## A Hyper-Parameter Tuning

If observations from the joint distribution of $(A, Y, Z, W)$ are available in both stages, we can tune the regularization parameters $\lambda_{1}, \lambda_{2}$ using the approach proposed in Singh et al. [30], Xu et al. [35]. Let the complete data of stage 1 and stage 2 be denoted as $\left(a_{i}, y_{i}, z_{i}, w_{i}\right)$ and $\left(\tilde{a}_{i}, \tilde{y}_{i}, \tilde{z}_{i}, \tilde{w}_{i}\right)$. Then, we can use the data not used in each stage to evaluate the out-of-sample performance of the other stage. Specifically, let Algorithm 1 converges at $t=T$, and the regularizing parameters are given by

$$
\begin{aligned}
& \lambda_{1}^{*}=\arg \min \mathcal{L}_{1-\mathrm{oos}}, \quad \mathcal{L}_{1-\mathrm{oos}}=\frac{1}{n} \sum_{j=1}^{n}\left\|\boldsymbol{\psi}_{\theta_{W}^{(T)}}\left(\tilde{w}_{j}\right)-\hat{\boldsymbol{V}}^{(T)}\left(\boldsymbol{\phi}_{\theta_{A(1)}^{(T)}}\left(\tilde{a}_{i}\right) \otimes \boldsymbol{\phi}_{\theta_{Z}^{(T)}}\left(\tilde{z}_{i}\right)\right)\right\|^{2}, \\
& \lambda_{2}^{*}=\arg \min \mathcal{L}_{2-\mathrm{oos}}, \\
& \quad \mathcal{L}_{2-\mathrm{oos}}=\frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-\left(\boldsymbol{u}^{(T)}\right)^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}^{(T)}}\left(a_{i}\right) \otimes\left(\hat{\boldsymbol{V}}^{(T)}\left(\boldsymbol{\phi}_{\theta_{A(1)}^{(T)}}\left(a_{i}\right) \otimes \boldsymbol{\phi}_{\theta_{Z}^{(T)}}\left(z_{i}\right)\right)\right)\right)^{2}\right.
\end{aligned}
$$

where $\theta_{Z}^{(T)}, \theta_{A(1)}^{(T)}, \theta_{W}^{(T)}, \theta_{A(2)}^{(T)}, \hat{\boldsymbol{V}}^{(T)}, \boldsymbol{u}^{(T)}$ are the learned parameters by Algorithm 1.

## B Identifiability

In this appendix, we prove propositions given in the main text. In the following, we assume that the spaces $\mathcal{U}, \mathcal{A}, \mathcal{Z}, \mathcal{W}$ are separable and completely metrizable topological spaces and equipped with Borel $\sigma$-algebras. In this section, we use the notation $P_{A \mid Z=z}$ to express the distribution of a random variable $A$ given another variable $Z=z$.

## B. 1 Existence of bridge function

First, we discuss conditions to guarantee the existence of the bridge function $h$. Let us consider the following operators:

$$
\begin{aligned}
E_{a}: L^{2}\left(P_{W \mid A=a}\right) & \rightarrow L^{2}\left(P_{Z \mid A=a}\right), E_{a} f:=\mathbb{E}[f(W) \mid A=a, Z=\cdot], \\
F_{a}: L^{2}\left(P_{Z \mid A=a}\right) & \rightarrow L^{2}\left(P_{W \mid A=a}\right), F_{a} g:=\mathbb{E}[g(Z) \mid A=a, W=\cdot],
\end{aligned}
$$

where the conditional expectations are identified as equivalent classes with natural inclusion into their corresponding $L^{2}$ spaces. Our goal is to show that $\mathbb{E}[Y \mid A=a, Z=\cdot]$ is in the range of $E_{a}$, i.e., we seek a solution of the inverse problem defined by

$$
\begin{equation*}
E_{a} h=\mathbb{E}[Y \mid A=a, Z=\cdot] \tag{8}
\end{equation*}
$$

This suffices to prove the existence of the function $h$, for if there exists a function $h_{a}^{*}$ for each $a \in \mathcal{A}$ such that

$$
\mathbb{E}\left[h_{a}^{*} \mid A=a, Z=\cdot\right]=\mathbb{E}[Y \mid A=a, Z=\cdot],
$$

we can define $h^{*}(a, w):=h_{a}^{*}(w)$. The existence of the bridge function corresponds to the results of Deaner [3, Lemma 1.1.b] and Miao et al. [20, Propisition 1]. For simplicity, in constrast to Deaner [3], we assume that there is no shared quantity between the proxy variables $W$ and $Z$. Following the proofs of the previous works [3, Appendix C, p. 57] and [20, Appendix 6], we aim to solve the integral equation in (8).

We will make use of the following theorem on the existence of a solution of a linear integral equation [12, Theorem 15.18].
Proposition 5 ([12, Theorem 15.18]). Let $\mathcal{X}$ and $\mathcal{Y}$ be Hilbert spaces. Let $E: \mathcal{X} \rightarrow \mathcal{Y}$ be a compact linear operator with singular system $\left\{\left(\mu_{n}, \varphi_{n}, g_{n}\right)\right\}_{n=1}^{\infty}$. The equation of the first kind

$$
E \varphi=f
$$

is solvable if and only if $f \in N\left(E^{*}\right)^{\perp}$ and

$$
\sum_{n=1}^{\infty} \frac{1}{\mu_{n}^{2}}\left|\left\langle f, g_{n}\right\rangle\right|^{2}<\infty
$$

Here, $N\left(E^{*}\right)$ denotes the null space of the operator $E^{*}$. Then a solution is given by

$$
\phi=\sum_{n=1}^{\infty} \frac{1}{\mu_{n}}\left\langle f, g_{n}\right\rangle \varphi_{n}
$$

To apply Proposition 5, we make the following additional assumptions.
Assumption 4. For each $a \in \mathcal{A}$, the operator $E_{a}$ is compact with singular system $\left\{\left(\mu_{a, n}, \varphi_{a, n}, g_{a, n}\right)\right\}_{n=1}^{\infty}$.
Assumption 5. For each $a \in \mathcal{A}$, the conditional expectation $f_{Y \mid a}:=\mathbb{E}[Y \mid A=a, Z=\cdot]$ satisfies

$$
\sum_{n=1}^{\infty} \frac{1}{\mu_{a, n}^{2}}\left|\left\langle f_{Y \mid a}, g_{a, n}\right\rangle_{L^{2}\left(P_{Z \mid A=a}\right)}\right|^{2}<\infty
$$

for a singular system $\left\{\left(\mu_{a, n}, \phi_{a, n}, g_{a, n}\right)\right\}_{n=1}^{\infty}$ given in Assumption 4.
Remark 1. Assumption 4 is a minimal requisite to apply Proposition 5. The existing works [3, 20] assume stronger conditions; the operator $E_{a}$ is assumed to be Hilbert-Schmidt, which implies the compactness. Deaner [3, Assumption A.1] assumes that the joint distribution of $W$ and $Z$ conditioned on $a$ is absolutely continuous with respect to the product measure of $P_{W \mid A=a}$ and $P_{Z \mid A=a}$, and its density is square integrable. The density function serves as the integral kernel of the operator $E_{a}$ whose $L^{2}\left(P_{W \mid A=a} \otimes P_{Z \mid A=a}\right)$-norm corresponds to the Hilbert-Schmidt norm of the operator, where we use $\otimes$ to denote the product measure. On the other hand, in the setting of Miao et al. [20], all probability distributions have densities with respect to the Lebesgue measure. The operator $E_{a}$ (denoted by $K_{x}$ in [20]) is defined by the relevant densities accordingly (see the paragraph after Lemma 2 of [20]). The compactness is, as in [3], established by the square-integrability of the integral kernel, which implies that the operator $E_{a}$ is Hilbert-Schmidt (see Condition A1 in [20]).

It is easy to see that Assumptions 4 and 5 are required for using Proposition 5. The remaining condition to show is that $\mathbb{E}[Y \mid A=a, Z=\cdot]$ is in $N\left(E_{a}^{*}\right)^{\perp}$. We show that the structural assumption (Assumption 1) and completeness assumption (Assumption 2) imply the required condition. The result below closely follows [3] in that we share the same completeness condition.
Lemma 1. Under Assumptions 1 and 2, the conditional expectation $\mathbb{E}[Y \mid A=a, Z=\cdot]$ is in the orthogonal complement of the null space $N\left(E_{a}^{*}\right)$.

Proof. We first show that the adjoint of $E_{a}$ is given by $F_{a}$. For the operator $E_{a}$, any $f \in L_{2}\left(P_{W \mid A=a}\right)$ and $g \in L_{2}\left(P_{Z \mid A=a}\right)$, we have

$$
\begin{aligned}
\left\langle E_{a} f, g\right\rangle_{L_{2}\left(P_{Z \mid A=a}\right)} & =\mathbb{E}_{Z \mid A=a}[\mathbb{E}[f(W) \mid A=a, Z] g(Z)] \\
& =\mathbb{E}_{Z \mid A=a}\left[\mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[f(W) \mid A=a, Z, U]] g(Z)\right] \\
& \stackrel{(a)}{=} \mathbb{E}_{Z \mid A=a}\left[\mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[f(W) \mid A=a, U]] g(Z)\right] \\
& =\mathbb{E}_{U, Z \mid A=a}[\mathbb{E}[f(W) \mid A=a, U] g(Z)] \\
& =\mathbb{E}_{U \mid A=a}[\mathbb{E}[f(W) \mid A=a, U] \mathbb{E}[g(Z) \mid A=a, U]]
\end{aligned}
$$

where (a) follows from $W \Perp(A, Z) \mid U$, which is from Assumption 1. Similarly,

$$
\begin{aligned}
& \left\langle f, F_{a} g\right\rangle_{L^{2}\left(P_{W \mid A=a}\right)}=\mathbb{E}_{W \mid A=a}[f(W) \mathbb{E}[g(Z) \mid A=a, W]] \\
& =\mathbb{E}_{W \mid A=a}\left[f(W) \mathbb{E}_{U \mid A=a, W}[\mathbb{E}[g(Z) \mid A=a, W, U]]\right] \\
& \stackrel{(\mathrm{b})}{=} \mathbb{E}_{W \mid A=a}\left[f(W) \mathbb{E}_{U \mid A=a, W}[\mathbb{E}[g(Z) \mid A=a, U]]\right] \\
& =\mathbb{E}_{W, U \mid A=a}[f(W) \mathbb{E}[g(Z) \mid A=a, U]] \\
& =\mathbb{E}_{U \mid A=a}[\mathbb{E}[f(W) \mid A=a, U] \mathbb{E}[g(Z) \mid A=a, U]] \\
& =\left\langle E_{a} f, g\right\rangle_{L_{2}\left(P_{Z \mid A=a}\right)} .
\end{aligned}
$$

Again, (b) is given by $W \Perp A, Z \mid U$ from Assumption 1. For any $f^{*} \in N\left(E_{a}^{*}\right)=N\left(F_{a}\right)$, by iterated expectations, we have

$$
\begin{align*}
0 & =\mathbb{E}\left[f^{*}(Z) \mid A=a, W=\cdot\right] \\
& =\mathbb{E}_{U}\left[\mathbb{E}\left[f^{*}(Z) \mid A, U, W\right] \mid A=a, W=\cdot\right] \\
& =\mathbb{E}_{U}\left[\mathbb{E}\left[f^{*}(Z) \mid A, U\right] \mid A=a, W=\cdot\right] . \tag{9}
\end{align*}
$$

