

Estimating heterogeneous treatment effects with survival outcomes via a deep survival learner

Yuming Sun · Jian Kang · Yi Li*

Received: date / Accepted: date

Abstract Estimating heterogeneous treatment effects in survival settings is complicated by right censoring as well as the time-varying nature of the estimand. While the conditional average treatment effect (CATE) provides a natural target, most existing approaches focus on a single prespecified time point and do not account for temporal trajectory, leading to instability in estimation. We propose a deep survival learner (DSL) for estimating heterogeneous treatment effects with right-censored outcomes. The method is based on a doubly robust pseudo-outcome whose conditional expectation identifies time-specific CATEs under standard assumptions. This construction remains unbiased if either the outcome model or the treatment assignment model is correctly specified, when properly accounting for censoring. To estimate CATEs over a clinically relevant time spectrum, DSL employs a multi-output deep neural network with shared representations, enabling joint estimation of treatment effect trajectories. From a theoretical perspective, we derive error bounds for both pointwise and joint estimation over time. We show that joint estimation can leverage temporal structure to control estimation error without incurring much additional approximation cost under smoothness conditions, leading to improved stability relative to separate estimation. Cross-fitting is incorporated to reduce overfitting and mitigate bias arising from flexible nuisance estimation. Simulation studies demonstrate favorable finite-sample performance, particularly under nuisance model misspecification. Applied to the

Yuming Sun
Department of Mathematics, William & Mary, Williamsburg, VA, United States E-mail: ysun30@wm.edu

Jian Kang
Department of Biostatistics, University of Michigan, Ann Arbor, MI, United States E-mail: jiankang@umich.edu

Yi Li
Department of Biostatistics, University of Michigan, Ann Arbor, MI, United States
E-mail: yili@umich.edu

Boston Lung Cancer Study, DSL reveals heterogeneity in the effects of peri-operative chemotherapy across patient characteristics and over time.

Keywords Causal Inference · Conditional Average Treatment Effect · Double Robustness · Deep Neural Network

1 Introduction

A substantial literature has examined heterogeneity in treatment effects, often summarized by the conditional average treatment effect (CATE) within the counterfactual framework (Imbens & Rubin, 2016). Modeling CATE as a function of pre-treatment covariates enables systematic assessment of variation across patient subgroups, a central goal in precision medicine, where large-scale studies have documented marked heterogeneity in patient characteristics, disease awareness, and clinical outcomes across populations (Dharmarajan et al., 2017). In survival settings, this task is further complicated by censoring and the time-to-event nature of the outcome. Moreover, biomedical outcomes frequently arise from complex, interdependent processes involving treatment, exposure, and patient characteristics, as illustrated in HPV-related cancers (Haddad et al., 2008) and lung cancer (Zhai et al., 2022). These features challenge the adequacy of traditional parametric and semiparametric approaches. Recently, metalearner frameworks, including the M-learner, R-learner, and X-learner, have been developed. These approaches leverage the flexibility of modern machine learning methods by decomposing CATE estimation into a sequence of regression problems that can be solved with off-the-shelf algorithms. For instance, the M-learner constructs pseudo-outcomes based on inverse probability weighting (Horvitz & Thompson, 1952). This modified-outcome formulation reduces the problem of estimating CATEs to a regression on covariates (Tian et al., 2014). The X-learner (Künzel et al., 2019) fits separate outcome models in the treated and control groups and uses them to impute individual treatment effects, while the R-learner (Nie & Wager, 2021) formulates the problem as a residual-on-residual regression. More broadly, doubly robust learners have been developed to improve statistical efficiency and to mitigate sensitivity to nuisance model misspecification (Kennedy, 2020).

