-valued F, set $\|F\|_{L_q(P)} = (P\|F\|_2^q)^{1/q}$, $q \in \{2,4\}$, using the appropriate marginal law for its arguments. The matrix norms $\|\cdot\|_F$ and $\|\cdot\|_{\mathrm{op}}$ are the Frobenius and operator norms. If the original data dimension varies with n, all laws and targets in each bound are evaluated at that same n; their additional index is suppressed.

Sampling and cross-fitting used throughout the proof. We use the independent and identically distributed sampling setup of the main text and invoke Assumption 2 for MAR and observation positivity. The additional sample-splitting requirements are as follows. The integers $K, R \geq 2$ are fixed, the partitions are fixed in advance or generated independently of the observations, and all outer evaluation and inner training/evaluation sample sizes tend to infinity. Both stages for outer fold k use only I_{-k}, and the final average uses only I_k. For every fixed deterministic $\omega \in \Omega_n$, the nuisance-fitting rule for fold (k,r) uses only $I_{-k,-r}$ and independent fitting randomness. All final refits use the same fitted proxy map $\widehat{\omega}_j^{(-k)}$.

Let $\mathcal{F}_{-k,j}$ be the sigma-field generated by the outer training observations, fold assignments, and fitting randomness. Let $\mathcal{F}_{-k,-r,j}$ contain the corresponding inner training information and the nuisance-fitting rules for fixed candidates ω. Conditional on the fold assignments, the outer evaluation sample is independent of $\mathcal{F}_{-k,j}$, and the inner evaluation sample is independent of the nuisance-fitting rule at each fixed ω. Uniform bounds below permit substitution of a data-selected $\widehat{\omega}_j^{(-k)}$; no inner-fold independence is asserted after that substitution. These sampling and fitting requirements are maintained in the proposition and theorem below.

Assumption D.1 (DFPV moments)**.** *The feature maps and outcome satisfy*

$$\sup_n \sup_{\omega\in\Omega_n} P\|\psi_\omega(W)\|_2^2 < \infty, \quad \sup_n \sup_{\vartheta\in\Theta_{1,n}} P\|\Phi_\vartheta(L)\|_2^2 < \infty, \quad \sup_n PY^2 < \infty.$$

All population risks used below are finite.

For a scalar or vector function s of the complete data with possible dependence on the fully observed variables, set $m_s(L^+) = \mathbb{E}(s \mid L^+)$. The oracle transformations T_j are

$$\begin{aligned} T_{\mathrm{imp}}(s) &= \delta_W s + (1-\delta_W)m_s, \\ T_{\mathrm{ipw}}(s) &= \delta_W s/e, \\ T_{\mathrm{aipw}}(s) &= m_s + \delta_W(s-m_s)/e. \end{aligned}$$

Nuisance functions here are evaluated at L^+. Under MAR, $\mathbb{E}\{T_j(s) \mid L\} = \mathbb{E}(s \mid L)$. Equation 81 is the estimated transformation $\widehat{T}_j$ evaluated at $s = \psi_\omega(W)$; the function argument is abbreviated to ψ_ω, as in the main text. No separate response symbol is introduced.

For $m_\omega(L^+) = \mathbb{E}\{\psi_\omega(W) \mid L^+\}$, define

$$\rho_{n,kr,\omega,j}(L^+) := \begin{cases} (1-e)(\widehat{m}_\omega^{(-k,-r)} - m_\omega), & j = \mathrm{imp}, \\ \dfrac{e - \widehat{e}^{(-k,-r)}}{\widehat{e}^{(-k,-r)}} m_\omega, & j = \mathrm{ipw}, \\ \dfrac{\widehat{e}^{(-k,-r)} - e}{\widehat{e}^{(-k,-r)}}(\widehat{m}_\omega^{(-k,-r)} - m_\omega), & j = \mathrm{aipw}. \end{cases} \tag{87}$$

Every nuisance term on the right is evaluated at L^+. Write

$$R_{n,j} := \max_{k,r} \sup_{\omega \in \Omega_n} \|\rho_{n,kr,\omega,j}\|_{L_2(P)}.$$

Assumption D.2 (Variant-specific nuisance recovery)**.** *For the selected transformation type j, $R_{n,j} = o_p(1)$.*

Assumption D.3 (Bounds for the fitted nuisances)**.** *For $j \in \{\text{ipw}, \text{aipw}\}$, fitted probabilities used as denominators are strictly positive, and there is a constant $c > 0$ such that*

$$\Pr\left\{\min_{k,r} \inf_{l^+ \in \mathcal{L}^+} \widehat{e}^{(-k,-r)}(l^+) \geq c\right\} \longrightarrow 1,$$

where $\mathcal{L}^+$ is the support of L^+. For $j \in \{\text{imp}, \text{aipw}\}$,

$$\max_{k,r} \sup_{\omega \in \Omega_n} \|\widehat{m}_\omega^{(-k,-r)}\|_{L_2(P)} = O_p(1).$$

For IMP, uniform $L_2(P)$ convergence of $\widehat{m}_\omega$ implies nuisance recovery. For IPW, the required quantity is $(\widehat{e} - e)m_\omega/\widehat{e}$. For AIPW, Hölder's inequality bounds the remainder by

$$R_{n,\text{aipw}} \leq c^{-1} \max_{k,r} \left\{\|\widehat{e}^{(-k,-r)} - e\|_{L_4(P)} \sup_{\omega \in \Omega_n} \|\widehat{m}_\omega^{(-k,-r)} - m_\omega\|_{L_4(P)}\right\}$$

on the fitted-probability event. It suffices that one factor is $o_p(1)$ and the other is $O_p(1)$ uniformly over the indicated indices. This is the product-remainder form of double robustness (Chernozhukov et al., 2018). Bounds for a fixed regression target, such as those of Farrell et al. (2021, Theorems 1–2), do not by themselves establish uniform recovery over a learned proxy-map class.

D.3 Conditional-mean targets for the two DFPV stages

D.3.1 Recovery of the Stage 1 conditional mean

Proposition D.1 (ProxMAR-DFPV Stage 1 target recovery)**.** *Under the sampling and cross-fitting setup above and Assumptions 2, D.1, and D.3, for each fixed deterministic $\omega \in \Omega_n$ and $i \in J_{kr}$,*

$$\mathbb{E}\{\widehat{T}_{j,i}^{(-k,-r)}(\psi_\omega) \mid L_i, \mathcal{F}_{-k,-r,j}\} = \mathbb{E}\{\psi_\omega(W_i) \mid L_i\} + r_{n,kr,\omega,j}(L_i), \tag{88}$$

where

$$r_{n,kr,\omega,j}(L) := \mathbb{E}\{\rho_{n,kr,\omega,j}(L^+) \mid L, \mathcal{F}_{-k,-r,j}\}.$$

Consequently, $\max_{k,r} \sup_{\omega \in \Omega_n} \|r_{n,kr,\omega,j}\|_{L_2(P)} \leq R_{n,j}$*, which is $o_p(1)$ under Assumption D.2.*

Proof. For fixed ω, condition on the inner training information. MAR gives

$$\mathbb{E}\{\delta_W \psi_\omega(W) \mid L^+, \mathcal{F}_{-k,-r,j}\} = e(L^+)m_\omega(L^+).$$

The three conditional means of the estimated transformations given L^+ are, respectively,

$$em_\omega + (1-e)\widehat{m}_\omega, \qquad (e/\widehat{e})m_\omega, \qquad \widehat{m}_\omega + (e/\widehat{e})(m_\omega - \widehat{m}_\omega).$$

Subtracting m_ω yields exactly Equation 87, including its signs. Iterated expectation given L proves Equation 88; conditional Jensen gives the norm bound. The AIPW remainder is zero if either fitted nuisance equals its target. This also gives the specialization of Lemma B.1 and Equation 26 used in the original appendix. The identity is proved for fixed ω; uniform control, rather than conditional independence at a learned ω, permits later substitution of $\widehat{\omega}_j^{(-k)}$. □

D.3.2 CONDITIONAL EXPECTATIONS FOR BRIDGE LEARNING AND EVALUATION

This subsection defines the conditional expectations used in the consistency proof in Appendix D.4. We follow the two-stage error decomposition in Xu et al. (2021, Appendix C, Lemmas 3–4 and Theorems 2–3), while keeping the estimated-transformation bias and evaluation-sample error explicit.

For $\eta \in \Xi_n$ and a fixed treatment value a, define $h_{\eta,a}(W, X) := h_\eta(W, a, X)$. The conditional mean of this bridge function is

$$\begin{aligned} g_{\eta,a}(L) &:= \mathbb{E}\{h_\eta(W, a, X) \mid L\} \\ &= \mathbb{E}\{T_j(h_{\eta,a}) \mid L\} \\ &= u^\top\{\psi_A(a) \otimes \phi_X(X) \otimes \mathbb{E}[\psi_\omega(W) \mid L]\}. \end{aligned} \tag{89}$$

The second equality follows from MAR recovery, and the third uses that X is a component of L. A candidate h_η need not be a valid bridge. Fix a jointly measurable valid bridge $h_0 \in L_2(P_{W,A,X})$ that is integrable under $P_{W,X}$ at the treatment values considered. Set $h_{0,a}(W, X) := h_0(W, a, X)$ and $g_{0,a}(L) := \mathbb{E}\{h_0(W, a, X) \mid L\}$. At those values, the proximal identification conditions give

$$\mathbb{E}(Y \mid L) = \mathbb{E}\{T_j(h_0) \mid L\}, \qquad \beta(a) = \mathbb{E}\big[\mathbb{E}\{T_j(h_{0,a}) \mid L\}\big] = \mathbb{E}\{g_{0,a}(L)\}.$$

Here $T_j(h_\eta)$ and $T_j(h_0)$ use the factual treatment A in the bridge argument, whereas $T_j(h_{\eta,a})$ and $T_j(h_{0,a})$ use the fixed evaluation value a. In both cases $L = (A, Z, X)$ retains the factual A.