From Assumption 2,

$$
\mathbb{E}[l(U) \mid A=a, W=w]=0 \quad \forall(a, w) \in \mathcal{A} \times \mathcal{W} \quad \Leftrightarrow \quad l(u)=0 \quad P_{U} \text {-a.e. }
$$

for all functions $l \in L^{2}\left(P_{U \mid A=a}\right)$. Note that $\mathbb{E}\left[f^{*}(Z) \mid A, U\right] \in L^{2}\left(P_{U \mid A=a}\right)$ since

$$
\begin{aligned}
\mathbb{E}_{U \mid A}\left[\mathbb{E}\left[f^{*}(Z) \mid A, U\right]^{2}\right] & \leq \mathbb{E}_{U \mid A}\left[\mathbb{E}\left[f^{*}(Z)^{2} \mid A, U\right]\right] \quad(\because \text { Jensen's Inequality }) \\
& =\mathbb{E}\left[f^{*}(Z)^{2} \mid A\right]<\infty
\end{aligned}
$$

Hence, (9) and Assumption 2 implies

$$
\mathbb{E}\left[f^{*}(Z) \mid A=a, U=\cdot\right]=0 \quad P_{U} \text {-a.s. }
$$

Then, the inner product between $f^{*}$ and $\mathbb{E}[Y \mid A=a, Z=\cdot]$ is given as follows:

$$
\begin{aligned}
\left\langle f^{*}, \mathbb{E}[Y \mid A=a, Z=\cdot]\right\rangle_{L^{2}\left(P_{Z \mid A=a}\right)} & =\mathbb{E}_{Z \mid A=a}\left[f^{*}(Z) \mathbb{E}[Y \mid A=a, Z]\right] \\
& =\mathbb{E}_{Z \mid A=a}\left[f^{*}(Z) \mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[Y \mid A=a, Z, U]]\right] \\
& \stackrel{(c)}{=} \mathbb{E}_{Z \mid A=a}\left[f^{*}(Z) \mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[Y \mid A=a, U]]\right] \\
& =\mathbb{E}_{U, Z \mid A=a}\left[f^{*}(Z) \mathbb{E}[Y \mid A=a, U]\right] \\
& =\mathbb{E}_{U \mid A=a}\left[\mathbb{E}_{Z \mid U, A=a}\left[f^{*}(Z)\right] \mathbb{E}[Y \mid A=a, U]\right] \\
& =0
\end{aligned}
$$

Again (c) holds from $Y \Perp Z \mid A, U$ in Assumption 1. Hence, we have

$$
\mathbb{E}[Y \mid A=a, Z=\cdot] \in N\left(E_{a}^{*}\right)^{\perp}
$$

Remark 2. The difference between the first condition in Assumption 2 and Condition 3 in [20] is in the approach to establishing that the conditional expectation $\mathbb{E}[Y \mid A=a, Z=\cdot]$ belongs to $N\left(F_{a}\right)^{\perp}$. More specifically, Condition 3 in [20] is equivalent to having $N\left(F_{a}\right)=\{0\}$ and so any nontrivial $L^{2}\left(P_{Z \mid A=a}\right)$-function is in the orthogonal complement. Assumption 2 does not require $N\left(F_{a}\right)=\{0\}$ but only implies

$$
N\left(F_{a}\right) \subset\left\{f \in L^{2}\left(P_{Z \mid A=a}\right): \mathbb{E}[f(Z) \mid A=a, U=\cdot]=0\right\}
$$

For our identifiability proof, we use the weaker condition, Assumption 2.
Now, we are able to apply Proposition 5, leading to the following lemma.
Lemma 2. Under Assumptions 1, 2, 4 and 5, for each $a \in \mathcal{A}$, there exists a function $h_{a}^{*} \in$ $L_{2}\left(P_{W \mid A=a}\right)$ such that

$$
\mathbb{E}[Y \mid A=a, Z=\cdot]=\mathbb{E}\left[h_{a}^{*}(W) \mid A=a, Z=\cdot\right]
$$

Proof. By Lemma 1, the regression function $\mathbb{E}[Y \mid A=a, Z=\cdot]$ is in $N\left(E_{a}^{*}\right)^{\perp}$. Therefore, by Proposition 5, under the given assumptions, there exists a solution to (8). Letting the solution be $h_{a}^{*}$ completes the proof.

## B. 2 Identifiability

Here, we show that the bridge function $h$ can be used to compute the various causal quantities. In addition to the assumptions required for the existence of bridge function listed in the previous section, we assume the conditional expectation of outcome given treatment is square integrable.

Assumption 6. For each $a \in \mathcal{A}$, under the observational distribution, we have

$$
\mathbb{E}\left[Y^{2} \mid A=a\right]<\infty
$$

Given the assumption, we can prove the following theorem.
Theorem 1 (Identifiability). Assume that there exists a function $h: \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R}$ such that for each $a \in \mathcal{A}$, the function $h(a, \cdot)$ is in $L^{2}\left(P_{W \mid A=a}\right)$ and satisfies

$$
\mathbb{E}[Y \mid A=a, Z=\cdot]=\mathbb{E}[h(a, W) \mid A=a, Z=\cdot] \quad P_{Z \mid A=a}-a . s
$$

Under Assumptions 1, 2 and 6, we have

$$
\begin{equation*}
\mathbb{E}_{W \mid U}[h(a, W)]=\mathbb{E}[Y \mid A=a, U] \tag{10}
\end{equation*}
$$

Proof. Note that we have

$$
\begin{aligned}
\mathbb{E}[h(a, W) \mid A=a, Z] & =\mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[h(a, W) \mid A=a, Z, U]] \\
& =\mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[h(a, W) \mid U]], \\
\mathbb{E}[Y \mid A=a, Z] & =\mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[Y \mid A=a, Z, U]] \\
& =\mathbb{E}_{U \mid A=a, Z}[\mathbb{E}[Y \mid A=a, U]],
\end{aligned}
$$

where the second line of each equation follows from Assumption 1. Moreover, by assumption, we have

$$
\begin{aligned}
\mathbb{E}_{U \mid A=a}\left[\mathbb{E}[h(a, W) \mid U]^{2}\right] & \leq \mathbb{E}_{W \mid A=a}\left[h(a, W)^{2}\right]<\infty, \text { and } \\
\mathbb{E}_{U \mid A=a}\left[\mathbb{E}[Y \mid A=a, U]^{2}\right] & \leq \mathbb{E}\left[Y^{2} \mid A=a\right]<\infty
\end{aligned}
$$

Note from Assumption 2, we have

$$
\mathbb{E}[l(U) \mid A=a, Z=z]=0 \quad \forall(a, z) \in \mathcal{A} \times \mathcal{Z} \quad \Leftrightarrow \quad l(u)=0 \quad P_{U} \text {-a.e. }
$$

Therefore, by setting $l(u)=\mathbb{E}[Y \mid A=a, U=u]-\mathbb{E}[h(a, W) \mid U=u]$, for all $a \in \mathcal{A}$, we have

$$
\begin{aligned}
& \mathbb{E}[h(a, W) \mid A=a, Z=\cdot]=\mathbb{E}[Y \mid A=a, Z=\cdot] \quad P_{Z \mid A=a} \text {-a.s. } \\
& \Leftrightarrow \mathbb{E}_{U \mid A=a, Z=.}[\mathbb{E}[h(a, W) \mid U]-\mathbb{E}[Y \mid A=a, U]]=0 \quad P_{Z \mid A=a} \text {-a.s. } \\
& \Leftrightarrow \mathbb{E}_{W \mid U=.}[h(a, W)]=\mathbb{E}[Y \mid A=a, U=\cdot] \quad P_{U} \text {-a.s. }
\end{aligned}
$$

Using Theorem 1, we can show following two corollaries, which are used in the main body.
Corollary 1. Let the assumptions in Theorem 1 hold. Given a bridge function $h^{*}$, we can estimate structural function $f_{\text {struct }}$ as

$$
\begin{aligned}
f_{\text {struct }}(a) & :=\mathbb{E}_{U}[\mathbb{E}[Y \mid A=a, U]] \\
& =\mathbb{E}_{W}\left[h^{*}(a, W)\right]
\end{aligned}
$$

Proof. From Theorem 1, we have

$$
\mathbb{E}_{U}[\mathbb{E}[Y \mid A=a, U]]=\mathbb{E}_{U}\left[\mathbb{E}_{W \mid U}\left[h^{*}(a, W)\right]\right]=\mathbb{E}_{W}\left[h^{*}(a, W)\right]
$$

Remark 3. The above corollary corresponds to the identifiablity results in obtained in the previous works [3, 20]. We follow the proof of Theorem 1.1.a in [3, Appendix B] (See also, Theorem 1 and Appendix 3 of [20]).
Corollary 2. Assume we are given a bridge function $h^{*}(a, w)$ that is jointly measurable. Suppose $C \Perp W \mid U$. With the assumptions in Theorem 1, we can write the value function of policy $\pi(C)$ as

$$
\begin{aligned}
v(\pi) & :=\mathbb{E}_{C, U}[\mathbb{E}[Y \mid A=\pi(C), U]] \\
& =\mathbb{E}_{C}\left[\mathbb{E}_{W \mid C}[h(\pi(C), W)]\right]
\end{aligned}
$$

Proof. From Theorem 1, we have

$$
\mathbb{E}_{C, U}[\mathbb{E}[Y \mid A=\pi(C), U]]=\mathbb{E}_{C, U}\left[\mathbb{E}_{W \mid U}\left[h^{*}(\pi(C), W)\right]\right]=\mathbb{E}_{C}\left[\mathbb{E}_{W \mid C}[h(\pi(C), W)]\right]
$$

where the last equality holds by the conditional independence $C \Perp W \mid U$.
Note that in the existence claim, the bridge function $h^{*}(a, w)$ is constructed by aggregating over $\left\{h_{a}^{*}\right\}_{a \in \mathcal{A}}$. This construction does not guarantee that the function is measurable with respect to $a$; the lack of measurability renders its expectation undefined. Thus, we have additionally assumed the joint measurability of the bridge function in Corollary 2. Validating this assumption is crucial theoretically, and we leave it for future work.

## C Consistency of DFPV algorithm

In this appendix, we prove consistency of the DFPV approach. Following Xu et al. [35], we establish consistency of the end-to-end procedure incorporating Stages 1 and 2. We establish the result by first
showing a Stage 1 consistency result (Lemma 3), and then establishing the consistency of Stage 2 with the empirical Stage 1 solution used as input (Lemma 4). The desired result then follows in Theorem 2. Here, we assume the bridge function $h^{*}(a, w)$ is jointly measurable so that the expectation of $h^{*}$ is defined.
Consistency results will be expressed in terms of the complexity of the function classes used in Stages 1 and 2, as encoded in the Rademacher complexity of the function space of these functions (see Proposition 6 below). Consistency for particular function classes can then be shown by establishing that the respective Rademacher complexities vanish. We leave for future work the task of demonstrating this property for individual function classes of interest.