Building on developments for continuous outcomes, recent work has extended metalearner approaches to survival settings. For example, Yang et al. (2025) propose a doubly robust learner that combines inverse probability of censoring weighting with nonparametric failure-time Bayesian additive regression trees to construct pseudo-outcomes for CATE estimation. Similarly, Bo et al. (2025) develop alternative pseudo-outcome constructions and extend a range of metalearners, including the X-learner, M-learner, and R-learner, to right-censored data. Beyond the metalearner framework, tree-based and forest-based methods have also been adapted to estimate heterogeneous treatment effects with survival outcomes. Cui et al. (2023) introduce causal survival forests, which extend causal forests to accommodate censoring, while Dandl

et al. (2024) develop related model-based forest approaches. In a different direction, Hu (2024) propose a random-intercept accelerated failure time model combined with Bayesian additive regression trees to estimate differences in log survival time between treatment groups. Finally, Curth et al. (2021) introduce SurvITE, a neural network approach that models treatment-specific discrete-time hazard functions to capture treatment effect heterogeneity.

Several methodological challenges remain. Many pseudo-outcome based approaches rely on complete-case information and do not fully account for censoring, which can lead to bias or loss of efficiency in CATE estimation (Xu et al., 2023). Survival outcomes are inherently time-dependent, and treatment effects may evolve over follow-up due to delayed onset, attenuation, or time-varying biological mechanisms. In many applications, the primary scientific target is therefore the trajectory of heterogeneous treatment effects over a clinically relevant time interval, rather than a contrast at a single prespecified time point. Most existing methods, however, estimate CATE at a fixed time and treat different time points as unrelated, without leveraging the underlying temporal structure. This can result in unstable estimation and reduced statistical efficiency.

To address these limitations, we propose a deep survival learner (DSL) for estimating heterogeneous treatment effects with right-censored survival outcomes. We adopt a pseudo-outcome formulation that enables the use of flexible function approximators, such as deep neural networks, to learn its conditional expectation. Our approach distinguishes from existing methods by unifying causal identification, temporal modeling, and robust estimation within a single framework. First, we introduce a doubly robust pseudo-outcome for survival settings whose conditional expectation identifies the CATE on survival probabilities at each time point. The construction, when properly accounting for right censoring, is unbiased if either the outcome model or the treatment assignment model is correctly specified. Second, we cast CATE estimation as a joint learning problem over time and develop a multi-output neural network to estimate the entire treatment effect trajectory simultaneously. By sharing representations across time points, the method exploits temporal structure in the CATE, enabling information borrowing and leading to smoother and more stable estimates than separate models fit at each time point. From a theoretical perspective, we derive error bounds for both pointwise estimation and joint estimation over a time grid. We show that, under mild smoothness conditions on the treatment effect trajectory, joint estimation can control estimation error without incurring much additional approximation cost, thereby improving stability relative to separate estimation. Third, we incorporate cross-fitting to separate nuisance estimation from target estimation, which mitigates overfitting and enhances robustness when more flexible methods are used.

The remainder of the paper is organized as follows. Section 2 introduces the data structure, notation, and target estimand, and presents the proposed DNN estimator. Section 3 establishes the error rates at both a fixed time point and over a grid of time points. Section 4 reports simulations under varying degrees of nuisance model misspecification and different choices of time grids.

Section 5 applies the proposed method to the Boston Lung Cancer Study, a large prospective cohort in the country, to examine heterogeneous effects of perioperative chemotherapy among patients with early-stage non-small cell lung cancer.

2 Methods

2.1 Data Structure, Estimand and Pseudo-outcome Construction

Let T and $\mathcal{C}$ denote the event time and censoring time, respectively. With right censoring, we observe $U = \min(T, \mathcal{C})$ and the event indicator $\Delta = \mathbb{I}(T \leq \mathcal{C})$. We observe i.i.d. data $\mathcal{D} = \{(U_i, \Delta_i, X_i, W_i) : i = 1, \dots, N\}$, where $X_i \in \mathbb{R}^d$ denotes pre-treatment covariates and $W_i \in \{0, 1\}$ denotes treatment assignment. Under the potential outcomes framework (Imbens & Rubin, 2016), let T_i^w denote the potential survival time under treatment $w \in \{0, 1\}$. The observed event time satisfies $T_i = W_i T_i^1 + (1 - W_i) T_i^0$. Our target estimand is the conditional average treatment effect (CATE) on the survival probability at any time t in the support of U , defined as

$$\tau_0(x; t) = P(T_i^1 > t \mid X_i = x) - P(T_i^0 > t \mid X_i = x). \quad (1)$$