The estimated scalar transformation is induced by the same nuisance regression used for the vector response:

$$\begin{aligned} \widehat{T}_j^{(-k,-r)}(h_{\eta,a}) &:= u^\top\{\psi_A(a) \otimes \phi_X(X) \otimes \widehat{T}_j^{(-k,-r)}(\psi_\omega)\}, \\ \widehat{m}_{h_{\eta,a}}^{(-k,-r)}(L^+) &:= u^\top\{\psi_A(a) \otimes \phi_X(X) \otimes \widehat{m}_\omega^{(-k,-r)}(L^+)\}. \end{aligned}$$

Thus the scalar correction agrees with the algorithm's tensor construction. It is not a separately fitted scalar nuisance regression.

We use the fitted conditional-expectation notation to describe the actual Stage 1 regression and the induced scalar prediction:

$$\begin{aligned} \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\psi_\omega) \mid L] &:= \widehat{V}_j^{(-k)}(\omega)\Phi_{\widehat{\vartheta}_j^{(-k)}(\omega)}(L) = \widehat{\mu}_{\omega,j}^{(-k)}(L), \\ \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(h_{\eta,a}) \mid L] &:= u^\top\{\psi_A(a) \otimes \phi_X(X) \otimes \widehat{\mu}_{\omega,j}^{(-k)}(L)\}. \end{aligned} \tag{90}$$

Inside $\widehat{\mathbb{E}}^{(-k)}$, the argument $\widehat{T}_j$ specifies the inner-fold response family in Equation 81: each training observation uses the nuisance fit from its complementary inner sample. The displayed expression is the tensor-induced regression estimate, not a literal empirical conditional expectation or an additional scalar regression. Its definition retains the candidate representation η; no representation-independent linear regression operator is assumed. Write $\widehat{h}_j^{(-k)} = h_{\widehat{\eta}_j^{(-k)}}$ and $\widehat{h}_{j,a}^{(-k)}(W, X) = \widehat{h}_j^{(-k)}(W, a, X)$.

The population regression target is written directly as a weighted sum of the inner-fold conditional means:

$$\begin{aligned} &\sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega) \mid L] \\ &\quad = \mathbb{E}[\psi_\omega(W) \mid L] + \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|} r_{n,kr,\omega,j}(L), \\ &\sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(h_{\eta,a}) \mid L] \\ &\quad = u^\top\left\{\psi_A(a) \otimes \phi_X(X) \otimes \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega) \mid L]\right\}. \end{aligned} \tag{91}$$

Every population expectation involving fitted functions in this and subsequent displays is evaluated on a fresh observation with all training data and fitted functions held fixed; equivalently, it also

conditions on $\mathcal{F}_{-k,j}$. This convention applies to fitted bridges as well as fitted nuisances. The weighted sum is the population target of the pooled Stage 1 loss. It does not replace the fold-specific responses used on the observed training sample. The proof compares these population conditional means with the fitted conditional expectations in Equation 90, and then with the oracle means in Equation 89.

D.3.3 TWO-STAGE LEARNING CONDITIONS

For $i \in J_{kr}$, let

$$\ell_{j,kr,i}(\omega, V, \vartheta) := \|\widehat{T}_{j,i}^{(-k,-r)}(\psi_\omega) - V\Phi_\vartheta(L_i)\|_2^2.$$

The notation $\ell_{j,kr}(\omega, V, \vartheta; O)$ denotes the same loss evaluated on a fresh observation using the fitted inner-fold nuisances.

Assumption D.4 (Uniform empirical risk control for Stage 1)**.** *The following measurable deviation converges to zero in probability:*

$$\Delta_{1,n,j} := \max_{k,r} \sup_{\omega,V,\vartheta} \left| \frac{1}{|J_{kr}|} \sum_{i \in J_{kr}} \ell_{j,kr,i}(\omega, V, \vartheta) - P\ell_{j,kr}(\omega, V, \vartheta; O) \right| = o_p(1),$$

where the supremum is over $\Omega_n \times \mathcal{V}_n \times \Theta_{1,n}$. For fixed candidates, the population term equals the conditional expectation given $\mathcal{F}_{-k,-r,j}$.

The compactness, continuity, and integrable-envelope requirements from the earlier formulation may be retained for each n. They are useful regularity conditions for fixed classes (Newey & McFadden, 1994; van der Vaart & Wellner, 1996), but do not alone prove a uniform law for classes growing with n and depending on fitted nuisances. The displayed uniform convergence is therefore a separate assumption; a sufficient complexity bound is given below.

Assumption D.5 (Uniform two-stage approximation and fitting)**.** *There is a constant $B < \infty$ such that, for all n,*

$$\operatorname*{ess\,sup}_{x \sim P_X} \sup_{\eta \in \Xi_n} \sup_{a \in \mathcal{A}} \sup_{\|v\|_2 \le 1} \left| u^\top \{\psi_A(a) \otimes \phi_X(x) \otimes v\} \right| \le B.$$

Let $U_n = \sup_{\eta \in \Xi_n} \|u\|_2 < \infty$ and $D_{V,n} = \sup_{V \in \mathcal{V}_n} \|V\|_F < \infty$. The approximation errors are

$$\kappa_{1,n} := \sup_{\omega \in \Omega_n} \inf_{(V,\vartheta) \in \mathcal{V}_n \times \Theta_{1,n}} \|\mathbb{E}\{\psi_\omega(W) \mid L\} - V\Phi_\vartheta\|_{L_2(P)}^2,$$

$$\kappa_{2,n} := \inf_{\eta \in \Xi_n} \|\mathbb{E}(Y \mid L) - \mathbb{E}\{T_j(h_\eta) \mid L\}\|_{L_2(P)},$$

and satisfy $\kappa_{1,n}, \kappa_{2,n} \to 0$. There are nonnegative $\xi_{1,n,j}, \xi_{2,n,j} = o_p(1)$ such that, with probability tending to one, the Stage 1 fits in Equation 82 are within $\xi_{1,n,j}$ of their penalized empirical infima, uniformly over k and $\omega \in \Omega_n$, and the Stage 2 fits in Equation 84 are within $\xi_{2,n,j}$ of their penalized empirical infima, uniformly over k.

The Stage 2 risk deviation satisfies

$$\Delta_{2,n,j} := \max_{k} \sup_{\eta \in \Xi_n} \left| (P_{n,-k} - P) \left[Y - \widehat{\mathbb{E}}^{(-k)}\{\widehat{T}_j(h_\eta) \mid L\} \right]^2 \right| = o_p(1). \tag{92}$$

All displayed risks are finite, and $\lambda_1 D_{V,n}^2 \to 0$ and $\lambda_2 U_n^2 \to 0$.

The supremum in Equation 92 includes the dependence of Stage 1 on the candidate proxy map. Because Stage 1 and Stage 2 reuse the outer training sample, this uniform risk condition is an explicit empirical-process requirement, not a consequence of conditioning on that same training sample. It can, for example, be obtained from a uniform law for a deterministic class containing all candidate squared losses. Similarly, the uniform Stage 1 optimization and approximation conditions permit a data-adaptive $\widehat{\omega}_j^{(-k)}$. A fixed number of alternating gradient updates does not by itself establish the required optimization tolerances.

For a sample $S = (o_1, \ldots, o_N)$, define

$$\widehat{\mathfrak{R}}_S(\mathcal{H}) := \mathbb{E}_\sigma \left[\sup_{f \in \mathcal{H}} \frac{1}{N} \sum_{i=1}^N \sigma_i f(o_i) \right],$$

where the σ_i are independent uniform signs. Conditional on the inner training information, use the loss class

$$\mathcal{H}^{kr}_{1,n,j} := \{o \mapsto \ell_{j,kr}(\omega, V, \vartheta; o) : (\omega, V, \vartheta) \in \Omega_n \times \mathcal{V}_n \times \Theta_{1,n}\}.$$

For Stage 2, use the deterministic enlarged class

$$\mathcal{H}_{2,n} := \left\{(y,l) \mapsto \left[y - u^\top\{\psi_A(A) \otimes \phi_X(X) \otimes \mu(l)\}\right]^2 : \eta \in \Xi_n,\ \mu \in \mathcal{G}_n\right\}.$$

It contains every fitted Stage 2 loss, including the dependence of the Stage 1 fit on ω. Suppose these two loss classes are bounded in $[0, M_{1,n}]$ and $[0, M_{2,n}]$, respectively, with deterministic bounds, almost surely. With $S_{kr} = (O_i : i \in J_{kr})$, $S_{-k} = (O_i : i \in I_{-k})$, and $c_\delta = \log\{4K(R+1)/\delta\}$, the two one-sided empirical Rademacher bounds and a union bound give, with probability at least $1 - \delta$,

$$\Delta_{1,n,j} \le \max_{k,r}\left\{2\widehat{\mathfrak{R}}_{S_{kr}}(\mathcal{H}^{kr}_{1,n,j}) + 3M_{1,n}\sqrt{\frac{c_\delta}{2|J_{kr}|}}\right\},$$
$$\Delta_{2,n,j} \le \max_{k}\left\{2\widehat{\mathfrak{R}}_{S_{-k}}(\mathcal{H}_{2,n}) + 3M_{2,n}\sqrt{\frac{c_\delta}{2|I_{-k}|}}\right\}.$$

For Stage 1 the bound is first conditional on each inner training sample, then integrated over that sample. A union bound does not require different folds to be independent. For Stage 2 the uniform law is applied to the deterministic enlarged class, so reusing the outer training observations in both stages is permitted. If the displayed upper bounds vanish in probability for a sequence $\delta = \delta_n \downarrow 0$, the required empirical risk deviations vanish. Bounded losses are a sufficient condition for this paragraph, not an additional requirement of the moment-based theorem below. In either formulation, control of class complexity and optimization error must be established or assumed; it is not guaranteed merely by choosing neural networks.