## C. 1 Operator view of DFPV

The goal of DFPV is to learn a bridge function $h^{*}$, which satisfies

$$
\begin{equation*}
\mathbb{E}_{Y \mid A=a, Z=z}[Y]=\mathbb{E}_{W \mid A=a, Z=z}\left[h^{*}(a, W)\right] \tag{11}
\end{equation*}
$$

We use the model

$$
h(a, w)=\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \boldsymbol{\psi}_{\theta_{W}}(w)\right)
$$

and denote the hypothesis spaces for $\boldsymbol{\psi}_{\theta_{W}}$ and $h$ as $\mathcal{H}_{\boldsymbol{\psi}_{\theta_{W}}}: \mathcal{W} \rightarrow \mathbb{R}^{d_{W}}$ and $\mathcal{H}_{h}: \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R}$, respectively. To learn the parameters, we minimize the following stage 2 loss:

$$
\begin{aligned}
& \hat{\boldsymbol{u}}, \hat{\theta}_{W}, \hat{\theta}_{A(2)}=\underset{\boldsymbol{u}, \theta_{W}, \theta_{A(2)}}{\arg \min } \hat{\mathcal{L}}_{2}\left(\boldsymbol{u}, \theta_{X}, \theta_{A(2)}\right) \\
& \hat{\mathcal{L}}_{2}=\frac{1}{n} \sum_{i=1}^{n}\left(\tilde{y}_{i}-\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}\left(\tilde{a}_{i}\right) \otimes \hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]\left(\tilde{a}_{i}, \tilde{z}_{i}\right)\right)\right)^{2} .
\end{aligned}
$$

We denote the resulting estimated bridge function as

$$
\hat{h}(a, w)=(\hat{\boldsymbol{u}})^{\top}\left(\boldsymbol{\psi}_{\hat{\theta}_{A(2)}}(a) \otimes \boldsymbol{\psi}_{\hat{\theta}_{W}}(w)\right)
$$

For simplicity, we set all regularization terms to zero. Here, $\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]$ is the empirical conditional expectation operator, which maps an element of $\mathcal{H}_{{\sigma_{\theta_{W}}}}$ to some function $\boldsymbol{g}(a, z)$ which is defined as

$$
\begin{aligned}
& \hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]=\underset{\boldsymbol{g} \in \mathcal{G}}{\arg \min } \hat{\mathcal{L}}_{1}\left(\boldsymbol{g} ; \boldsymbol{\psi}_{\theta_{W}}\right) \\
& \hat{\mathcal{L}}_{1}=\frac{1}{n} \sum_{i=1}^{m}\left\|\boldsymbol{\psi}_{\theta_{W}}\left(w_{i}\right)-\boldsymbol{g}\left(a_{i}, z_{i}\right)\right\|^{2}
\end{aligned}
$$

where $\|\cdot\|$ is the $\ell_{2}$-norm, and $\mathcal{G}$ is an arbitrary function space. In DFPV, we specify $\mathcal{G}$ as the set consisting of functions $\boldsymbol{g}$ of the form

$$
\boldsymbol{g}=\boldsymbol{V}\left(\boldsymbol{\phi}_{\theta_{A(1)}}\left(a_{i}\right) \otimes \boldsymbol{\phi}_{\theta_{Z}}\left(z_{i}\right)\right) .
$$

Note that this formulation is equivalent to the one introduced in Section 3. With a slight abuse of notation, for $h(a, w)=\boldsymbol{u}^{\top} \boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \boldsymbol{\psi}_{\theta_{W}}(w) \in \mathcal{H}_{h}$, we define $\hat{\mathbb{E}}_{W \mid A, Z}[h]$ to be

$$
\hat{\mathbb{E}}_{W \mid A, Z}[h(A, W)](a, z)=\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)\right)
$$

since this is the empirical estimate of $\mathbb{E}_{W \mid A, Z}[h(A, W)]$.

## C. 2 Generalization errors for regression

Here, we bound the generalization errors of both stages using Rademacher complexity bounds [24].
Proposition 6. [Theorem 3.3 24, with slight modification] Let $\mathcal{S}$ be a measurable space and $\mathcal{H}$ be a family of functions mapping from $\mathcal{S}$ to $[0, M]$. Given fixed dataset $S=\left(s_{1}, s_{2}, \ldots, s_{n}\right) \in \mathcal{S}^{n}$, the empirical Rademacher complexity is given by

$$
\hat{\mathfrak{R}}_{S}(\mathcal{H})=\mathbb{E}_{\boldsymbol{\sigma}}\left[\frac{1}{n} \sup _{h \in \mathcal{H}} \sum_{i=1}^{n} \sigma_{i} h\left(s_{i}\right)\right]
$$

where $\boldsymbol{\sigma}=\left(\sigma_{1}, \ldots, \sigma_{n}\right)$, with $\sigma_{i}$ independent random variables taking values in $\{-1,+1\}$ with equal probability. Then, for any $\delta>0$, with probability at least $1-\delta$ over the draw of an i.i.d sample
$S$ of size $n$, each of following holds for all $h \in \mathcal{H}$ :

$$
\begin{aligned}
& \mathbb{E}[h(s)] \leq \frac{1}{n} \sum_{i=1}^{n} h\left(s_{i}\right)+2 \hat{\mathfrak{R}}_{S}+3 M \sqrt{\frac{\log 2 / \delta}{2 n}} \\
& \frac{1}{n} \sum_{i=1}^{n} h\left(s_{i}\right) \leq \mathbb{E}[h(s)]+2 \hat{\mathfrak{R}}_{S}+3 M \sqrt{\frac{\log 2 / \delta}{2 n}}
\end{aligned}
$$

We list the assumptions below.

## Assumption 7. The following hold:

1. Bounded outcome variable $|Y| \leq M$.
2. Bounded stage 1 hypothesis space: $\forall a \in \mathcal{A}, z \in \mathcal{Z},\|\boldsymbol{g}(a, z)\| \leq 1$.
3. Bounded stage 2 feature map $\forall w \in \mathcal{W},\left\|\boldsymbol{\psi}_{\theta_{W}}(w)\right\| \leq 1$.
4. Bounded stage 2 weight: $\forall\|\boldsymbol{v}\| \leq 1, \forall a \in \mathcal{A},\left|\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \boldsymbol{v}\right)\right| \leq M$.

Note that unlike Xu et al. [35], we consider the case where the bridge function $h^{*}$ and the conditional expectation $\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right]$ might not be included in the hypothesis spaces $\mathcal{H}_{h}$, $\mathcal{G}$, respectively.
As in Xu et al. [35], we leave aside questions of optimization. Thus, we assume that the optimization procedure over $\left(\theta_{A(1)}, \theta_{Z}, \boldsymbol{V}\right)$ is sufficient to recover $\hat{\mathbb{E}}_{W \mid A, Z}\left[\psi_{\theta_{W}}(W)\right]$, and that the optimization procedure over $\left(\theta_{A(2)}, \theta_{W}, \boldsymbol{u}\right)$ is sufficient to recover $\hat{h}$ (which requires, in turn, the correct $\hat{\mathbb{E}}_{W \mid A, Z}$, for this $\psi_{\theta_{W}}$ ). We emphasize that Algorithm 1 does not guarantee these properties. Based on these assumptions, we derive the generalization error in terms of $L_{2}$-norm with respect to $\mathbb{P}(A, W)$.
Following the discussion in Xu et al. [35, Lemma 1], we can show the following lemma that provides a generalization error bound for Stage 1 regression.
Lemma 3. Under Assumption 7, and given stage 1 data $S_{1}=\left\{\left(a_{i}, z_{i}, w_{i}\right)\right\}_{i=1}^{m}$, for any $\delta>0$, with at least probability $1-2 \delta$, we have

$$
\left\|\hat{\mathbb{E}}_{W \mid A, Z}[h(A, W)]-\mathbb{E}_{W \mid A, Z}[h(A, W)]\right\|_{P(A, Z)} \leq M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}
$$

for any $\psi_{\theta_{W}} \in \mathcal{H}_{\psi_{\theta_{W}}}$, where hypothesis space $\mathcal{H}_{1}$ is defined as

$$
\begin{equation*}
\mathcal{H}_{1}=\left\{(w, a, z) \in \mathcal{W} \times \mathcal{A} \times \mathcal{Z} \mapsto\left\|\boldsymbol{\psi}_{\theta_{W}}(w)-\boldsymbol{g}(a, z)\right\|^{2} \in \mathbb{R} \mid \boldsymbol{g} \in \mathcal{G}, \boldsymbol{\psi}_{\theta_{W}} \in \mathcal{H}_{\boldsymbol{\psi}_{\theta_{W}}}\right\} \tag{12}
\end{equation*}
$$

and $\kappa_{1}$ is the misspecificaiton error in the stage 1 regression defined as

$$
\begin{equation*}
\kappa_{1}=\max _{\boldsymbol{\psi}_{\theta_{W}} \in \mathcal{H}_{\psi_{\theta_{W}}}} \min _{\boldsymbol{g} \in \mathcal{G}} \mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\boldsymbol{g}(A, Z)\right\|^{2}\right] . \tag{13}
\end{equation*}
$$

Proof. From Assumption 7, we have

$$
\begin{aligned}
& \left\|\hat{\mathbb{E}}_{W \mid A, Z}[h(A, W)]-\mathbb{E}_{W \mid A, Z}[h(A, W)]\right\|_{P(A, Z)} \\
& =\left\|\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes\left(\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)-\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)\right)\right)\right\|_{P(A, Z)} \\
& =\sqrt{\mathbb{E}_{A, Z}\left[\left|\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes\left(\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)-\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)\right)\right)\right|^{2}\right]} \\
& \leq \sqrt{\mathbb{E}_{A, Z}\left[M^{2}\left\|\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)-\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)\right\|^{2}\right]} \\
& =M\left\|\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)-\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}\right](a, z)\right\|_{P(A, Z)}
\end{aligned}
$$

From Proposition 6, we have

$$
\begin{aligned}
& \mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right](A, Z)\right\|^{2}\right] \\
& \quad \leq \hat{\mathcal{L}}_{1}\left(\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]\right)+2 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+12 \sqrt{\frac{\log 2 / \delta}{2 m}}
\end{aligned}
$$

with probability $1-\delta$, since all functions $f \in \mathcal{H}_{1}$ satisfy $\|f\| \leq 4$ from $\left\|\boldsymbol{\psi}_{\theta_{W}}\right\| \leq 1,\|\boldsymbol{g}\| \leq 1$. Now, let $\boldsymbol{g}_{\theta_{W}}^{*}$ be

$$
\boldsymbol{g}_{\theta_{W}}^{*}=\underset{\boldsymbol{g} \in \mathcal{G}}{\arg \min } \mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\boldsymbol{g}(A, Z)\right\|^{2}\right]
$$

Then, again from Proposition 6, we have

$$
\hat{\mathcal{L}}_{1}\left(\boldsymbol{g}_{\theta_{W}}^{*}\right) \leq \mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\boldsymbol{g}_{\theta_{W}}^{*}(A, Z)\right\|^{2}\right]+2 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+12 \sqrt{\frac{\log 2 / \delta}{2 m}}
$$

From the optimality of $\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]$, we have $\hat{\mathcal{L}}_{1}\left(\boldsymbol{g}_{\theta_{W}}^{*}\right) \geq \hat{\mathcal{L}}_{1}\left(\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]\right)$, thus

$$
\begin{align*}
& \mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right](A, Z)\right\|^{2}\right] \\
& \quad \leq \mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\boldsymbol{g}_{\theta_{W}}^{*}(A, Z)\right\|^{2}\right]+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}} \tag{14}
\end{align*}
$$

holds with probability $1-2 \delta$. Since we have

$$
\begin{aligned}
& \mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-g(A, Z)\right\|^{2}\right] \\
& =\mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]\right\|^{2}\right]+\mathbb{E}_{A, Z}\left[\left\|g(A, Z)-\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]\right\|^{2}\right]
\end{aligned}
$$

for all $\boldsymbol{g} \in \mathcal{G}$, by subtracting $\mathbb{E}_{W, A, Z}\left[\left\|\boldsymbol{\psi}_{\theta_{W}}(W)-\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]\right\|^{2}\right]$ from both sides of (14), we have

$$
\begin{aligned}
& \mathbb{E}_{A, Z}\left[\left\|\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]-\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right](A, Z)\right\|^{2}\right] \\
& \quad \leq \mathbb{E}_{A, Z}\left[\left\|\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]-\boldsymbol{g}_{\theta_{W}}^{*}(A, Z)\right\|^{2}\right]+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}} \\
& \quad \leq \kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}
\end{aligned}
$$