Since the potential survival times T_i^0 and T_i^1 in (1) cannot be observed simultaneously (even in the absence of censoring), direct estimation of CATE is infeasible. We propose to construct a pseudo-outcome whose expectation conditional on X_i equals the target estimand:

$$\begin{aligned} \varphi_i(t) &= \pi_i^1(t) - \pi_i^0(t), \text{ with} \\ \pi_i^w(t) &= S^w(X_i; t) + \frac{\mathbb{I}(W_i = w)}{e^w(X_i)} \left\{ \frac{Y_i(t)}{G^w(X_i; t)} - S^w(X_i; t) \right\}, \quad w \in \{0, 1\}, \end{aligned} \quad (2)$$

where $S^w(x; t)$, $G^w(x; t)$, and $e^w(x)$ denote the nuisance functions corresponding to the conditional survival function $P(T_i > t \mid X_i = x, W_i = w)$, the conditional censoring survival function $P(\mathcal{C}_i > t \mid X_i = x, W_i = w)$, and the propensity score $P(W_i = w \mid X_i = x)$, respectively. Here, for a fixed time point t , the at-risk indicator $Y_i(t) = \mathbb{I}(U_i > t)$. In practice, these nuisance functions are unknown and must be estimated from the data, and may be subject to model misspecification.

Despite this, the pseudo-outcome in (2) remains unbiased under partial misspecification due to its double robustness. Building on this property, we estimate the pseudo-outcome by evaluating it over a prespecified grid of time points, reflecting that the target CATE is indexed by time and typically of interest over an interval rather than at a single time point. We then model its conditional expectation given baseline covariates using a deep neural network (DNN) introduced in the next subsection and obtain a flexible estimator of the time-varying CATE function.

2.2 Deep Neural Network: Architecture and Implementation

Let $0 \leq t_{\min} < t_{\max}$, where $[t_{\min}, t_{\max}]$ is a compact subset of the support of U , which is specified to balance clinical relevance and data support. In our motivating lung cancer study, $t_{\min} = 0$ represents the time of diagnosis, whereas $t_{\max}$ is chosen to ensure that a sufficient number of patients remain at risk at that time for stable estimation. We estimate $\tau_0(x; t)$ on a prespecified grid of J time points, $t_{\min} = t_1 < t_2 < \dots < t_J = t_{\max}$.

To model the smoothness of the treatment effects over time, we represent the time-varying CATE using a fixed basis; common choices include B-splines or indicator basis functions (de Boor, 2001; Schumaker, 2007). With $\{B_k(t)\}_{k=1}^m$ being $m(\leq J)$ basis functions on $[t_{\min}, t_{\max}]$, we approximate the CATE by

$$\tau(x; t) = \sum_{k=1}^m f_k(x) B_k(t), \quad t \in [t_{\min}, t_{\max}] \quad (3)$$

which provides a concrete low-dimensional representation of temporal variation, facilitating our later theoretical analysis. To estimate (3), we consider a fully connected feedforward network with L hidden layers. Let $h^{(0)}(x) = x$ and, for $\ell = 1, \dots, L$,

$$h^{(\ell)}(x) = \sigma(W_\ell h^{(\ell-1)}(x) + b_\ell),$$

where $\sigma(u) = \max(u, 0)$ is the ReLU activation applied componentwise. The m spline coefficient functions in (3), i.e., $f(x) = (f_1(x), f_2(x), \dots, f_m(x))^\top$, are jointly estimated by the output layer

$$f(x) = W_{L+1} h^{(L)}(x) + b_{L+1} \in \mathbb{R}^m.$$

Let $\mathcal{F}_{n,m}$ denote the class of such networks subject to the following structural constraints:

- (Depth)** $L \leq L_n$;
- (Width)** $\max_{1 \leq \ell \leq L} p_\ell \leq p_n$;
- (Sparsity)** the total number of nonzero parameters satisfies $\sum_{\ell=1}^{L+1} \|W_\ell\|_0 + \sum_{\ell=1}^{L+1} \|b_\ell\|_0 \leq s_n$;
- (Magnitude constraint)** all entries of $\{W_\ell, b_\ell\}$ are bounded in absolute value by B_n .