Here $\kappa_{1,n}$ is the squared approximation error for the conditional mean. For any candidate vector regression μ,

$$\mathbb{E}\|\psi_\omega(W) - \mu(L)\|_2^2 = \mathbb{E}\|\psi_\omega(W) - \mathbb{E}\{\psi_\omega(W) \mid L\}\|_2^2 + \|\mathbb{E}\{\psi_\omega(W) \mid L\} - \mu(L)\|^2_{L_2(P_L)}.$$

The first term is irreducible feature noise. Only the approximation error in the second term is required to vanish. The Stage 2 argument likewise subtracts the irreducible outcome noise. The proof below writes the conditional-expectation objects and the sums of error terms explicitly.

Let $\mathcal{H}_0$ denote the set of jointly measurable functions $h \in L_2(P_{W,A,X})$ satisfying the factual conditional-moment restriction

$$\mathbb{E}\{h(W,A,X) - Y \mid L\} = 0 \quad \text{a.s.},$$

as well as the outcome-bridge equation and integrability requirements in Assumption 1(iii) at the treatment values specified below. When treatment is continuous, also impose the pointwise conditions used for identification in Appendix A.1. The treatment-specific requirements apply at every value covered by the desired conclusion: the specified a for pointwise consistency, or P_A-almost every a for consistency of the causal curve. The set $\mathcal{H}_0$ is fixed after these treatment values are specified and is nonempty; the valid bridge h_0 introduced above is one of its elements. By proximal identification, every $h \in \mathcal{H}_0$ satisfies

$$\mathbb{E}\{h(W,a,X)\} = \beta(a)$$

at the treatment values covered by the conclusion. The set $\mathcal{H}_0$ is a population comparison set; its elements need not belong to the candidate class $\{h_\eta : \eta \in \Xi_n\}$.

Define the distance to this set by

$$d_2(f, \mathcal{H}_0) := \inf_{h \in \mathcal{H}_0} \|f - h\|_{L_2(P_{W,A,X})}, \qquad f \in L_2(P_{W,A,X}). \tag{93}$$

The infimum need not be attained, and no nearest bridge is selected. If the complete-data law varies with n, the set and the distance are understood under that same law, with the additional index suppressed.

Assumption D.6 (Restricted inverse stability)**.** *There are deterministic $1 \le \tau_n < \infty$ such that, for every $\eta \in \Xi_n$,*

$$d_2(h_\eta, \mathcal{H}_0) \le \tau_n \left\| \mathbb{E}[T_j(h_\eta) \mid L] - \mathbb{E}(Y \mid L) \right\|_{L_2(P_L)}.$$

The following rate condition holds:

$$\tau_n \left\{ \begin{array}{l} \kappa_{2,n} + 2B\sqrt{\kappa_{1,n}} + 4BR_{n,j} \\ \quad + 2B\sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2} + \sqrt{2\Delta_{2,n,j} + \xi_{2,n,j} + \lambda_2 U_n^2} \end{array} \right\} = o_p(1). \tag{94}$$

This condition controls the distance of each candidate bridge to the set of valid bridges. It permits multiple valid bridges and does not require the fitted bridge to approach the particular representative h_0. In particular, a candidate that belongs to $\mathcal{H}_0$ has both zero set distance and zero conditional-moment error. The constant τ_n may grow with the candidate class, so Equation 94 is retained to ensure that the set distance tends to zero. Neither this stability inequality nor its vanishing-error requirement follows from proximal identification alone.

Assumption D.7 (Transfer from the valid-bridge set to the causal mean)**.** *Use the clause corresponding to the desired conclusion.*

(i) For the $L_2(P_A)$ causal curve, assume $P_{W,X} \ll P_{W,X|A=a}$ for P_A-almost every a, and

$$r_a := \frac{dP_{W,X}}{dP_{W,X|A=a}}, \qquad C_{\mathrm{av}} := \sup_n \operatorname*{ess\,sup}_{a \sim P_A} \|r_a\|_{L_2(P_{W,X|A=a})} < \infty.$$

(ii) At a specified fixed treatment a, proximal identification holds for every $h \in \mathcal{H}_0$, and there are deterministic $0 \le C_{a,n} < \infty$ such that

$$|\mathbb{E}\{h_\eta(W, a, X)\} - \beta(a)| \le C_{a,n}\, d_2(h_\eta, \mathcal{H}_0), \qquad \eta \in \Xi_n,$$

and

$$C_{a,n}\tau_n \left\{ \begin{array}{l} \kappa_{2,n} + 2B\sqrt{\kappa_{1,n}} + 4BR_{n,j} \\ \quad + 2B\sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2} + \sqrt{2\Delta_{2,n,j} + \xi_{2,n,j} + \lambda_2 U_n^2} \end{array} \right\} = o_p(1).$$

A sufficient condition for clause (ii)'s norm inequality is that the same $C_{a,n}$ satisfies

$$|\mathbb{E}\{h_\eta(W, a, X) - h(W, a, X)\}| \le C_{a,n} \|h_\eta - h\|_{L_2(P_{W,A,X})}$$

for every $\eta \in \Xi_n$ and every $h \in \mathcal{H}_0$. Every comparison bridge has mean $\beta(a)$, so taking the infimum over $h \in \mathcal{H}_0$ gives the stated inequality. A bound involving only the original fixed bridge h_0 would not justify this step. For discrete treatment with $P(A = a) > 0$, absolute continuity and $\|r_a\|_{L_2(P_{W,X|A=a})} < \infty$ give this sufficient condition with $C_{a,n} = \|r_a\|_{L_2(P_{W,X|A=a})} / \sqrt{P(A = a)}$. Clause (i) similarly yields the causal-curve bound by applying the change-of-measure inequality to each $h \in \mathcal{H}_0$, integrating over a, and then taking the infimum over h. For continuous treatment, clause (i) does not imply clause (ii). The fixed-treatment transfer inequality and the displayed vanishing-error condition remain separate requirements; replacing a fixed bridge by a set does not establish them automatically. The pointwise identification conditions in Appendix A.1 are also maintained.

D.4 Consistency of ProxMAR-DFPV

Theorem D.1 (Consistency from conditional-expectation error bounds)**.** *Fix $j \in \mathcal{J}$. Use the sampling and cross-fitting setup in Appendix D.2. Suppose Assumptions 1, 2, D.1, D.2, D.3, D.4, and D.5 hold. With probability tending to one, the fitted Stage 1 conditional expectation satisfies*

$$\begin{aligned} \max_k \sup_{\omega \in \Omega_n} \left\| \begin{array}{l} \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\psi_\omega) \mid L] \\ \quad - \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|} \mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega) \mid L] \end{array} \right\|_{L_2(P_L)} \\ \le \sqrt{\kappa_{1,n}} + R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2}, \end{aligned} \tag{95}$$

and its error relative to the oracle conditional mean satisfies

$$\begin{aligned}\max_k \sup_{\omega\in\Omega_n} \left\|\widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\psi_\omega)\mid L]-\mathbb{E}[T_j(\psi_\omega)\mid L]\right\|_{L_2(P_L)} \\ \le \sqrt{\kappa_{1,n}}+2R_{n,j}+\sqrt{2\Delta_{1,n,j}+\xi_{1,n,j}+\lambda_1 D_{V,n}^2}.\end{aligned} \tag{96}$$

The fitted bridge has the conditional-moment error bound

$$\begin{aligned}\max_k \left\|\mathbb{E}(Y\mid L)-\mathbb{E}[T_j(\widehat{h}_j^{(-k)})\mid L]\right\|_{L_2(P_L)} \\ \le \kappa_{2,n}+2B\sqrt{\kappa_{1,n}}+4BR_{n,j} \\ +2B\sqrt{2\Delta_{1,n,j}+\xi_{1,n,j}+\lambda_1 D_{V,n}^2}+\sqrt{2\Delta_{2,n,j}+\xi_{2,n,j}+\lambda_2 U_n^2}.\end{aligned} \tag{97}$$

Under Assumption D.6, the fitted bridges satisfy

$$\begin{aligned}&\max_k d_2(\widehat{h}_j^{(-k)},\mathcal{H}_0) \\ &\le \tau_n\left\{\begin{array}{l}\kappa_{2,n}+2B\sqrt{\kappa_{1,n}}+4BR_{n,j} \\ \quad +2B\sqrt{2\Delta_{1,n,j}+\xi_{1,n,j}+\lambda_1 D_{V,n}^2}+\sqrt{2\Delta_{2,n,j}+\xi_{2,n,j}+\lambda_2 U_n^2}\end{array}\right\}=o_p(1).\end{aligned} \tag{98}$$

Retain Assumption D.6. Under clause (i) of Assumption D.7, with the valid-bridge set defined for P_A-almost every treatment value,

$$\|\widehat{\beta}_{\mathrm{DFPV},j}-\beta\|_{L_2(P_A)} \xrightarrow{p} 0. \tag{99}$$

Under clause (ii), at the specified fixed treatment value a,

$$\widehat{\beta}_{\mathrm{DFPV},j}(a) \xrightarrow{p} \beta(a). \tag{100}$$

For continuous treatment, the corresponding pointwise identification conditions in Appendix A.1 are maintained. These conclusions require convergence to the valid-bridge set, not to a preselected element of that set.

Proof. All population expectations below use fresh observations with the fitted functions held fixed, as specified after Equation 91. We first bound the regression errors, then apply the three-term conditional-expectation decomposition, and finally average over an outer evaluation sample.