By taking the square root of both sides, we have

$$
\left\|\hat{\mathbb{E}}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]-\mathbb{E}_{W \mid A, Z}\left[\boldsymbol{\psi}_{\theta_{W}}(W)\right]\right\|_{P(A, Z)} \leq M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}
$$

The generalization error for Stage 2 can be shown by similar reasoning as in Xu et al. [35, Lemma 2].
Lemma 4. Under Assumption 7, given stage 1 data $S_{1}=\left\{\left(a_{i}, z_{i}, w_{i}\right)\right\}_{i=1}^{m}$, stage 2 data $S_{2}=$ $\left\{\left(\tilde{a}_{i}, \tilde{z}_{i}, \tilde{y}_{i}\right)\right\}_{i=1}^{n}$, and the estimated structural function $\hat{h}(a, z)=(\hat{\boldsymbol{u}})^{\top}\left(\boldsymbol{\psi}_{\hat{\theta}_{A(2)}}(a) \otimes \boldsymbol{\psi}_{\hat{\theta}_{W}}(w)\right)$, then for any $\delta>0$, with at least probability $1-4 \delta$, we have

$$
\begin{aligned}
& \left\|\mathbb{E}_{Y \mid A, Z}[Y]-\hat{\mathbb{E}}_{W \mid A, Z}[\hat{h}(A, W)]\right\|_{P(A, Z)} \\
& \quad \leq \kappa_{2}+M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{2}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}}
\end{aligned}
$$

where $\mathcal{H}_{1}$ is defined in (12), and $\mathcal{H}_{2}$ is defined as

$$
\begin{align*}
\mathcal{H}_{2}=\{(y, a, z) & \in \mathcal{Y} \times \mathcal{A} \times \mathcal{Z} \\
& \left.\mapsto\left(y-\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \boldsymbol{g}(a, z)\right)\right)^{2} \in \mathbb{R} \mid \boldsymbol{g} \in \mathcal{G}, \boldsymbol{u}, \theta_{A(2)}\right\} \tag{15}
\end{align*}
$$

and $\kappa_{2}$ is the misspecification error in Stage 2 defined as

$$
\begin{equation*}
\kappa_{2}=\min _{h \in \mathcal{H}_{h}}\left\|\mathbb{E}_{W \mid A, Z}[h(A, W)](A, Z)-\mathbb{E}_{Y \mid A, Z}[Y](A, Z)\right\|_{P(A, Z)} \tag{16}
\end{equation*}
$$

Proof. From Proposition 6, we have

$$
\mathbb{E}_{Y, A, Z}\left[\left|Y-\hat{\mathbb{E}}_{W \mid A, Z}[\hat{h}(A, W)](A, Z)\right|^{2}\right] \leq \hat{\mathcal{L}}_{2}(\hat{h})+2 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+12 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}
$$

with probabiltiy $1-\delta$, since all functions $f \in \mathcal{H}_{2}$ satisfy $\|f\| \leq 4 M^{2}$ from $\|Y\| \leq$ $M,\left|\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \boldsymbol{g}\right)\right| \leq M$. Let $\tilde{h}$ be

$$
\tilde{h}=\underset{h \in \mathcal{H}_{h}}{\arg \min }\left\|\mathbb{E}_{W \mid A, Z}[h(A, W)](A, Z)-\mathbb{E}_{Y \mid A, Z}[Y](A, Z)\right\|_{P(A, Z)}
$$

Again from Proposition 6, we have

$$
\hat{\mathcal{L}}_{2}(\tilde{h}) \leq \mathbb{E}_{Y, A, Z}\left[\left|Y-\hat{\mathbb{E}}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)\right|^{2}\right]+2 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+12 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}
$$

with probabiltiy $1-\delta$. From the optimality of $\hat{h}$, we have $\hat{\mathcal{L}}_{2}(\tilde{h}) \geq \hat{\mathcal{L}}_{2}(\hat{h})$, hence

$$
\begin{aligned}
& \mathbb{E}_{Y, A, Z}\left[\left|Y-\hat{\mathbb{E}}_{W \mid A, Z}[\hat{h}(A, W)](A, Z)\right|^{2}\right] \\
& \leq \mathbb{E}_{Y, A, Z}\left[\left|Y-\hat{\mathbb{E}}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)\right|^{2}\right]+4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}
\end{aligned}
$$

By subtracting $\mathbb{E}_{Y, A, Z}\left[\left\|Y-\mathbb{E}_{Y \mid A, Z}[Y]\right\|^{2}\right]$ from both sides, we have

$$
\begin{aligned}
& \mathbb{E}_{A, Z}\left[\left|\mathbb{E}_{Y \mid A, Z}[Y]-\hat{\mathbb{E}}_{W \mid A, Z}[\hat{h}(A, W)](A, Z)\right|^{2}\right] \\
& \leq \mathbb{E}_{A, Z}\left[\left|\mathbb{E}_{Y \mid A, Z}[Y]-\hat{\mathbb{E}}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)\right|^{2}\right]+4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}
\end{aligned}
$$

with probability $1-2 \delta$. By taking the square root of both sides, with probability $1-2 \delta$, we have

$$
\begin{aligned}
& \left\|\mathbb{E}_{Y \mid A, Z}[Y]-\hat{\mathbb{E}}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)\right\|_{P(A, Z)} \\
& \leq \sqrt{\mathbb{E}_{A, Z}\left[\left|\mathbb{E}_{Y \mid A, Z}[Y]-\hat{\mathbb{E}}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)\right|^{2}\right]+4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}} \\
& \leq\left\|\mathbb{E}_{Y \mid A, Z}[Y]-\hat{\mathbb{E}}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)\right\|_{P(A, Z)}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}} \\
& \stackrel{(a)}{\leq}\left\|\mathbb{E}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)-\hat{\mathbb{E}}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)\right\|_{P(A, Z)} \\
& \quad+\| \mathbb{E}_{W \mid A, Z}[\tilde{h}(A, W)](A, Z)-\mathbb{E}_{Y \mid A, Z} \\
& \begin{array}{l}
(b) \\
\leq
\end{array} \kappa_{2}+M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{2}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}}
\end{aligned}
$$

where (a) holds from the triangular inequality and (b) holds from Lemma 3.

Given the generalization errors in both stages, we can bound the error in (11) for the estimated bridge function $\hat{h}$. We need Assumption 3 to connect the error in (11) and the error $\left\|\hat{h}-h^{*}\right\|_{\mathbb{P}(A, W)}$. Let us restate Assumption 3.
Assumption 3 (Completeness Assumption on Outcome-Inducing Proxy [3, 18]). Let $l: \mathcal{W} \rightarrow \mathbb{R}$ be any square integrable function $\|l\|_{\mathbb{P}(W)}<\infty$, We assume the following condition:

$$
\mathbb{E}[l(W) \mid A=a, Z=z]=0 \quad \forall(a, z) \in \mathcal{A} \times \mathcal{Z} \quad \Leftrightarrow \quad l(w)=0 \quad \mathbb{P}(W) \text {-a.e. }
$$

Given this assumption, we can consider the following constant $\tau_{a}$.

$$
\tau_{a}=\max _{h \in L^{2}(\mathbb{P}(W \mid A=a)), h \neq h^{*}} \frac{\left\|h^{*}(a, W)-h(a, W)\right\|_{\mathbb{P}(W \mid A=a)}}{\left\|\mathbb{E}_{W \mid A=a, Z}\left[h^{*}(a, W)\right]-\mathbb{E}_{W \mid A=a, Z}[h(a, W)]\right\|_{\mathbb{P}(Z \mid A=a)}}
$$

Note that Assumption 3 ensures $\tau_{a}<\infty$. We can bound the error using the supremum of this constant, $\tau=\sup _{a} \tau_{a}$.
Theorem 2. Let Assumptions 3 and 7 hold. Given stage 1 data $S_{1}=\left\{\left(w_{i}, a_{i}, z_{i}\right)\right\}_{i=1}^{m}$ and stage 2 data $S_{2}=\left\{\left(\tilde{y}_{i}, \tilde{a}_{i}, \tilde{z}_{i}\right)\right\}_{i=1}^{n}$, for any $\delta>0$, with at least probability of $1-6 \delta$, we have

$$
\begin{aligned}
& \left\|h^{*}(A, W)-\hat{h}(A, W)\right\|_{P(A, W)} \\
& \quad \leq \tau\left(\kappa_{2}+2 M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}}\right)
\end{aligned}
$$

where $\tau=\sup _{a} \tau_{a}$ and $\mathcal{H}_{1}, \kappa_{1}, \mathcal{H}_{2}, \kappa_{2}$ are defined in (12), (13), (15), (16), respectively.

Proof.

$$
\begin{aligned}
& \left\|h^{*}(A, W)-\hat{h}(A, W)\right\|_{P(A, W)} \\
& \leq \sqrt{\mathbb{E}_{A}\left[\left\|h^{*}(a, W)-\hat{h}(a, W)\right\|_{\mathbb{P}(W \mid A=a)}^{2}\right]} \\
& \leq \sqrt{\mathbb{E}_{A}\left[\tau_{a}^{2}\left\|\mathbb{E}_{W \mid A=a, Z}\left[h^{*}(a, W)-\hat{h}(a, W)\right]\right\|_{\mathbb{P}(Z \mid A=a)}^{2}\right]} \\
& \leq \tau \sqrt{\mathbb{E}_{A}\left[\left\|\mathbb{E}_{W \mid A=a, Z}\left[h^{*}(a, W)-\hat{h}(a, W)\right]\right\|_{\mathbb{P}(Z \mid A=a)}^{2}\right]} \\
& \leq \tau\left\|\mathbb{E}_{Y \mid A, Z}[Y]-\mathbb{E}_{W \mid A, Z}[\hat{h}(A, W)]\right\|_{\mathbb{P}(A, Z)} \\
& \leq \tau\left\|\mathbb{E}_{Y \mid A, Z}[Y]-\hat{\mathbb{E}}_{W \mid A, Z}[\hat{h}(A, W)]\right\|_{\mathbb{P}(A, Z)}+\tau\left\|\mathbb{E}_{W \mid A, Z}[\hat{h}(A, W)]-\hat{\mathbb{E}}_{W \mid A, Z}[\hat{h}(A, W)]\right\|_{\mathbb{P}(A, Z)}
\end{aligned}
$$

Using Lemmas 3 and 4, the result thus follows.

From this result, we obtain the following corollary.
Corollary 3. Let Assumption 7 hold and $\kappa_{1}, \kappa_{2}=0$. If $\hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right) \rightarrow 0$ and $\hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right) \rightarrow 0$ in probability as the dataset size increases, $\hat{h}$ converges to $h^{*}$ in probability with respect to $\|\cdot\|_{\mathbb{P}(A, W)}$.