The corresponding DNN class of time-varying CATE functions is

$$\mathcal{T}_{n,m} = \left\{ (x, t) \mapsto \sum_{k=1}^m f_k(x) B_k(t) : f = (f_1, \dots, f_m)^\top \in \mathcal{F}_{n,m} \right\}.$$

Taking $m = 1$ and $B_1(t) = 1$ recovers the scalar network for estimating $\tau_0(x; t)$ at a single t . When $J > 1$, let $\mathbf{B} = [B_k(t_j)]_{J \times m}$, with $m > 1$, denote the

fixed spline evaluation matrix. Then the CATE vector over the time grid is represented by

$$\boldsymbol{\tau}(x) = (\tau(x; t_1), \dots, \tau(x; t_J))^\top = \mathbf{B}f(x).$$

Thus, the spline formulation allows information to be shared across time through a common basis. As characterized in the subsequent Theorem 6, joint estimation can reduce approximation error when the CATE trajectory is sufficiently smooth in time and admits a low-complexity representation; this advantage may not hold when such temporal structure is weak or misspecified.

2.3 Estimation

Let $\hat{\boldsymbol{\varphi}}_i = (\hat{\varphi}_i(t_1), \dots, \hat{\varphi}_i(t_J))^\top$ denote the vector of estimated pseudo-outcomes for subject i on a prespecified grid of J time points, $t_{\min} = t_1 < t_2 < \dots < t_J = t_{\max}$. To reduce overfitting and permit flexible nuisance estimation, we employ K -fold cross-fitting throughout. Let the observed data be $\mathcal{D} = \{(U_i, \Delta_i, X_i, W_i)\}_{i=1}^N$, and randomly partition $\mathcal{D}$ into K approximately equal-sized, disjoint folds $\mathcal{D}_1, \dots, \mathcal{D}_K$. For each fold k , let $\mathcal{D}_{-k} = \mathcal{D} \setminus \mathcal{D}_k$ denote the corresponding training sample, and use $\mathcal{D}_k$ as the held-out sample.

Within each training sample $\mathcal{D}_{-k}$, we estimate the nuisance functions needed for pseudo-outcome construction. Specifically, for each treatment group $w \in \{0, 1\}$, we estimate $S^w(x; t)$, and $G^w(x; t)$ using only subjects in $\mathcal{D}_{-k}$ with $W_i = w$. These nuisance components may be estimated using Cox-type models, or, when greater flexibility is desired, using nonparametric survival learners such as random survival forests. Also, the propensity score $e^w(x)$ is estimated on $\mathcal{D}_{-k}$ using a classification method, such as logistic regression or a more flexible tree-based ensemble learner.

These estimators are then evaluated only on observations in the held-out fold $\mathcal{D}_k$. Thus, for each subject $i \in \mathcal{D}_k$, the pseudo-outcome vector $\hat{\boldsymbol{\varphi}}_i$ is constructed using nuisance estimates obtained from data that do not include subject i . Repeating this procedure over all K folds yields cross-fitted pseudo-outcomes for every subject in the sample.

We estimate the CATE at the evaluation times as

$$\hat{\tau}(x; t_j) = \sum_{k=1}^m \hat{f}_k(x) B_k(t_j), \quad j = 1, \dots, J,$$

where

$$\hat{f} = \arg \min_{f \in \mathcal{F}_{n,m}} \frac{1}{NJ} \sum_{i=1}^N \|\hat{\boldsymbol{\varphi}}_i - \mathbf{B}f(X_i)\|_2^2. \quad (4)$$

That is, the parameters $\{W_\ell, b_\ell\}_{\ell=1}^{L+1}$ associated with f are estimated by minimizing the empirical squared loss over subjects and time points. The resulting procedure, referred to as the Deep Survival Learner (DSL), is summarized in Algorithm 1.

Regardless of the specific choice of nuisance estimators, our theoretical analysis shows that the error of the DNN estimator decomposes into nuisance estimation error as well as approximation error and statistical error arising from fitting the DNN.