Stage 1: estimating the conditional mean of the corrected response. Fix an outer fold k and a deterministic candidate ω. By conditional variance decomposition, for each inner fold the population risk differs from

$$\left\|\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega)\mid L]-V\Phi_\vartheta(L)\right\|_{L_2(P_L)}^2$$

by a term independent of (V,ϑ). After averaging the risks over the inner folds, their parameter-dependent part is

$$\left\|\sum_{r=1}^R \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega)\mid L]-V\Phi_\vartheta(L)\right\|_{L_2(P_L)}^2. \tag{101}$$

The omitted additive term includes both conditional response noise and the dispersion of the inner-fold conditional means. It is constant in (V,ϑ) for this fixed ω, and therefore cancels in the following comparison. Uniform empirical risk control, approximate penalized minimization, and nonnegativ-

ity of the penalty give

$$\begin{aligned}
&\left\| \begin{array}{l} \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\psi_\omega) \mid L] \\ \quad - \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|} \mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega) \mid L] \end{array} \right\|_{L_2(P_L)}^2 \\
&\leq \inf_{V,\vartheta} \left\| \begin{array}{l} V\Phi_\vartheta(L) \\ \quad - \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|} \mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega) \mid L] \end{array} \right\|_{L_2(P_L)}^2 + 2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2 \\
&\leq (\sqrt{\kappa_{1,n}} + R_{n,j})^2 + 2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2.
\end{aligned}$$

The last inequality follows from the definition of $\kappa_{1,n}$ and Proposition D.1, which yields

$$\left\| \begin{array}{l} \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|} \mathbb{E}[\widehat{T}_j^{(-k,-r)}(\psi_\omega) \mid L] \\ \quad - \mathbb{E}[T_j(\psi_\omega) \mid L] \end{array} \right\|_{L_2(P_L)} \leq R_{n,j}.$$

Taking square roots proves Equation 95. Adding this last bias bound proves Equation 96. Both statements are uniform in ω, so they also hold at $\widehat{\omega}_j^{(-k)}$.

The tensor multiplier bound in Assumption D.5 transfers these vector errors to the scalar bridge expressions. Uniformly over η, k, and fixed a,

$$\begin{aligned}
&\left\| \begin{array}{l} \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(h_{\eta,a}) \mid L] \\ \quad - \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|} \mathbb{E}[\widehat{T}_j^{(-k,-r)}(h_{\eta,a}) \mid L] \end{array} \right\|_{L_2(P_L)} \\
&\leq B\{\sqrt{\kappa_{1,n}} + R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2}\},
\end{aligned} \tag{102}$$

and

$$\left\| \begin{array}{l} \sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|} \mathbb{E}[\widehat{T}_j^{(-k,-r)}(h_{\eta,a}) \mid L] \\ \quad - \mathbb{E}[T_j(h_{\eta,a}) \mid L] \end{array} \right\|_{L_2(P_L)} \leq BR_{n,j}. \tag{103}$$

The same inequalities hold with the fixed bridge argument a replaced by the factual treatment A. This follows from the uniform tensor multiplier bound, not from a general extrapolation of a regression fitted only at factual treatment values.

Stage 2: controlling the conditional outcome residual. For each realization of the fits and each candidate η, variance decomposition gives

$$\begin{aligned}
&\mathbb{E}\left[\{Y - \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(h_\eta) \mid L]\}^2\right] \\
&\quad = \mathbb{E}\left[\{Y - \mathbb{E}(Y \mid L)\}^2\right] + \left\|\mathbb{E}(Y \mid L) - \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(h_\eta) \mid L]\right\|_{L_2(P_L)}^2.
\end{aligned}$$

The first term is independent of η. The Stage 2 empirical risk bound and approximate penalized minimization therefore imply

$$\begin{aligned}
&\left\|\mathbb{E}(Y \mid L) - \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)}^2 \\
&\leq \inf_{\eta \in \Xi_n} \left\|\mathbb{E}(Y \mid L) - \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(h_\eta) \mid L]\right\|_{L_2(P_L)}^2 + 2\Delta_{2,n,j} + \xi_{2,n,j} + \lambda_2 U_n^2 \\
&\leq \left[\kappa_{2,n} + B\{\sqrt{\kappa_{1,n}} + 2R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2}\}\right]^2 + 2\Delta_{2,n,j} + \xi_{2,n,j} + \lambda_2 U_n^2.
\end{aligned}$$

The last comparison uses the oracle conditional-mean approximation defining $\kappa_{2,n}$ and the Stage 1 error bound. Taking square roots gives

$$\begin{aligned}&\left\|\mathbb{E}(Y \mid L) - \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)} \\ &\quad \le \kappa_{2,n} + B\{\sqrt{\kappa_{1,n}} + 2R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2}\} + \sqrt{2\Delta_{2,n,j} + \xi_{2,n,j} + \lambda_2 U_n^2}.\end{aligned} \tag{104}$$

Conditional-moment consistency: the three-term conditional-expectation decomposition. The bridge equation and oracle recovery give

$$\begin{aligned}&\left\|\mathbb{E}[\widehat{h}_j^{(-k)}(W, A, X) - h_0(W, A, X) \mid L]\right\|_{L_2(P_L)} \\ &\qquad = \left\|\mathbb{E}(Y \mid L) - \mathbb{E}[T_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)}.\end{aligned}$$

Insert the fitted conditional expectation and the population conditional mean of the estimated transformations. The triangle inequality gives exactly

$$\begin{aligned}&\left\|\mathbb{E}(Y \mid L) - \mathbb{E}[T_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)} \\ &\quad \le \underbrace{\left\|\mathbb{E}(Y \mid L) - \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)}}_{\text{(1) Stage 2 conditional-mean residual}} \\ &\qquad + \underbrace{\left\|\sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\widehat{h}_j^{(-k)}) \mid L] - \widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)}}_{\text{(2) Regression error for the estimated transformations}} \\ &\qquad + \underbrace{\left\|\sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\widehat{h}_j^{(-k)}) \mid L] - \mathbb{E}[T_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)}}_{\text{(3) Conditional bias of the estimated transformations}}.\end{aligned} \tag{105}$$

Term (1) is bounded by Equation 104. Term (2) is bounded by Equation 102, evaluated at the fitted bridge and factual A. Term (3) is bounded by Equation 103 at the same arguments. Adding the three explicit bounds yields

$$\begin{aligned}&\left\|\mathbb{E}(Y \mid L) - \mathbb{E}[T_j(\widehat{h}_j^{(-k)}) \mid L]\right\|_{L_2(P_L)} \\ &\quad \le \kappa_{2,n} + 2B\sqrt{\kappa_{1,n}} + 4BR_{n,j} \\ &\qquad + 2B\sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2} + \sqrt{2\Delta_{2,n,j} + \xi_{2,n,j} + \lambda_2 U_n^2}.\end{aligned}$$

This proves Equation 97. By Assumptions D.2, D.4, and D.5, its right-hand side converges to zero in probability. Hence

$$\max_k \left\|\mathbb{E}[T_j(\widehat{h}_j^{(-k)}) \mid L] - \mathbb{E}(Y \mid L)\right\|_{L_2(P_L)} \xrightarrow{p} 0.$$

Assumption D.6 further gives

$$\max_k d_2(\widehat{h}_j^{(-k)}, \mathcal{H}_0) \le \tau_n \max_k \left\|\mathbb{E}[T_j(\widehat{h}_j^{(-k)}) \mid L] - \mathbb{E}(Y \mid L)\right\|_{L_2(P_L)}.$$

Combining this inequality with Equation 97 and the vanishing-error condition in Equation 94 proves Equation 98. No bridge attaining the infimum is needed. The next step transfers this set-distance convergence to the causal mean using Assumption D.7.

Causal-mean evaluation: averaging the conditional expectations. For a fixed treatment value a, Equations 85–86 implement

$$\widehat{\beta}_{k,j}(a) = P_{n,k}\Big\{\widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_{j,a}^{(-k)}) \mid L]\Big\}, \qquad \beta(a) = \mathbb{E}[\mathbb{E}\{T_j(h_{0,a}) \mid L\}].$$

Adding and subtracting the corresponding population averages gives the exact decomposition

$$\begin{aligned}
&\widehat{\beta}_{k,j}(a) - \beta(a) \\
&\quad = (P_{n,k} - P)\Big\{\widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_{j,a}^{(-k)}) \mid L]\Big\} \\
&\qquad + \mathbb{E}\left[\begin{array}{l}\widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_{j,a}^{(-k)}) \mid L] \\ \quad - \displaystyle\sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\widehat{h}_{j,a}^{(-k)}) \mid L]\end{array}\right] \\
&\qquad + \mathbb{E}\left[\begin{array}{l}\displaystyle\sum_{r=1}^{R} \frac{|J_{kr}|}{|I_{-k}|}\mathbb{E}[\widehat{T}_j^{(-k,-r)}(\widehat{h}_{j,a}^{(-k)}) \mid L] \\ \quad - \mathbb{E}[T_j(\widehat{h}_{j,a}^{(-k)}) \mid L]\end{array}\right] \\
&\qquad + \mathbb{E}\left[\begin{array}{l}\mathbb{E}[T_j(\widehat{h}_{j,a}^{(-k)}) \mid L] \\ \quad - \mathbb{E}[T_j(h_{0,a}) \mid L]\end{array}\right].
\end{aligned} \tag{106}$$

The four terms are, respectively, the outer-sample averaging error, the fitted conditional-expectation error, the conditional transformation bias, and the effect of bridge estimation on the target mean. The last term equals

$$\mathbb{E}\{\widehat{h}_j^{(-k)}(W, a, X) - h_0(W, a, X)\}$$

by oracle recovery and iterated expectation. Thus the conditional function and its marginal average are kept distinct.