## C. 3 Consistency Result for Causal Parameters

In this section, we develop a consistency result for causal parameters discussed in the main body. First, we consider the structural function. Given estimated bridge function $\hat{h}(a, w)=\hat{\boldsymbol{u}}^{\top}\left(\boldsymbol{\psi}_{\hat{\theta}_{A(2)}}(a) \otimes\right.$ $\boldsymbol{\psi}_{\hat{\theta}_{W}}(w)$ ), we estimate the structural function by taking the empirical mean over $W$. To make the discussion simple, we assume the access to an additional data sample $S_{W}=\left\{w_{i}^{\text {extra }}\right\}_{i=1}^{n_{W}}$, such that the estimated structural function is given as

$$
\begin{equation*}
\hat{f}_{\text {struct }}(a)=\hat{\boldsymbol{u}}^{\top}\left(\boldsymbol{\psi}_{\hat{\theta}_{A(2)}}(a) \otimes \boldsymbol{\mu}_{\hat{\theta}_{W}}\right) \tag{17}
\end{equation*}
$$

where

$$
\boldsymbol{\mu}_{\hat{\theta}_{W}}=\sum_{i=1}^{n_{W}} \boldsymbol{\psi}_{\hat{\theta}_{W}}\left(w_{i}^{\text {extra }}\right)
$$

Note that empirically, we can use outcome-proxy data in $S_{1}$ instead of $S_{W}$.
To bound the deviation from the true structural function, we need the following assumption.
Assumption 8. Let $\rho_{W}(w), \rho_{W \mid A}(w \mid a)$ be the respective density functions of probability distributions $\mathbb{P}(W), \mathbb{P}(W \mid A)$. The densities satisfy

$$
\eta_{a}:=\left\|\frac{\rho_{W}(W)}{\rho_{W \mid A}(W \mid a)}\right\|_{\mathbb{P}(W \mid A=a)}<\infty
$$

Given Assumption 8, we can bound the error in structural function estimation. Before stating the theorem, let us introduce a useful concentration inequality for multi-dimensional random variables. ${ }^{2}$

Lemma 5. Let $\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{n} \in[-1,1]^{d}$ be independent random variables. Then, with probability at least $1-\delta$, we have

$$
\left\|\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{x}_{i}-\boldsymbol{\mu}\right\| \leq \sqrt{\frac{2 d \log 2 d / \delta}{n}}
$$

where $\boldsymbol{\mu}=\mathbb{E}\left[\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{x}_{i}\right]$.

Proof. Let $j$-th coordinate of $\boldsymbol{x}_{i}$ be denoted as $\left(\boldsymbol{x}_{i}\right)_{j}$. Then,

$$
\begin{aligned}
& \mathbb{P}\left(\left\|\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{x}_{i}-\boldsymbol{\mu}\right\| \leq \varepsilon\right) \\
& =\mathbb{P}\left(\sum_{j=1}^{d}\left(\frac{1}{n} \sum_{i=1}^{n}\left(\boldsymbol{x}_{i}\right)_{j}-(\boldsymbol{\mu})_{j}\right)^{2} \leq \varepsilon^{2}\right) \\
& \leq \mathbb{P}\left(\bigcap_{j=1}^{d}\left|\frac{1}{n} \sum_{i=1}^{n}\left(\boldsymbol{x}_{i}\right)_{j}-(\boldsymbol{\mu})_{j}\right| \leq \frac{\varepsilon}{\sqrt{d}}\right) \\
& \leq 1-2 d \exp \left(-\frac{n \varepsilon^{2}}{2 d}\right)
\end{aligned}
$$

Where the last inequality holds from Hoeffding's inequality. The claim follows by solving $\delta=$ $d \exp \left(-\frac{n \varepsilon^{2}}{2 d}\right)$.

Given this lemma, we can prove the following result.
Theorem 3. Let Assumptions 3, 7 and 8 hold. Let $\eta=\sup _{a \in \mathcal{A}} \eta_{a}$. Given stage 1 data $S_{1}=$ $\left\{\left(w_{i}, a_{i}, z_{i}\right)\right\}_{i=1}^{m}$, stage 2 data $S_{2}=\left\{\left(\tilde{y}_{i}, \tilde{a}_{i}, \tilde{z}_{i}\right)\right\}_{i=1}^{n}$, additional outcome-proxy variable data $S_{W}=$ $\left\{w_{i}^{\text {extra }}\right\}_{i=1}^{n_{W}}$, then with probability at least $1-7 \delta$, we have
$\left\|f_{\text {struct }}^{*}-\hat{f}_{\text {struct }}\right\|_{\mathbb{P}(A)}$

$$
\leq \sqrt{\frac{2 d_{W} \log 2 d_{W} / \delta}{n_{W}}}+\eta \tau\left(\kappa_{2}+2 M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}}\right)
$$

where $\hat{f}_{\text {struct }}$ is given in (17) and $d_{W}$ is the dimension of $\psi_{\theta_{W}}$.

Proof. From the relationship between structural function and bridge function, we have

$$
\begin{aligned}
& \left\|f_{\text {struct }}^{*}-\hat{f}_{\text {struct }}\right\|_{\mathbb{P}(A)} \\
& =\left\|\mathbb{E}_{W}\left[h^{*}(A, W)\right]-\hat{f}_{\text {struct }}\right\|_{\mathbb{P}(A)} \\
& \leq\left\|\mathbb{E}_{W}\left[h^{*}(A, W)\right]-\mathbb{E}_{W}[\hat{h}(A, W)]\right\|_{\mathbb{P}(A)}+\left\|\mathbb{E}_{W}[\hat{h}(A, W)]-\hat{f}_{\text {struct }}\right\|_{\mathbb{P}(A)} .
\end{aligned}
$$

[^0]We can bound each term as follows. For the first term, we have

$$
\begin{aligned}
& \left\|\mathbb{E}_{W}\left[h^{*}(\cdot, W)\right]-\mathbb{E}_{W}[\hat{h}(\cdot, W)]\right\|_{\mathbb{P}(A)} \\
& =\sqrt{\int\left(\mathbb{E}_{W}\left[h^{*}(a, W)-\hat{h}(a, W)\right]\right)^{2} \rho_{A}(a) \mathrm{d} a} \\
& =\sqrt{\int\left(\mathbb{E}_{W \mid A=a}\left[\left(h^{*}(a, W)-\hat{h}(a, W)\right) \frac{\rho_{W}(W)}{\rho_{W \mid A}(W \mid a)}\right]\right)^{2} \rho_{A}(a) \mathrm{d} a} \\
& \leq \sqrt{\int\left\|h^{*}(a, W)-\hat{h}(a, W)\right\|_{\mathbb{P}(W \mid A=a)}^{2}\left\|\frac{\rho_{W}(W)}{\rho_{W \mid A}(W \mid a)}\right\|_{\mathbb{P}(W \mid A=a)}^{2} \rho_{A}(a) \mathrm{d} a} \quad \because \text { Cauchy-Schwarz inequality } \\
& \leq \sqrt{\int \eta_{a}^{2} \mathbb{E}_{W \mid A=a}\left[\left(h^{*}(a, W)-\hat{h}(a, W)\right)^{2}\right] \rho_{A}(a) \mathrm{d} a} \quad \because \text { Assumption } 8 \\
& \leq \eta \sqrt{\mathbb{E}_{W, A}\left[\left(h^{*}(A, W)-\hat{h}(A, W)\right)^{2}\right]} \\
& =\eta\left\|h^{*}-\hat{h}\right\|_{\mathbb{P}(A, W)}
\end{aligned}
$$

From Theorem 2, with probability at least $1-6 \delta$, we have

$$
\begin{aligned}
& \left\|\mathbb{E}_{W}\left[h^{*}(\cdot, W)\right]-\mathbb{E}_{W}[\hat{h}(\cdot, W)]\right\|_{\mathbb{P}(A)} \\
& \leq \eta \tau\left(\kappa_{2}+2 M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}}\right)
\end{aligned}
$$

For the second term, from Assumption 7, we have with probability at least $1-\delta$,

$$
\begin{aligned}
\left\|\mathbb{E}_{W}[\hat{h}(A, W)]-\hat{f}_{\text {struct }}\right\|_{\mathbb{P}(A)} & \leq M\left\|\boldsymbol{\mu}_{\hat{\theta}_{W}}-\mathbb{E}_{W}\left[\boldsymbol{\psi}_{\hat{\theta}_{W}}\right]\right\| \\
& \leq \sqrt{\frac{2 d_{W} \log 2 d_{W} / \delta}{n_{W}}}
\end{aligned}
$$

The last inequality holds from Lemma 5. Using the uniform inequality, we have shown the claim.
We can evaluate the error in estimating a value function as well. Given $S_{3}=\left\{\check{w}_{i}, \check{c}_{i}\right\}_{i=1}^{n^{\prime}}$, we estimate the value function as

$$
\hat{v}(\pi)=\frac{1}{n^{\prime}} \sum_{i=1}^{n} \hat{\boldsymbol{u}}^{\top}\left(\boldsymbol{\psi}_{\hat{\theta}_{A(2)}}\left(\pi\left(\check{c}_{i}\right)\right) \otimes \boldsymbol{\psi}_{\hat{\theta}_{W}}\left(\check{w}_{i}\right)\right)
$$

Furthermore, we assume the following relationship between distributions of $A$ and $C$
Assumption 9. There exists a constant $\sigma$ such that

$$
\|l(\pi(C), W)\|_{\mathbb{P}(C, W)} \leq \sigma\|l(A, W)\|_{\mathbb{P}(A, W)}
$$

for all square integrable functions $l: \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R},\|l\|<\infty$.
Given these assumptions, we have the following theorem.
Theorem 4. Let Assumptions 3, 7, 8, 9 hold. Given stage 1 data $S_{1}=\left\{\left(w_{i}, a_{i}, z_{i}\right)\right\}_{i=1}^{m}$, stage 2 data $S_{2}=\left\{\left(\tilde{y}_{i}, \tilde{a}_{i}, \tilde{z}_{i}\right)\right\}_{i=1}^{n}$, stage 3 data $S_{3}=\left\{\check{w}_{i}, \check{c}_{i}\right\}_{i=1}^{n^{\prime}}$, with at least probability $1-7 \delta$, we have

$$
\begin{aligned}
& |v(\pi)-\hat{v}(\pi)| \\
& \leq \sigma \tau\left(\kappa_{2}+2 M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}}\right) \\
& \quad+\sqrt{\frac{M^{2} \log 2 / \delta}{n^{\prime}}}
\end{aligned}
$$



Figure 5: Causal graph with observable confounder

Proof. We have

$$
\begin{aligned}
|v(\pi)-\hat{v}(\pi)|= & \left|\mathbb{E}_{W \mid C}\left[h^{*}(\pi(C), W)\right]-\hat{v}(\pi)\right| \\
\leq & \left|\mathbb{E}_{W, C}[h(\pi(C), W)]-\mathbb{E}_{W, C}[\hat{h}(\pi(C), W)]\right| \\
& +\left|\mathbb{E}_{W, C}[\hat{h}(\pi(C), W)]-\hat{v}(\pi)\right| .
\end{aligned}
$$

We bound each term as follows. For the first term, we have

$$
\begin{aligned}
& \left|\mathbb{E}_{W, C}\left[h^{*}(\pi(C), W)\right]-\mathbb{E}_{W, C}[\hat{h}(\pi(C), W)]\right| \\
& \leq\left\|h^{*}(\pi(C), W)-\hat{h}(\pi(C), W)\right\|_{\mathbb{P}(C, W)} \quad \because \text { Jensen's inequality } \\
& \leq \sigma\left\|h^{*}(A, W)-\hat{h}(A, W)\right\|_{\mathbb{P}(A, W)} \quad \because \text { Assumption } 9
\end{aligned}
$$

Hence, from Theorem 2, we have

$$
\begin{aligned}
& \left|\mathbb{E}_{W, C}\left[h^{*}(\pi(C), W)\right]-\mathbb{E}_{W, C}[\hat{h}(\pi(C), W)]\right| \\
& \leq \sigma \tau\left(\kappa_{2}+2 M \sqrt{\kappa_{1}+4 \hat{\mathfrak{R}}_{S_{1}}\left(\mathcal{H}_{1}\right)+24 \sqrt{\frac{\log 2 / \delta}{2 m}}}+\sqrt{4 \hat{\mathfrak{R}}_{S_{2}}\left(\mathcal{H}_{2}\right)+24 M^{2} \sqrt{\frac{\log 2 / \delta}{2 n}}}\right)
\end{aligned}
$$

with probability at least $1-6 \delta$.
For the second term, we have

$$
\left|\mathbb{E}_{W, C}\left[\hat{\boldsymbol{u}}^{\top}\left(\boldsymbol{\psi}_{\hat{\theta}_{A(2)}}(\pi(C)) \otimes \boldsymbol{\psi}_{\hat{\theta}_{W}}(W)\right)\right]-\hat{v}(\pi)\right| \leq \sqrt{\frac{M^{2} \log 2 / \delta}{n^{\prime}}}
$$

with probability at least $1-\delta$, from Hoeffding's inequality. The result is obtained by taking the union bound.