By Cauchy–Schwarz and Equations 102–103, the sum of the absolute values of the second and third terms is at most

$$B\{\sqrt{\kappa_{1,n}} + 2R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2}\}.$$

This upper bound is $o_p(1)$ under the maintained nuisance recovery and two-stage learning conditions. Because the scalar regression bounds are uniform in the fixed treatment argument a, the same upper bound controls the sum of the $L_2(P_A)$ norms of these two error functions for the causal-curve conclusion.

For the first term, the outer evaluation sample is independent of the fit conditional on $\mathcal{F}_{-k,j}$. The moment condition, conditional Jensen inequality, and Equation 96 give

$$\begin{aligned}
&\max_k \left\|\widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\psi_{\widehat{\omega}_j^{(-k)}}) \mid L]\right\|_{L_2(P_L)} \\
&\quad \le \sup_{\omega\in\Omega_n} \|\psi_\omega(W)\|_{L_2(P_W)} + \sqrt{\kappa_{1,n}} + 2R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2} = O_p(1).
\end{aligned}$$

Consequently,

$$\begin{aligned}
&\mathbb{E}\left[\left\{(P_{n,k} - P)\left(\widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\widehat{h}_{j,a}^{(-k)}) \mid L]\right)\right\}^2 \middle| \mathcal{F}_{-k,j}\right] \\
&\qquad \le \frac{B^2}{|I_k|}\left\|\widehat{\mathbb{E}}^{(-k)}[\widehat{T}_j(\psi_{\widehat{\omega}_j^{(-k)}}) \mid L]\right\|_{L_2(P_L)}^2.
\end{aligned}$$

The displayed conditional variance converges to zero in probability because $\min_k |I_k| \to \infty$ and the squared norm is $O_p(1)$. Conditional Chebyshev inequality, followed by averaging the resulting probability bound truncated at one, therefore shows that the first term is $o_p(1)$. Since K is fixed, this conclusion holds simultaneously for all outer folds. This argument uses the bounded second moment of the entire feature vector, not only bounds for each coordinate, and does not require independence between different outer-fold fits.

For the causal-curve conclusion, integrate the displayed conditional second-moment bound over $a \sim P_A$. Its upper bound is independent of a and converges to zero in probability. Conditional Markov inequality applied to the squared $L_2(P_A)$ norm then gives an $o_p(1)$ outer-sample error in that norm as well.

For the fourth term, every $h \in \mathcal{H}_0$ has the same target mean as h_0. Consequently, at each treatment value under consideration,

$$\begin{aligned}\mathbb{E}\{\widehat{h}_j^{(-k)}(W,a,X) - h_0(W,a,X)\} &= \mathbb{E}\{\widehat{h}_j^{(-k)}(W,a,X)\} - \beta(a) \\ &= \mathbb{E}\{\widehat{h}_j^{(-k)}(W,a,X) - h(W,a,X)\}.\end{aligned}$$

Thus h_0 may remain in the exact decomposition without being the bridge to which the fitted function must converge.

For the causal-curve conclusion, fix any $h \in \mathcal{H}_0$. Clause (i) of Assumption D.7 and Cauchy–Schwarz give

$$\begin{aligned}&\int \left|\mathbb{E}\{\widehat{h}_j^{(-k)}(W,a,X)\} - \beta(a)\right|^2 dP_A(a) \\ &= \int \left|\mathbb{E}\Big[r_a(W,X)\{\widehat{h}_j^{(-k)}(W,a,X) - h(W,a,X)\} \mid A=a\Big]\right|^2 dP_A(a) \\ &\le C_{\mathrm{av}}^2 \int \mathbb{E}\Big[\{\widehat{h}_j^{(-k)}(W,a,X) - h(W,a,X)\}^2 \mid A=a\Big]\, dP_A(a) \\ &= C_{\mathrm{av}}^2 \|\widehat{h}_j^{(-k)} - h\|^2_{L_2(P_{W,A,X})}.\end{aligned}$$

The left-hand side does not depend on the comparison bridge h. Taking square roots and then the infimum over $h \in \mathcal{H}_0$ yields

$$\Big\|a \mapsto \mathbb{E}\{\widehat{h}_j^{(-k)}(W,a,X)\} - \beta(a)\Big\|_{L_2(P_A)} \le C_{\mathrm{av}}\, d_2(\widehat{h}_j^{(-k)}, \mathcal{H}_0).$$

The infimum is taken after the inequality for the integrated norm; it does not require a treatment-dependent choice of a closest bridge.

At a specified fixed treatment value a, clause (ii) instead gives

$$\left|\mathbb{E}\{\widehat{h}_j^{(-k)}(W,a,X)\} - \beta(a)\right| \le C_{a,n}\, d_2(\widehat{h}_j^{(-k)}, \mathcal{H}_0).$$

Combining Equation 98 with clause (ii)'s vanishing-error condition proves that this term is $o_p(1)$, uniformly over the fixed number of outer folds. Set-distance convergence alone would not suffice when $C_{a,n}$ increases with n.

Combining all four terms in Equation 106 gives, for the fixed-treatment conclusion,

$$\begin{aligned}\max_k |\widehat{\beta}_{k,j}(a) - \beta(a)| &\le C_{a,n} \max_k d_2(\widehat{h}_j^{(-k)}, \mathcal{H}_0) \\ &\quad + B\{\sqrt{\kappa_{1,n}} + 2R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2}\} + o_p(1) \\ &= o_p(1).\end{aligned}$$

For the causal-curve conclusion, the same argument yields

$$\begin{aligned}\max_k \|\widehat{\beta}_{k,j} - \beta\|_{L_2(P_A)} &\le C_{\mathrm{av}} \max_k d_2(\widehat{h}_j^{(-k)}, \mathcal{H}_0) \\ &\quad + B\{\sqrt{\kappa_{1,n}} + 2R_{n,j} + \sqrt{2\Delta_{1,n,j} + \xi_{1,n,j} + \lambda_1 D_{V,n}^2}\} + o_p(1) \\ &= o_p(1).\end{aligned}$$

The final estimator is a convex combination of the fold-specific estimators. Therefore,

$$\begin{aligned}|\widehat{\beta}_{\mathrm{DFPV},j}(a) - \beta(a)| &\le \max_k |\widehat{\beta}_{k,j}(a) - \beta(a)|, \\ \|\widehat{\beta}_{\mathrm{DFPV},j} - \beta\|_{L_2(P_A)} &\le \max_k \|\widehat{\beta}_{k,j} - \beta\|_{L_2(P_A)}.\end{aligned}$$

Applying the relevant bound proves Equations 100 and 99, under their respective clauses of Assumption D.7. □

E EXPERIMENTAL DETAILS AND ADDITIONAL RESULTS

E.1 DEMAND BENCHMARK

E.1.1 CAUSAL STRUCTURAL MODEL

The Demand design follows the structural model used in the original DFPV and NMMR studies (Xu et al., 2021; Kompa et al., 2022). There are no observed covariates, so $X = \varnothing$. The scalar U is an unobserved demand factor, A is treatment, W is the outcome-inducing proxy, and $Z = (Z_1, Z_2)$ contains two treatment-inducing proxies. Define

$$\psi(t) = 2\left\{\frac{(t-5)^4}{600} + \exp[-4(t-5)^2] + \frac{t}{10} - 2\right\}. \tag{107}$$

For each observation, generate

$$\begin{aligned} U_i &\sim \mathrm{Uniform}(0, 10), \\ Z_{1i} &= 2\sin(2\pi U_i/10) + \varepsilon_{Z1,i}, \\ Z_{2i} &= 2\cos(2\pi U_i/10) + \varepsilon_{Z2,i}, \\ A_i &= 35 + (Z_{1i} + 3)\psi(U_i) + Z_{2i} + \varepsilon_{A,i}, \\ W_i &= 7\psi(U_i) + 45 + \varepsilon_{W,i}, \\ Y_i &= \min\left\{\exp\left(\frac{W_i - A_i}{10}\right), 5\right\} A_i - 5\psi(U_i) + \varepsilon_{Y,i}. \end{aligned} \tag{108}$$

The five noise terms are mutually independent $N(0,1)$ variables and are independent of U_i. An outcome bridge is

$$h_D(w, a) = a\min\left\{\exp\left(\frac{w-a}{10}\right), 5\right\} - \frac{5}{7}(w - 45), \tag{109}$$

because $Y - h_D(W, A) = \varepsilon_Y + (5/7)\varepsilon_W$ has conditional mean zero given (A, Z). The causal curve is

$$\beta(a) = \mathbb{E}\left[a\min\left\{\exp\left(\frac{W-a}{10}\right), 5\right\} - 5\psi(U)\right]. \tag{110}$$

We evaluate ten equally spaced treatment values in $[10, 30]$. The common reference curve is calculated from 10^6 draws from Equation 108 using seed 42.

E.1.2 MISSING MECHANISM AND EVALUATION PROTOCOL

Only W is incomplete. For $V \in \{A, Z_1, Z_2, Y\}$, define the standardized variable $V_i^\circ = (V_i - \mu_V)/s_V$. In the order (A, Z_1, Z_2, Y), the fixed centers and scales are

$$\begin{aligned} (\mu_A, \mu_{Z_1}, \mu_{Z_2}, \mu_Y) &\approx (27.14, 0.001, -0.001, 43.74), \\ (s_A, s_{Z_1}, s_{Z_2}, s_Y) &\approx (6.63, 1.73, 1.73, 12.19). \end{aligned}$$

For target missing rate $r \in \{0.3, 0.5, 0.7\}$, set

$$\begin{aligned} \eta_{r,i} &= c_r + 1.6A_i^\circ + 0.8Z_{1i}^\circ - 0.8Z_{2i}^\circ + 1.2Y_i^\circ, \\ e_{r,i} &:= \Pr(\delta_{W,i} = 1 \mid A_i, Z_i, Y_i) = \frac{\exp(\eta_{r,i})}{1 + \exp(\eta_{r,i})}, \\ (c_{0.3}, c_{0.5}, c_{0.7}) &= (1.33, 0.25, -1.00). \end{aligned} \tag{111}$$

The untruncated logistic observation model is not uniformly bounded away from zero, so this benchmark also assesses performance beyond the uniform-positivity condition used in our consistency analysis. The centers, scales, coefficients, and intercepts are fixed before the reported replications. Independently draw $B_i \sim \mathrm{Uniform}(0, 1)$ and set $\delta_{W,i} = \mathbf{1}\{B_i < e_{r,i}\}$. The values $r \in \{0.3, 0.5, 0.7\}$ index the target missing rates, while the realized rate varies across replications.