## D DFPV algorithm with observable confounders

In this appendix, we formulate the DFPCL method when observable confounders are present, building on Tchetgen et al. [31] and Mastouri and Zhu et al. [18]. Here, we consider the causal graph given in Figure 5. In addition to variables $(A, Y, Z, W)$, we have an observable confounder $X \in \mathcal{X}$. The structural function $f_{\text {struct }}$ we aim to learn is

$$
f_{\text {struct }}(a)=\mathbb{E}_{X, U}[\mathbb{E}[Y \mid A=a, X, U]] .
$$

The structural assumption and completeness assumption including observable confounders are given as follows.

Assumption 10 (Structural Assumption [18]). We assume $Y \Perp Z \mid A, U, X$, and $W \Perp(A, Z) \mid U, X$.
Assumption 11 (Completeness Assumption [18]). Let $l: \mathcal{U} \rightarrow \mathbb{R}$ be any square integrable function $\|l\|_{\mathbb{P}(U)}$. We assume the following:

$$
\begin{aligned}
& \mathbb{E}[l(U) \mid A=a, Z=z, X=x]=0 \quad \forall(a, z, x) \in \mathcal{A} \times \mathcal{Z} \times \mathcal{X} \quad \Leftrightarrow \quad l(u)=0 \text { a.s. } \\
& \mathbb{E}[l(U) \mid A=a, Z=z, X=x]=0 \quad \forall(a, z, x) \in \mathcal{A} \times \mathcal{Z} \times \mathcal{X} \quad \Leftrightarrow \quad l(u)=0 \text { a.s. }
\end{aligned}
$$

Following similar reasoning as in Section 2, we can estimate the bridge function $\hat{h}: \mathcal{A} \times \mathcal{X} \times \mathcal{W} \rightarrow \mathbb{R}$ by minimizing the following loss:

$$
\hat{h}=\underset{h \in \mathcal{H}_{h}}{\arg \min } \tilde{\mathcal{L}}(h), \quad \tilde{\mathcal{L}}(h)=\mathbb{E}_{Y, A, Z, X}\left[\left(Y-\mathbb{E}_{W \mid Z, A, X}[h(A, X, W)]\right)^{2}\right]+\Omega(h)
$$

Given bridge function, we can estimate the structural function by

$$
f_{\text {struct }}(a)=\mathbb{E}_{X, W}\left[h^{*}(a, X, W)\right]
$$

Similar to Mastouri and Zhu et al. [18], we model

$$
\begin{aligned}
& \left.\mathbb{E}_{W \mid a, z}\left[\boldsymbol{\psi}_{\theta_{W}}(w)\right]=\boldsymbol{V}\left(\boldsymbol{\phi}_{\theta_{A(1)}}(A) \otimes \boldsymbol{\phi}_{\theta_{Z}}(Z)\right) \otimes \boldsymbol{\phi}_{\theta_{X(1)}}(X)\right) \\
& h(a, x, w)=\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \boldsymbol{\psi}_{\theta_{X(2)}}(x) \otimes \boldsymbol{\psi}_{\theta_{W}}(w)\right)
\end{aligned}
$$

where $\phi_{\theta_{X(1)}}(X), \boldsymbol{\psi}_{\theta_{X(2)}}(X)$ are the feature maps of $X$ parameterized by $\theta_{X(1)}, \theta_{X(2)}$, respectively. Then, in stage 1 , we learn $\left(\boldsymbol{V}, \theta_{A(1)}, \theta_{Z}, \theta_{X(1)}\right)$ by minimizing
$\hat{\mathcal{L}}_{1}\left(\boldsymbol{V}, \theta_{A(1)}, \theta_{Z}, \theta_{X(1)}\right)=\frac{1}{m} \sum_{i=1}^{m}\left\|\boldsymbol{\psi}_{\theta_{W}}\left(w_{i}\right)-\boldsymbol{V}\left(\boldsymbol{\phi}_{\theta_{A(1)}}\left(a_{i}\right) \otimes \boldsymbol{\phi}_{\theta_{Z}}\left(z_{i}\right) \otimes \boldsymbol{\phi}_{\theta_{X}}\left(x_{i}\right)\right)\right\|^{2}+\lambda_{1}\|\boldsymbol{V}\|^{2}$,
which estimates the conditional expectation $\mathbb{E}_{W}\left[\boldsymbol{\psi}_{W}(W) \mid A=a, Z=z, X=x\right]$. Let $\left(\hat{\boldsymbol{V}}, \hat{\theta}_{A(1)}, \hat{\theta}_{Z}, \hat{\theta}_{X(1)}\right)$ be the minimizer of $\hat{\mathcal{L}}_{1}$. Then, in stage 2 , we can learn $\left(\boldsymbol{u}, \theta_{A(2)}, \theta_{Z}, \theta_{X(2)}\right)$ using
$\hat{\mathcal{L}}_{2}\left(\boldsymbol{u}, \theta_{A(2)}, \theta_{W}, \theta_{X(2)}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(\tilde{y}_{i}-\boldsymbol{u}^{\top}\left(\boldsymbol{\psi}_{\theta_{A(2)}}\left(\tilde{a}_{i}\right) \otimes \boldsymbol{\psi}_{\theta_{X(2)}}\left(\tilde{x}_{i}\right) \otimes \hat{\boldsymbol{V}} \boldsymbol{v}_{1}\left(\tilde{a}_{i}, \tilde{x}_{i}, \tilde{z}_{i}\right)\right)^{2}+\lambda_{2}\|\boldsymbol{u}\|^{2}\right.$,
where we denote $\boldsymbol{v}_{1}(a, x, z)=\left(\boldsymbol{\phi}_{\hat{\theta}_{A(1)}}\left(\tilde{a}_{i}\right) \otimes \boldsymbol{\phi}_{\hat{\theta}_{X}}\left(\tilde{x}_{i}\right) \otimes \boldsymbol{\phi}_{\hat{\theta}_{Z}}\left(\tilde{z}_{i}\right)\right)$.
In DFPV, we first fix parameters $\boldsymbol{\theta}=\left(\theta_{A(1)}, \theta_{Z}, \theta_{X(1)}, \theta_{A(2)}, \theta_{Z}, \theta_{X(2)}\right)$ and obtain weights. This is given as

$$
\begin{equation*}
\hat{\boldsymbol{V}}(\boldsymbol{\theta})=\Psi_{1}^{\top} \Phi_{1}\left(\Phi_{1}^{\top} \Phi_{1}+m \lambda_{1} I\right)^{-1}, \quad \hat{\boldsymbol{u}}(\boldsymbol{\theta})=\left(\Phi_{2}^{\top} \Phi_{2}+n \lambda_{2} I\right)^{-1} \Phi_{2}^{\top} \boldsymbol{y}_{2} \tag{20}
\end{equation*}
$$

where we denote $\boldsymbol{\theta}=\left(\theta_{A(1)}, \theta_{Z}, \theta_{A(2)}, \theta_{W}\right)$ and define matrices as follows:

$$
\begin{array}{ll}
\Psi_{1}=\left[\boldsymbol{\psi}_{\theta_{W}}\left(w_{1}\right), \ldots, \boldsymbol{\psi}_{\theta_{W}}\left(w_{m}\right)\right]^{\top}, & \Phi_{1}=\left[\boldsymbol{v}_{1}\left(a_{1}, z_{1}\right), \ldots, \boldsymbol{v}_{1}\left(a_{m}, z_{m}\right)\right]^{\top} \\
\boldsymbol{y}_{2}=\left[\tilde{y}_{1}, \ldots, \tilde{y}_{n}\right]^{\top}, & \Phi_{2}=\left[\boldsymbol{v}_{2}\left(\tilde{a}_{1}, \tilde{z}_{1}\right), \ldots, \boldsymbol{v}_{2}\left(\tilde{a}_{n}, \tilde{z}_{n}\right)\right]^{\top}, \\
\boldsymbol{v}_{1}(a, x, z)=\boldsymbol{\phi}_{\theta_{A(1)}}(a) \otimes \boldsymbol{\phi}_{\theta_{X(1)}}(x) \otimes \boldsymbol{\phi}_{\theta_{Z}}(z), & \boldsymbol{v}_{2}(a, z)=\boldsymbol{\psi}_{\theta_{A(2)}}(a) \otimes \boldsymbol{\psi}_{\theta_{X(2)}}(x) \otimes\left(\hat{\boldsymbol{V}}(\boldsymbol{\theta}) \boldsymbol{v}_{1}(a, x, z)\right) .
\end{array}
$$

We learn the parameters by minimizing the following:

$$
\begin{aligned}
& \hat{\mathcal{L}}_{1}^{\mathrm{DFPV}}(\boldsymbol{\theta})=\frac{1}{m} \sum_{i=1}^{m}\left\|\boldsymbol{\psi}_{\theta_{W}}\left(w_{i}\right)-\hat{\boldsymbol{V}}(\boldsymbol{\theta}) \boldsymbol{v}_{1}\left(a_{i}, x_{i}, z_{i}\right)\right\|^{2}+\lambda_{1}\|\hat{\boldsymbol{V}}(\boldsymbol{\theta})\|^{2}, \\
& \hat{\mathcal{L}}_{2}^{\mathrm{DFPV}}(\boldsymbol{\theta})=\frac{1}{n} \sum_{i=1}^{n}\left(\tilde{y}_{i}-\hat{\boldsymbol{u}}(\boldsymbol{\theta})^{\top} \boldsymbol{v}_{2}\left(\tilde{a}_{i}, \tilde{x}_{i}, \tilde{z}_{i}\right)\right)^{2}+\lambda_{2}\|\hat{\boldsymbol{u}}(\boldsymbol{\theta})\|^{2} .
\end{aligned}
$$

The algorithm is given in Algorithn 2.