The missing-rate sweep fixes $n = 5000$ and varies $r \in \{0.3, 0.5, 0.7\}$. The sample-size sweep fixes $r = 0.5$ and varies $n \in \{1000, 2500, 5000\}$. Each cell contains 100 replications. Within

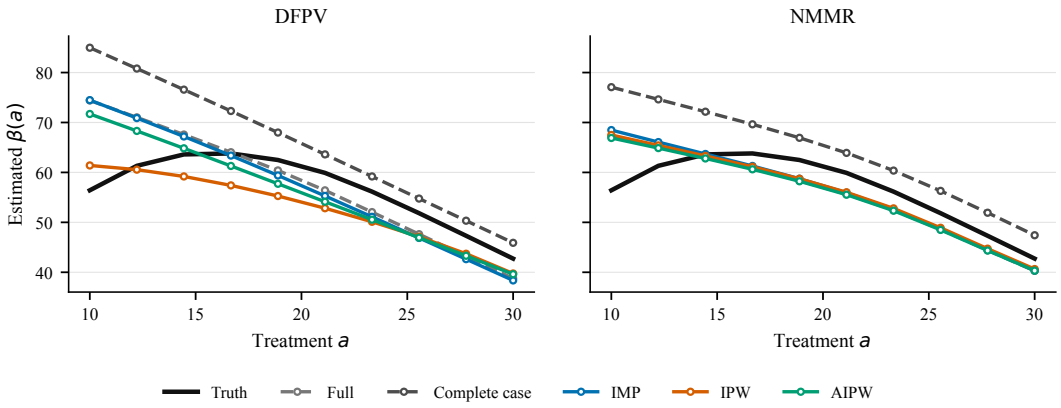


Figure 5: Demand response curves at $n = 5000$ and 50% target missingness. Lines show mean fitted curves over 100 replications.

a cell and replication, all methods receive the same complete data and missingness mask. Their outer folds and method-specific initialization seeds are also fixed. The sample-size cells are nested prefixes of a common complete-data draw and missingness-uniform draw within each replication. The two sweeps use separate prescribed seed sequences at $(n, r) = (5000, 0.5)$ and are summarized separately.

For replication b, method m, and version v, define

$$\mathrm{CMAE}_{b,m,v} = \frac{1}{10} \sum_{j=1}^{10} \left| \widehat{\beta}_{b,m,v}(a_j) - \beta(a_j) \right|. \tag{112}$$

The tables report the mean of the 100 replication-level CMAEs and its Monte Carlo standard error.

E.1.3 Implementation

DFPV uses ten outer folds and five inner folds; NMMR uses five outer and five inner folds. DFPV reserves 20% of each training fold for validation and selects the parameters with the lowest validation loss over at most 2,000 alternating epochs. Validation is evaluated every 20 epochs after epoch 200. Each epoch contains 20 Stage 1 updates, and the proxy network is pretrained for 50 epochs. The Stage 1 penalty, Stage 2 penalty, and weight decay are each 0.1.

NMMR represents the bridge by a three-layer width-80 network and performs 3,000 full-batch Adam updates with learning rate 0.003, moment-kernel length scale 1, and penalty 3×10^{-6}. For the conditional-mean regressions used by NMMR, the Gaussian bandwidth is selected from $\{0.03125, 0.0625, 0.125, 0.25, 0.5, 1, 2\}$ by leave-one-out squared prediction error over complete cases in the corresponding training sample.

The IPW and AIPW variants use an observation-probability network with two width-32 ReLU hidden layers, trained for 400 epochs with learning rate 0.001 and weight decay 10^{-4}. Fitted probabilities are truncated to $[0.001, 0.999]$ before inversion.

E.1.4 Performances with changing missing rate

Figure 5 compares the mean fitted curves at $n = 5000$ and 50% target missingness. DFPV-IPW and the three NMMR corrections closely track one another across the treatment grid. The complete-case curves depart more strongly from the reference curve for DFPV and NMMR.

Table 1 and Figure 6 show that every DFPV and NMMR correction reduces mean CMAE at 50% and 70% missingness. At 30%, DFPV-IPW and DFPV-AIPW also reduce mean CMAE.

Method	Version	30%	50%	70%
DFPV	Full data	5.403 (0.027)	5.403 (0.027)	5.403 (0.027)
	Complete case	5.655 (0.029)	9.088 (0.072)	16.353 (0.111)
	IMP	5.685 (0.020)	5.849 (0.016)	6.035 (0.015)
	IPW	**4.732 (0.030)**	**4.897 (0.052)**	**5.168 (0.058)**
	AIPW	5.540 (0.026)	5.449 (0.020)	6.293 (0.062)
NMMR	Full data	3.913 (0.029)	3.913 (0.029)	3.913 (0.029)
	Complete case	3.384 (0.054)	7.479 (0.097)	15.484 (0.186)
	IMP	3.918 (0.026)	3.958 (0.022)	3.812 (0.024)
	IPW	3.836 (0.026)	**3.716 (0.028)**	**3.496 (0.033)**
	AIPW	3.926 (0.031)	3.927 (0.031)	4.028 (0.088)

Table 1: Demand CMAE across target missing rates at $n = 5000$: mean (Monte Carlo SE) over 100 replications. The full-data fit is shared across missing rates.

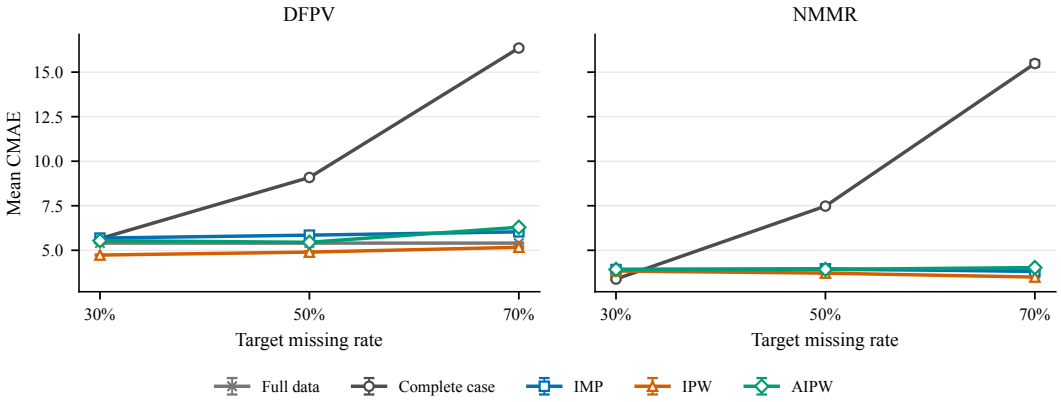


Figure 6: Demand CMAE across target missing rates at $n = 5000$. Points are means over 100 replications; error bars are Monte Carlo standard errors.

E.1.5 Performances with changing sample size

Every NMMR correction reduces mean CMAE at each sample size. DFPV-IPW reduces mean CMAE and attains the best MAR-aware DFPV result at $n = 2500$ and $n = 5000$.

Method	Version	$n = 1000$	$n = 2500$	$n = 5000$
DFPV	Full data	5.739 (0.023)	5.695 (0.022)	5.332 (0.027)
	Complete case	4.373 (0.093)	8.779 (0.090)	9.199 (0.061)
	IMP	5.783 (0.029)	5.881 (0.020)	5.823 (0.017)
	IPW	5.813 (0.084)	**4.963 (0.055)**	**4.771 (0.048)**
	AIPW	6.919 (0.087)	5.357 (0.023)	5.494 (0.022)
NMMR	Full data	3.865 (0.041)	3.917 (0.035)	3.985 (0.027)
	Complete case	8.656 (0.198)	7.951 (0.131)	7.515 (0.091)
	IMP	3.919 (0.042)	3.896 (0.028)	3.961 (0.026)
	IPW	**3.279 (0.073)**	**3.469 (0.042)**	**3.771 (0.030)**
	AIPW	4.325 (0.186)	3.914 (0.032)	3.989 (0.031)

Table 2: Demand CMAE across sample sizes at 50% target missingness: mean (Monte Carlo SE) over 100 replications.

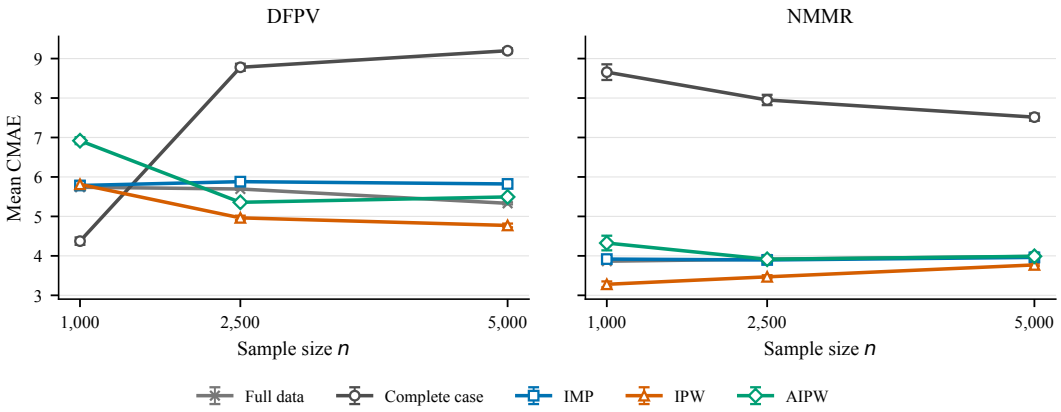


Figure 7: Demand CMAE across sample sizes at 50% target missingness. Points are means over 100 replications; error bars are Monte Carlo standard errors.