## E Additional Experiments

In this appendix, we report the results of two additional experiments. One is a synthetic setting introduced in Mastouri et al. [18], which has a simpler data generating process. The other is based on the real-world setting introduced by Deaner [3]. In both setting, DFPV performs similarly to or better than existing methods.

```
Algorithm 2 Deep Feature Instrumental Variable with Observable Confounder
Input: Stage 1 data \(\left(a_{i}, z_{i}, w_{i}, x_{i}\right)\), Stage 2 data \(\left(\tilde{a}_{i} \tilde{z}_{i}, \tilde{y}_{i}, \tilde{x}_{i}\right)\), Regularization parameters \(\left(\lambda_{1}, \lambda_{2}\right)\).
    Initial values \(\boldsymbol{\theta}^{(0)}=\left(\theta_{A(1)}^{(0)}, \theta_{Z}^{(0)}, \theta_{A(2)}^{(0)}, \theta_{W}^{(0)}, \theta_{X(1)}^{(0)}, \theta_{X(2)}^{(0)}\right)\). Learning rate \(\alpha\), additional data
    \(\left(x_{i}^{\text {extra }}, w_{i}^{\text {extra }}\right)\)
Output: Estimated structural function \(\hat{f}_{\text {struct }}(a)\)
    \(t \leftarrow 0\)
    repeat
        Compute \(\hat{\boldsymbol{V}}(\boldsymbol{\theta}), \hat{\boldsymbol{u}}(\boldsymbol{\theta})\) in (20)
        Update parameters in features \(\boldsymbol{\theta}^{(t+1)} \leftarrow\left(\theta_{A(1)}^{(t+1)}, \theta_{Z}^{(t+1)}, \theta_{A(2)}^{(t+1)}, \theta_{W}^{(t+1)}, \theta_{X(1)}^{(t+1)}, \theta_{X(2)}^{(t+1)}\right)\) by
        \(\theta_{A(1)}^{(t+1)} \leftarrow \theta_{A(1)}^{(t)}-\left.\alpha \nabla_{\theta_{A(1)}} \hat{\mathcal{L}}_{1}^{\mathrm{DFPV}}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(t)}}, \quad \theta_{Z}^{(t+1)} \leftarrow \theta_{Z}^{(t)}-\left.\alpha \nabla_{\theta_{Z}} \hat{\mathcal{L}}_{1}^{\mathrm{DFPV}}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(t)}}\)
        \(\theta_{X(1)}^{(t+1)} \leftarrow \theta_{X(1)}^{(t)}-\left.\alpha \nabla_{\theta_{X(1)}} \hat{\mathcal{L}}_{1}^{\mathrm{DFPV}}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(t)}}, \quad \theta_{X(2)}^{(t+1)} \leftarrow \theta_{X(2)}^{(t)}-\left.\alpha \nabla_{\theta_{X(2)}} \hat{\mathcal{L}}_{2}^{\mathrm{DFPV}}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(t)}}\)
        \(\theta_{A(2)}^{(t+1)} \leftarrow \theta_{A(2)}^{(t)}-\left.\alpha \nabla_{\theta_{A(2)}} \hat{\mathcal{L}}_{2}^{\mathrm{DFPV}}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(t)}}, \quad \theta_{W}^{(t+1)} \leftarrow \theta_{W}^{(t)}-\left.\alpha \nabla_{\theta_{W}} \hat{\mathcal{L}}_{2}^{\mathrm{DFPV}}(\boldsymbol{\theta})\right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(t)}}\)
```

        Increment counter \(t \leftarrow t+1\);
    until convergence
    Compute \(\hat{\boldsymbol{u}}\left(\boldsymbol{\theta}^{(t)}\right)\) from (20)
    8: Compute mean feature for $W$ using stage 1 dataset

$$
\boldsymbol{\mu}_{\theta_{X(2)} \otimes \theta_{W}} \leftarrow \frac{1}{n} \sum \boldsymbol{\psi}_{\theta_{X(2)}^{(t)}}\left(x_{i}^{\text {extra }}\right) \otimes \boldsymbol{\psi}_{\theta_{W}^{(t)}}\left(w_{i}^{\text {extra }}\right)
$$

9: return $\hat{f}_{\text {struct }}(a)=\left(\hat{\boldsymbol{u}}^{(t)}\right)^{\top}\left(\boldsymbol{\psi}_{\hat{\theta}_{A(2)}^{(t)}}(a) \otimes \boldsymbol{\mu}_{\theta_{X(2)} \otimes \theta_{W}}\right)$

## E. 1 Experiments with Simpler Data Generating Process

Here, we show the result for the synthetic setting proposed in Mastouri et al. [18]. The data generating process for each variable is given as follows:

$$
\begin{aligned}
U & :=\left[U_{1}, U_{2}\right], \quad U_{2} \sim \operatorname{Unif}[-1,2] \quad U_{1}, \sim \operatorname{Unif}[0,1]-\mathbb{1}\left[0 \leq U_{2} \leq 1\right] \\
Z & :=\left[U_{1}+\operatorname{Unif}[-1,1], U_{2}+\mathcal{N}(0,3)\right] \\
W & :=\left[U_{1}+\mathcal{N}(0,3), U_{2}+\operatorname{Unif}[-1,1]\right] \\
A & :=U_{2}+\mathcal{N}(0,0.05) \\
Y & :=U_{2} \cos \left(2\left(A+0.3 U_{1}+0.2\right)\right)
\end{aligned}
$$

From observations of $(Y, W, Z, A)$, we estimate $\hat{f}_{\text {struct }}$ by PCL. For each estimated $\hat{f}_{\text {struct }}$, we measure out-of-sample error as the mean square error of $\hat{f}$ versus true $f_{\text {struct }}$ obtained from MonteCarlo simulation. Specifically, we consider 20 evenly spaced values of $A \in[0.0,1.0]$ as the test data. The results with data size $n=m=\{500,1000\}$ are shown in Figure 6.
From Figure 6, we can see that DFPV and CEVAE methods perform worse and have larger variances than KPV and PMMR methods. This is not surprising, since DFPV tends to require more data than KPV and PMMR, as needed to learn the neural net feature maps (rather than using fixed pre-defined kernel features). Hence, we can say that we should favor KPV and PMMR over DFPV when the data is low-dimensional and the relations between the variables are smooth. We would like to note, however, DFPV outperforms CEVAE, which shows that the proxy setting is still required.

## E. 2 Experiments using Grade Retention dataset

To test the performance of DFPV in a more realistic setting, we conducted the experiment on the Grade Retention dataset introduced by Deaner [3]. This aims to estimate the effect of grade retention based on the score of math and reading on the long-term cognitive outcomes, in which we use scores in elementary school as a treatment-inducing proxy $(\mathrm{Z})$ and cognitive test scores from Kindergarten


Figure 6: Result of structural function experiment in the setting in Mastouri et al. [18]

|  | DFPV | CEVAE | KPV | PMMR |
| :---: | :---: | :---: | :---: | :---: |
| Math | $0.023(0.001)$ | $0.054(0.007)$ | $0.043(0.000)$ | $0.032(0.001)$ |
| Reading | $0.027(0.002)$ | $0.082(0.007)$ | $0.028(0.000)$ | $0.022(0.000)$ |

Table 1: Results of grade retension dataset
as the an outcome-inducing proxy (W). Following Mastouri et al. [18], we generate a synthetic "ground truth" by fitting a generalized additive model to learn a structured causal model (SCM), and a Gaussian mixture model to learn unmeasured confounder based on the learned SCM. Note, this is needed since for real-world data there is no measured ground truth.

Table 1 shows the result for this dataset. In this setting, the performance of DFPV matches KPV and PMMR. As in the experiment described in the revious section, the setting is low-dimensional (one-dim treatment variable, three-dim treatment-inducing proxy, four-dim outcome-inducing proxy) and the generative model is smooth (the "ground truth" being a generalized additive model and a Gaussian mixture model). For these reasons, we might again expect this data to favor kernel methods, such as KPV and PMMR; nonetheless, our method matches them. DFPV again outperforms CEVAE in this setting.

## F Experiment Details

In this section, we present the data generation process of experiments and the detailed settings of hyper-parameters.

## F. 1 Demand Design Experiment

Here, we introduce the details of demand design experiments. The observations are generated from the following causal model,

$$
Y=P\left(\exp \left(\frac{V-P}{10}\right) \wedge 5\right)-5 g(D)+\varepsilon, \quad \varepsilon \sim \mathcal{N}(0,1)
$$

where $Y$ represents sales, $P$ is the treatment variable (price), and these are confounded by potential demand $D$. Here we denote $a \wedge b=\min (a, b)$, and the function $g$ as

$$
g(d)=2\left(\frac{(d-5)^{4}}{600}+\exp \left(-4(d-5)^{2}\right)+\frac{d}{10}-2\right)
$$

To correct this confounding bias, we introduce cost-shifter $C_{1}, C_{2}$ as a treatment-inducing proxy, and views $V$ of the reservation page as the outcome-inducing proxy. Data is sampled as

$$
\begin{aligned}
& D \sim \operatorname{Unif}[0,10] \\
& C_{1} \sim 2 \sin (2 D \pi / 10)+\varepsilon_{1} \\
& C_{2} \sim 2 \sin (2 D \pi / 10)+\varepsilon_{2} \\
& V \sim 7 g(D)+45+\varepsilon_{3} \\
& P=35+\left(C_{1}+3\right) h(D)+C_{2}+\varepsilon_{4}
\end{aligned}
$$

where $\varepsilon_{1}, \varepsilon_{2}, \varepsilon_{3}, \varepsilon_{4} \sim \mathcal{N}(0,1)$. From observations of $\left(Y, P, C_{1}, C_{2}, V\right)$, we estimate $\hat{f}_{\text {struct }}$ by PCL. For each estimated $\hat{f}_{\text {struct }}$, we measure out-of-sample error as the mean square error of $\hat{f}$ versus true $f_{\text {struct }}$ obtained from Monte-Carlo simulation. Specifically, we consider 10 evenly spaced values of $p \in[10,30]$ as the test data.

## F. 2 dSprite Experiment

Here, we describe the data generation process for the dSprites dataset experiment. This is an image dataset parametrized via five latent variables (shape, scale, rotation, posX and posY). The images are $64 \times 64=4096$-dimensional. In this experiment, we fixed the shape parameter to heart, i.e. we only used the heart-shaped images. The other latent parameters take values of scale $\in[0.5,1]$, rotation $\in[0,2 \pi]$, posX $\in[0,1]$, posY $\in[0,1]$.
From this dataset, we generate the treatment variable $A$ and outcome $Y$ as follows:

1. Uniformly samples latent parameters (scale, rotation, posX, posY).
2. Generate treatment variable $A$ as

$$
A=\operatorname{Fig}(\text { scale }, \text { rotation, posX }, \text { posY })+\boldsymbol{\eta}_{A}
$$

3. Generate outcome variable $Y$ as

$$
Y=\frac{1}{12}(\operatorname{posY}-0.5) \frac{\|B A\|_{2}^{2}-5000}{1000}+\varepsilon, \quad \varepsilon \sim \mathcal{N}(0,0.5)
$$

Here, function Fig returns the corresponding image for the latent parameters, and $\boldsymbol{\eta}, \varepsilon$ are noise variables generated from $\boldsymbol{\eta}_{A} \sim \mathcal{N}(0.0,0.1 I)$ and $\varepsilon \sim \mathcal{N}(0.0,0.5)$. Each element of the matrix $B \in \mathbb{R}^{10 \times 4096}$ is generated from $\operatorname{Unif}(0.0,1.0)$ and fixed throughout the experiment. From the data generation process, we can see that $A$ and $Y$ are confounded by posY. Treatment variable $A$ is given as a figure corrupted with Gaussian random noise. The variable posY is not revealed to the model, and there is no observable confounder. The structural function for this setting is

$$
f_{\text {struct }}(A)=\frac{\|B A\|_{2}^{2}-5000}{1000}
$$

To correct this confounding bias, we set up the following PCL setting. We define the treatmentinducing variable $Z=$ (scale, rotation, $\operatorname{posX}) \in \mathbb{R}^{3}$, and the outcome-inducing variable by another figure that shares the same posY, with the remaining latent parameters fixed as follows:

$$
W=\operatorname{Fig}(0.8,0,0.5, \operatorname{posY})+\boldsymbol{\eta}_{W},
$$

where $\boldsymbol{\eta}_{W} \sim \mathcal{N}(0.0,0.1 I)$.
We use 588 test points for measuring out-of-sample error, which are generated from the grid points of latent variables. The grids consist of 7 evenly spaced values for posX, posY, 3 evenly spaced values for scale, and 4 evenly spaced values for orientation.