E.2 High-dimensional proxy design

E.2.1 Causal Structural Model

There are no observed covariates. Let $U_i, Z_i \in \mathbb{R}^2$ and $W_i \in \mathbb{R}^{p_W}$. We generate

$$\begin{aligned} U_i &\sim N(0, I_2), \\ Z_i &= U_i + \varepsilon_{Z,i}, & \varepsilon_{Z,i} &\sim N(0, 0.5^2 I_2), \\ W_i &= 2Q_{p_W} U_i + \varepsilon_{W,i}, & \varepsilon_{W,i} &\sim N(0, I_{p_W}), \\ A_i &= \gamma_Z^\top Z_i + \gamma_U^\top U_i + \varepsilon_{A,i}, & \varepsilon_{A,i} &\sim N(0, 0.7^2), \\ Y_i &= \tau(A_i) + \beta_U^\top U_i + \varepsilon_{Y,i}, & \varepsilon_{Y,i} &\sim N(0,1), \end{aligned} \tag{113}$$

where the noise terms are mutually independent,

$$\gamma_Z = (0.8, -0.5)^\top, \qquad \gamma_U = (0.7, -0.4)^\top, \qquad \beta_U = (1.0, -0.8)^\top,$$

and

$$\tau(a) = 2\sin(a/2) + 0.3a. \tag{114}$$

For each p_W, we generate a fixed Gaussian $p_W \times 2$ matrix and take the orthonormal factor Q_{p_W} from its thin QR decomposition. We fix the sign of each column by a deterministic rule. Hence

$$(2Q_{p_W})^\top (2Q_{p_W}) = 4I_2$$

for every proxy dimension. Thus the strength of the relationship between W and U is the same for every p_W, while the number of variables in W increases.

Let $c_W = Q_{p_W}\beta_U/2$. Since $c_W^\top(2Q_{p_W}) = \beta_U^\top$, an outcome bridge and the corresponding causal curve are

$$h_0(w, a) = \tau(a) + c_W^\top w, \qquad \beta(a) = \mathbb{E}\{h_0(W, a)\} = \tau(a). \tag{115}$$

We evaluate ten equally spaced treatment values in $[-3, 3]$.

E.2.2 Missing mechanism and evaluation protocol

Only W is incomplete. Write $L_i^+ = (A_i, Z_{1i}, Z_{2i}, Y_i)^\top$ and standardize its coordinates as $\widetilde{L}_{ij}^+ = (L_{ij}^+ - \mu_j)/s_j$, where

$$\mu = (0.002, -0.001, -0.003, 0.005)^\top, \qquad s = (1.944, 1.119, 1.118, 3.194)^\top.$$

For target missing rate $r \in \{0.1, 0.3, 0.5\}$, define

$$\eta_{r,i} = c_r + 1.6\widetilde{A}_i + 0.8\widetilde{Z}_{1i} - 0.8\widetilde{Z}_{2i} + 1.2\widetilde{Y}_i,$$
$$e_{r,i} = \min\left[0.95, \max\left\{0.05, \frac{\exp(\eta_{r,i})}{1+\exp(\eta_{r,i})}\right\}\right],$$
$$(c_{0.1}, c_{0.3}, c_{0.5}) = (6.420, 2.475, -0.001). \tag{116}$$

The centers, scales, and intercepts are fixed from an independent reference draw of size 10^6. Independently draw $B_i \sim \mathrm{Uniform}(0,1)$ and set $\delta_{W,i} = \mathbf{1}\{B_i < e_{r,i}\}$.

We vary one factor at a time. The missing-rate experiment fixes $(n, p_W) = (5000, 200)$ and varies $r \in \{0.1, 0.3, 0.5\}$. The proxy-dimension experiment fixes $(n, r) = (5000, 0.5)$ and varies $p_W \in \{100, 200, 300\}$. The sample-size experiment fixes $(p_W, r) = (200, 0.5)$ and varies $n \in \{1000, 5000, 10000\}$. Increasing p_W increases the input dimension of the bridge networks. The nuisance regressions always use the four variables in $L^+ = (A, Z_1, Z_2, Y)$ as inputs. Each of the seven distinct settings contains 50 replications. CMAE is defined as in Equation 112; the tables report its mean and Monte Carlo standard error.

E.2.3 Implementation

DFPV and NMMR compare full-data, complete-case, IMP, IPW, and AIPW fitting. For DFPV, the proxy feature satisfies $\psi_\omega(W) \in \mathbb{R}^8$. The W-encoder in NMMR also has eight output coordinates. Each estimator uses five outer folds and five inner folds.

DFPV uses a summed squared-error objective in Stage 1, runs for at most 2,000 alternating epochs, and selects parameters by validation loss. NMMR minimizes the V-statistic objective for 600 updates. For NMMR, the Gaussian regression bandwidth is selected from $\{0.03125, 0.0625, 0.125, 0.25, 0.5, 1, 2\}$ by leave-one-out squared prediction error on the corresponding training observations. The IPW and AIPW versions estimate observation probabilities with a neural classifier and replace fitted probabilities below 0.05 by 0.05 before inversion.

E.2.4 Performances with changing missing rate

Tables 3–5 report the seven settings. DFPV-IPW improves on complete-case DFPV throughout, and DFPV-IMP improves in six settings. Each NMMR correction improves on complete-case NMMR in every setting. NMMR-AIPW attains the lowest corrected CMAE at 10% missingness, NMMR-IPW at 30% and 50% missingness, NMMR-IMP at $p_W = 300$ and $n = 1000$, and NMMR-IPW at $n = 5000$ and $n = 10000$.

Method	Version	10%	30%	50%
DFPV	Full data	0.346 (0.009)	0.346 (0.009)	0.346 (0.009)
	Complete case	0.394 (0.010)	0.408 (0.005)	0.624 (0.007)
	Mean value imputation	0.484 (0.007)	0.410 (0.008)	0.439 (0.009)
	IMP	**0.334 (0.009)**	0.437 (0.008)	**0.359 (0.007)**
	IPW	0.380 (0.010)	**0.374 (0.008)**	0.392 (0.009)
	AIPW	0.341 (0.009)	0.513 (0.009)	0.615 (0.011)
NMMR	Full data	0.194 (0.006)	0.194 (0.006)	0.194 (0.006)
	Complete case	0.218 (0.006)	0.412 (0.004)	0.641 (0.004)
	Mean value imputation	0.318 (0.006)	0.592 (0.005)	0.617 (0.004)
	IMP	0.199 (0.007)	0.211 (0.006)	0.215 (0.007)
	IPW	0.199 (0.007)	**0.180 (0.005)**	**0.188 (0.005)**
	AIPW	**0.196 (0.007)**	0.197 (0.006)	0.219 (0.005)

Table 3: High-dimensional CMAE across target missing rates at $(n, p_W) = (5000, 200)$: mean (Monte Carlo SE) over 50 replications.

E.2.5 PERFORMANCES WITH CHANGING PROXY DIMENSIONS

Method	Version	$p_W = 100$	$p_W = 200$	$p_W = 300$
DFPV	Full data	0.331 (0.008)	0.346 (0.009)	0.351 (0.011)
	Complete case	0.617 (0.007)	0.624 (0.007)	0.620 (0.008)
	Mean value imputation	0.461 (0.009)	0.439 (0.009)	0.450 (0.011)
	IMP	**0.401 (0.009)**	**0.359 (0.007)**	**0.348 (0.010)**
	IPW	0.402 (0.012)	0.392 (0.009)	0.387 (0.011)
	AIPW	0.611 (0.009)	0.615 (0.011)	0.614 (0.013)
NMMR	Full data	0.357 (0.006)	0.194 (0.006)	0.099 (0.004)
	Complete case	0.667 (0.008)	0.641 (0.004)	0.635 (0.005)
	Mean value imputation	0.668 (0.004)	0.617 (0.004)	0.574 (0.005)
	IMP	0.392 (0.007)	0.215 (0.007)	**0.143 (0.004)**
	IPW	**0.258 (0.008)**	**0.188 (0.005)**	0.232 (0.007)
	AIPW	0.341 (0.007)	0.219 (0.005)	0.261 (0.007)

Table 4: High-dimensional CMAE across proxy dimensions at $(n, r) = (5000, 0.5)$: mean (Monte Carlo SE) over 50 replications.

E.2.6 PERFORMANCES WITH CHANGING SAMPLE SIZE

Method	Version	$n = 1000$	$n = 5000$	$n = 10000$
DFPV	Full data	0.549 (0.014)	0.346 (0.009)	0.263 (0.007)
	Complete case	0.671 (0.017)	0.624 (0.007)	0.651 (0.008)
	Mean value imputation	0.417 (0.014)	0.439 (0.009)	0.386 (0.009)
	IMP	0.520 (0.013)	**0.359 (0.007)**	0.324 (0.008)
	IPW	**0.516 (0.017)**	0.392 (0.009)	**0.307 (0.011)**
	AIPW	0.590 (0.016)	0.615 (0.011)	0.563 (0.009)
NMMR	Full data	0.150 (0.008)	0.194 (0.006)	0.193 (0.007)
	Complete case	0.600 (0.009)	0.641 (0.004)	0.640 (0.004)
	Mean value imputation	0.570 (0.007)	0.617 (0.004)	0.635 (0.004)
	IMP	**0.146 (0.007)**	0.215 (0.007)	0.219 (0.007)
	IPW	0.342 (0.012)	**0.188 (0.005)**	**0.152 (0.003)**
	AIPW	0.314 (0.011)	0.219 (0.005)	0.194 (0.004)

Table 5: High-dimensional CMAE across sample sizes at $(p_W, r) = (200, 0.5)$: mean (Monte Carlo SE) over 50 replications.

Within each experimental design, neural-network architectures and training hyperparameters are held fixed across sample sizes, without sample-size-specific tuning; this may partly explain the flat or non-monotonic estimation-error trends observed for some methods.

E.3 RHC

E.3.1 SEMI-SYNTHETIC RIGHT HEART CATHETERIZATION STUDY

We use 5,735 critically ill patients from the SUPPORT Right Heart Catheterization (RHC) study (Connors et al., 1996). The patient records, treatment, outcome, and proxy measurements are observed; the experiment imposes missingness on the outcome-proxy block. Table 6 defines the variables. The treatment A indicates receipt of right heart catheterization. The outcome Y is the number of days survived during the first 30 days, capped at 30. The treatment-inducing proxy Z contains the arterial oxygenation ratio and arterial carbon-dioxide pressure, and the outcome-inducing proxy W contains arterial pH and hematocrit. We set $X = \varnothing$. The proxy allocation follows earlier proximal analyses of this cohort (Tchetgen et al., 2020; Cui et al., 2024).

Role	Symbol	RHC column(s)	Definition
Treatment	A	`swang1`	receipt of right heart catheterization
Outcome	Y	`t3d30`	days survived during the first 30 days, capped at 30
Treatment proxy	Z	`pafi1`, `paco21`	PaO_2/FiO_2 ratio and $PaCO_2$
Outcome proxy	W	`ph1`, `hema1`	arterial pH and hematocrit
Covariates	X	none	no additional observed covariates

Table 6: Variables in the RHC study. The four proxy measurements are from the initial 24-hour physiological panel.

Only the two components of W are deleted. Let $L_i^+ = (A_i, Z_{1i}, Z_{2i}, Y_i)^\top$ and let $\widetilde{L}_i^+$ denote its coordinatewise standardized form. For target missing rate $r \in \{0.3, 0.5\}$, we set

$$e_{r,i} = \min\left[0.95, \max\left\{0.05, \frac{\exp\{\alpha_0(r) + \widetilde{L}_i^{+\top}\alpha\}}{1 + \exp\{\alpha_0(r) + \widetilde{L}_i^{+\top}\alpha\}}\right\}\right], \qquad \alpha = (0, 0.4, -0.4, 0.8)^\top, \tag{117}$$

where $\alpha_0(r)$ is chosen so that the cohort average of $e_{r,i}$ equals $1-r$. A common uniform draw within each replication generates the observation indicators at both rates, giving nested observed samples. Each rate has 50 replications. All methods use the same data split and observation indicators within a replication and use five outer folds and five inner folds.

No counterfactual reference is available for RHC. We therefore measure recovery of each method's full-data estimate of the causal mean. For method m, version $v \in \{\mathrm{CC}, \mathrm{MAR}\}$, and replication b, define

$$\mathrm{CMAE}_{b,m,v}(r) = \frac{1}{2}\sum_{a\in\{0,1\}} \left|\widehat{\beta}_{b,m,v}(a; r) - \widehat{\beta}_{b,m,\mathrm{full}}(a)\right|. \tag{118}$$

Table 7 reports the mean recovery error and its Monte Carlo standard error. The final column reports the mean of the complete-case error minus the MAR-aware error and its 95% t interval.

Method	Missing	CC error	MAR-aware error	CC − MAR-aware [95% interval]
DFPV (IMP)	30%	2.435 (0.034)	**0.084 (0.010)**	2.351 [2.285, 2.416]
DFPV (IMP)	50%	3.465 (0.042)	**0.109 (0.012)**	3.355 [3.272, 3.439]
DFPV (IPW)	30%	2.435 (0.034)	0.147 (0.013)	2.288 [2.210, 2.365]
DFPV (IPW)	50%	3.465 (0.042)	0.212 (0.019)	3.252 [3.158, 3.347]
DFPV (AIPW)	30%	2.435 (0.034)	0.109 (0.012)	2.326 [2.254, 2.399]
DFPV (AIPW)	50%	3.465 (0.042)	0.160 (0.016)	3.305 [3.214, 3.396]
NMMR (IMP)	30%	2.273 (0.021)	**0.243 (0.018)**	2.030 [1.975, 2.085]
NMMR (IMP)	50%	3.276 (0.025)	**0.387 (0.020)**	2.890 [2.826, 2.954]
NMMR (IPW)	30%	2.273 (0.021)	0.249 (0.025)	2.024 [1.961, 2.087]
NMMR (IPW)	50%	3.276 (0.025)	0.491 (0.038)	2.785 [2.708, 2.862]
NMMR (AIPW)	30%	2.273 (0.021)	0.253 (0.018)	2.020 [1.960, 2.080]
NMMR (AIPW)	50%	3.276 (0.025)	0.465 (0.025)	2.811 [2.740, 2.881]

Table 7: RHC recovery of method-specific full-data estimates over 50 replications. Error entries are means with Monte Carlo standard errors in parentheses. Boldface marks the lowest MAR-aware error within each estimator family and missing rate.

Every MAR-aware version has lower recovery error than its complete-case fit in both missing-rate settings, and every 95% interval in the final column lies above zero. DFPV-IMP and NMMR-IMP attain the lowest recovery error within their respective families at each missing rate. Figure 8 shows the replication-level differences summarized by the table.

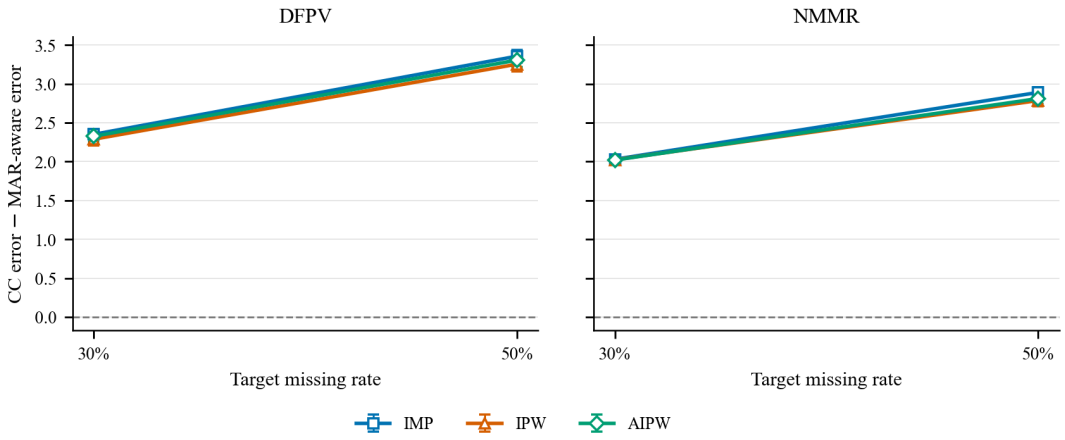


Figure 8: RHC recovery improvement across 50 replications. The vertical axis is complete-case minus MAR-aware error relative to each method's full-data fit. Hollow markers and bars show means and 95% t intervals.

DFPV uses 100 alternating epochs, and NMMR uses 3,000 updates. The imputation-based NMMR versions select Gaussian regression bandwidths from $\{0.03125, 0.0625, 0.125, 0.25, 0.5, 1, 2\}$ by leave-one-out squared prediction error within the corresponding training sample. Neural classifiers estimate the observation probabilities for IPW and AIPW.

E.4 DFPV NUISANCE-SPECIFICATION DIAGNOSTIC

We isolate recovery of the DFPV Stage 1 conditional feature target from proxy-representation learning. The diagnostic uses 50 replications of a high-dimensional proxy design with $n = 5000$, $p_W = 500$, and 50% missingness. We fix the scalar proxy feature at $\psi(W) = c_W^\top W$ and define

$$\mu(L) = \mathbb{E}\{\psi(W) \mid L\}, \qquad m_\omega(L^+) = \mathbb{E}\{\psi(W) \mid L^+\},$$

where $L = (A, Z)$ and $L^+ = (A, Z, Y)$. Correct specification uses the corresponding data-generating quantities. Under misspecification, the observation probability is replaced by the constant observed fraction and the feature regression is replaced by zero, its intercept-only value under the centered design. The four regimes use the same datasets, missingness indicators, and folds.

For each replication, signed pseudo-outcome bias is the sample average of $\phi_{\text{aipw}} - \mu(L)$, and Stage 1 target MAE is the sample average of $|\widehat{\mu}(L) - \mu(L)|$. Table 8 reports their across-replication means; the Monte Carlo standard error is the sample standard deviation of the 50 replication-level target MAEs divided by $\sqrt{50}$.

Observation model e	Feature regression m_ω	Mean signed pseudo-outcome bias	Stage 1 target MAE (1 SD)
Correct	Correct	−0.002	0.049 (0.002)
Correct	Misspecified	0.002	0.118 (0.007)
Misspecified	Correct	−0.001	0.026 (0.001)
Misspecified	Misspecified	0.809	0.808 (0.003)

Table 8: DFPV fixed-feature Stage 1 nuisance-specification diagnostic over 50 replications.

The signed pseudo-outcome bias remains close to zero and the Stage 1 target MAE remains small when either nuisance model is correctly specified. When both nuisance models are misspecified, the signed bias and target MAE are both approximately 0.81, matching the product-remainder pattern in Equation 26.