## F. 3 Policy Evaluation Experiments

We use the same data $\left(Y, P, C_{1}, C_{2}, V\right)$ in demand design for policy evaluation experiments. We consider two policies. One is a policy depends on costs $C_{1}, C_{2}$ which is

$$
\pi_{C_{1}, C_{2}}\left(C_{1}, C_{2}\right)=23+C_{1} C_{2} .
$$

To conduct offline-policy evaluation, we use data $\left(C_{1}, C_{2}, V\right)$ to compute the empirical average of $h\left(\pi_{C_{1}, C_{2}}\left(C_{1}, C_{2}\right), V\right)$. In our second experiment, the policy depends on current price $P$, which is

Table 2: Network structures of DFPV for demand design experiments. For the input layer, we provide the input variable. For the fully-connected layers (FC), we provide the input and output dimensions.

| Stage $\mathbf{1}$ Treatment Feature $\phi_{\theta_{A(1)}}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | Input $(P)$ |
| 2 | FC(1, 32), ReLU |
| 3 | $\mathrm{FC}(32,16), \operatorname{ReLU}$ |
| 4 | $\mathrm{FC}(16,8), \operatorname{ReLU}$ |

Stage 2 Treatment Feature $\psi_{\theta_{A(2)}}$

| Layer | Configuration |
| :---: | :---: |
| 1 | $\operatorname{Input}(P)$ |
| 2 | $\mathrm{FC}(1,32), \operatorname{ReLU}$ |
| 3 | $\mathrm{FC}(32,16), \operatorname{ReLU}$ |
| 4 | $\mathrm{FC}(16,8), \operatorname{ReLU}$ |


| Treatment-inducing Proxy Feature $\phi_{\theta_{Z}}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | $\operatorname{Input}\left(C_{1}, C_{2}\right)$ |
| 2 | $\mathrm{FC}(2,32), \operatorname{ReLU}$ |
| 3 | $\mathrm{FC}(32,16), \operatorname{ReLU}$ |
| 4 | $\mathrm{FC}(16,8), \operatorname{ReLU}$ |

Outcome-inducing Proxy Feature $\psi_{\theta_{W}}$

| Layer | Configuration |
| :---: | :---: |
| 1 | $\operatorname{Input}(V)$ |
| 2 | $\mathrm{FC}(1,32), \operatorname{ReLU}$ |
| 3 | $\mathrm{FC}(32,16), \operatorname{ReLU}$ |
| 4 | $\mathrm{FC}(16,8), \operatorname{ReLU}$ |

given as

$$
\pi_{P}(P)=\max (0.7 P, 10)
$$

Again, we use data $(P, V)$ to compute the empirical average of $h\left(\pi_{P}(P), V\right)$.

## F. 4 Hyper-parameters and network architectures

Here, we describe the network architecture and hyper-parameters of all experiments.
For KPV and PMMR method, we used the Gaussian kernel where the bandwidth is determined by the median trick. We follow the procedure for hyper-parameter tuning proposed in Mastouri et al. [18] in selecting the regularizers $\lambda_{1}, \lambda_{2}$.
For DFPV, we optimize the model using Adam [10] with learning rate $=0.001, \beta_{1}=0.9, \beta_{2}=0.999$ and $\varepsilon=10^{-8}$. Regularizers $\lambda_{1}, \lambda_{2}$ are both set to 0.1 as a result of the tuning procedure described in Appendix A. Network structure is given in Tables 2, 3.

In CEVAE, we attempt to reconstruct the latent variable $L$ from $(A, Z, W, Y)$ using a VAE. Here, we use a 20-dim latent variable $L$, whose the conditional distribution is specified as follows:

$$
q(L \mid A, Z, W, Y)=\mathcal{N}\left(\boldsymbol{V}_{1} \boldsymbol{\psi}_{q}(A, Z, W, Y), \operatorname{diag}\left(\boldsymbol{V}_{2} \boldsymbol{\psi}_{q}(A, Z, W, Y) \vee 0.1\right)\right)
$$

where $\boldsymbol{\psi}_{q}$ is a neural net and $\boldsymbol{V}_{1}, \boldsymbol{V}_{2}$ are matrices to be learned, and we denote $\boldsymbol{a} \vee 0.1=$ $\left(\max \left(a_{i}, 0.1\right)\right)_{i}$. Furthermore, we specify the likelihood distribution as follows:

$$
\begin{aligned}
& p(W, Z \mid L)=\mathcal{N}\left(\boldsymbol{V}_{3} \boldsymbol{\psi}_{p(W, Z \mid L)}(L), \operatorname{diag}\left(\boldsymbol{V}_{4} \boldsymbol{\psi}_{p(W, Z \mid L)}(L) \vee 0.1\right)\right) \\
& p(A \mid L)=\mathcal{N}\left(\boldsymbol{V}_{5} \boldsymbol{\psi}_{p(A \mid L)}(L), \operatorname{diag}\left(\boldsymbol{V}_{6} \boldsymbol{\psi}_{p(A \mid L))}(L) \vee 0.1\right)\right) \\
& p(Y \mid A, L)=\mathcal{N}\left(\boldsymbol{\mu}_{p(Y \mid A, L)}(A, L), 0.5\right)
\end{aligned}
$$

Here, $\boldsymbol{\psi}_{p(W, Z \mid L)}, \boldsymbol{\psi}_{p(A \mid L)}, \boldsymbol{\mu}_{p(Y \mid A, L)}$ are neural networks. We provide the structure of neural nets in Table 4 and 5. Following the orignal work [17], we train all neural nets by Adamax [10] with a learning rate of 0.01 , which was annealed with an exponential decay schedule. We also performed early stopping according to the lower bound on a validation set. To predict structural function, we obtain $q(L)$ by marginalizing $q(L \mid A, Z, W, Y)$ by observed data $A, Z, W, Y$. We then output $\hat{f}_{\text {struct }}(a)=\mathbb{E}_{L \sim q(L)}\left[\mathbb{E}_{Y \sim p(Y \mid A=a, L)}[Y]\right]$.

Table 3: Network structures of DFPV for dSprite experiments. For the input layer, we provide the input variable. For the fully-connected layers (FC), we provide the input and output dimensions. SN denotes Spectral Normalization [23]. BN denotes Batch Normalization.

| Stage 1 Treatment Feature $\phi_{\theta_{A(1)}}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | $\operatorname{Input}(A)$ |
| 2 | $\mathrm{FC}(4096,1024)$, SN, ReLU |
| 3 | $\mathrm{FC}(1024,512)$, SN, ReLU, BN |
| 4 | $\mathrm{FC}(512,128)$, SN, ReLU |
| 5 | $\mathrm{FC}(128,32)$, SN, ReLU |

Stage 2 Treatment Feature $\psi_{\theta_{A(2)}}$

| Layer | Configuration |
| :---: | :---: |
| 1 | $\operatorname{Input}(A)$ |
| 2 | FC(4096, 1024), SN, ReLU |
| 3 | FC(1024, 512), SN, ReLU, BN |
| 4 | FC $(512,128)$, SN, ReLU |
| 5 | FC $(128,32)$, SN, ReLU |


| Treatment-inducing Proxy Feature $\phi_{\theta_{Z}}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | $\operatorname{Input}(Z)$ |
| 2 | $\mathrm{FC}(3,128)$, ReLU |
| 3 | $\mathrm{FC}(128,64)$, ReLU |
| 4 | $\mathrm{FC}(64,32), \operatorname{ReLU}$ |

Outcome-inducing Proxy Feature $\psi_{\theta_{W}}$

| Layer | Configuration |
| :---: | :---: |
| 1 | $\operatorname{Input}(W)$ |
| 2 | FC(4096, 1024), SN, ReLU |
| 3 | FC(1024, 512), SN, ReLU, BN |
| 4 | FC(512, 128), SN, ReLU |
| 5 | FC $(128,32)$, SN, ReLU |

Table 4: Network structures of CEVAE for demand design experiment. For the input layer, we provide the input variable. For the fully-connected layers (FC), we provide the input and output dimensions.

| Structure of $\psi_{q}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | Input $\left(P, Y, C_{1}, C_{2}, V\right)$ |
| 2 | $\mathrm{FC}(5,128), \operatorname{ReLU}$ |
| 3 | $\mathrm{FC}(128,64)$, ReLU |
| 4 | $\mathrm{FC}(64,32), \operatorname{ReLU}$ |


| Structure of $\psi_{p(A \mid L)}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | Input $(L)$ |
| 2 | $\mathrm{FC}(20,64), \operatorname{ReLU}$ |
| 3 | $\mathrm{FC}(64,32), \operatorname{ReLU}$ |
| 4 | $\mathrm{FC}(32,16)$, ReLU |


| Structure of $\psi_{p(W, Z \mid L)}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | Input $(L)$ |
| 2 | $\mathrm{FC}(20,64)$, ReLU |
| 3 | $\mathrm{FC}(64,32)$, ReLU |
| 4 | $\mathrm{FC}(32,16)$, ReLU |

Structure of $\boldsymbol{\mu}_{p(Y \mid A, L)}$

| Layer | Configuration |
| :---: | :---: |
| 1 | Input $(L, P)$ |
| 2 | FC $(21,64)$, ReLU |
| 3 | FC $(64,32)$, ReLU |
| 4 | FC( 32,16$),$ ReLU |
| 5 | FC $(16,1)$ |

Table 5: Network structures of CEVAE for dSprite experiment. For the input layer, we provide the input variable. For the fully-connected layers (FC), we provide the input and output dimensions. SN denotes Spectral Normalization [23]. BN denotes Batch Normalization.

|  | Structure of $\psi_{q}$ |
| :---: | :---: |
| Layer | Configuration |
| 1 | Input( $W, A, Z, Y)$ |
| 2 | FC(8196, 1024), SN, ReLU |
| 3 | FC(1024, 512), SN, ReLU, BN |
| 4 | FC(512, 128), SN, ReLU |
| 5 | FC(128, 32), SN, ReLU |


| Structure of $\psi_{p(W, Z \mid L)}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | Input $(L)$ |
| 2 | FC $(20,64)$, ReLU |
| 3 | FC( 64,128$),$ ReLU |
| 4 | FC $(128,256)$, ReLU |

Structure of $\boldsymbol{\mu}_{p(Y \mid A, L)}$

| Structure of $\psi_{p(A \mid L)}$ |  |
| :---: | :---: |
| Layer | Configuration |
| 1 | $\operatorname{Input}(L)$ |
| 2 | FC $(20,64)$, ReLU |
| 3 | FC(64, 128), ReLU |
| 4 | FC $(128,256), \operatorname{ReLU}$ |


| Layer | Configuration |
| :---: | :---: |
| 1 | $\operatorname{Input}(L, A)$ |
| 2 | FC(4116, 1024), SN, ReLU, BN |
| 3 | FC(1024, 512), SN, ReLU, BN |
| 4 | FC(512, 128), SN, ReLU |
| 5 | FC(128, 32), SN, ReLU |
| 6 | FC(32, 1) |


[^0]:    ${ }^{2}$ Lemma 5 is discussed in MathOverflow (https://mathoverflow.net/questions/186097/ hoeffdings-inequality-for-vector-valued-random-variables, accessed on June 2nd 2021)