Algorithm 1 Deep Survival Learner (DSL)

Require: Observed data $\mathcal{D} = \{(U_i, \Delta_i, X_i, W_i)\}_{i=1}^N$, time grid $\{t_1, \dots, t_J\}$, spline basis $\{B_1, \dots, B_m\}$, number of folds K
Ensure: Estimated CATE $\hat{\tau}(\cdot; t)$ on $\{t_1, \dots, t_J\}$
1: Evaluate the spline basis on the time grid to calculate $\mathbf{B} = [B_k(t_j)]_{1 \leq j \leq J, 1 \leq k \leq m}$
2: Randomly partition $\mathcal{D}$ into K disjoint folds $\mathcal{D}_1, \dots, \mathcal{D}_K$
3: **for** $k = 1, \dots, K$ **do**
4: Let $\mathcal{D}_{-k} = \mathcal{D} \setminus \mathcal{D}_k$
5: Estimate nuisance functions $\hat{S}^w(x; t)$, $\hat{G}^w(x; t)$, and $\hat{e}^w(x)$ using $\mathcal{D}_{-k}$
6: **for** each $i \in \mathcal{D}_k$ **do**
7: Construct estimated pseudo-outcomes $\hat{\varphi}_i = (\hat{\varphi}_i(t_1), \dots, \hat{\varphi}_i(t_J))^\top$ using the fitted nuisance models
8: **end for**
9: **end for**
10: Fit a m -output DNN $\hat{f} \in \mathcal{F}_{n,m}$ by solving (4)

3 Theoretical Properties

Notation: For $w \in \{0, 1\}$ and $t \in [t_{\min}, t_{\max}]$, define

$$S_0^w(x; t) = P(T_i > t \mid X_i = x, W_i = w), e_0^w(x) = P(W_i = w \mid X_i = x),$$

$G_0^w(x; t) = P(\mathcal{C}_i > t \mid X_i = x, W_i = w)$. Let $\hat{S}^w(x; t)$, $\hat{e}^w(x)$, and $\hat{G}^w(x; t)$ denote the corresponding estimators. For subject i , let $O_i = (U_i, \Delta_i, X_i, W_i)$ denote the observed data. For any measurable function g , define the empirical measure by $\mathbb{P}_n g = \frac{1}{n} \sum_{i=1}^n g(O_i)$. Further, for a random function $\hat{g}(\cdot)$ constructed independently of O_i , the quantities $\mathbb{E}\{\hat{g}(O_i)\}$ and $\mathbb{E}\{\hat{g}(O_i) \mid X_i\}$ are understood as expectations taken with respect to O_i only, or $\mathbb{E}\{\hat{g}(O_i)\} = \int \hat{g}(o) dP_O(o)$, and $\mathbb{E}\{\hat{g}(O_i) \mid X_i = x\} = \int \hat{g}(o) dP_{O|X}(o|x)$. For a vector $x = (x_1, \dots, x_d)^\top \in \mathbb{R}^d$, define its ℓ_q -norm by $\|x\|_q = \left(\sum_{j=1}^d |x_j|^q\right)^{1/q}$, $1 \leq q < \infty$, and $\|x\|_\infty = \max_{1 \leq j \leq d} |x_j|$. Let P_X denote the probability distribution of the covariate vector X on $\mathcal{X}$. For a measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$, define $\|f\|_{L_q(P_X)} = \left(\int_{\mathcal{X}} |f(x)|^q dP_X(x)\right)^{1/q}$, $1 \leq q < \infty$, and $\|f\|_{L_\infty(P_X)} = \sup_{x \in \mathcal{X}} |f(x)|$.

We impose the following assumptions to establish the statistical properties of the proposed estimator.

Assumption 1 (Consistency) For each individual, $T_i = T_i^{W_i}$.

Assumption 2 (Unconfoundedness) $\{T_i^0, T_i^1\} \perp\!\!\!\perp W_i \mid X_i$.

Assumption 3 (Overlap) Assume the support of X is $\mathcal{X} = [0, 1]^d$, and P_X has a density bounded away from 0 and ∞ on $\mathcal{X}$. There exists a constant $c > 0$ such that $e_0^w(x) \geq c$, for all $x \in \mathcal{X}$ and $w \in \{0, 1\}$.

Assumption 4 (Noninformative censoring) For each $w \in \{0, 1\}$, we assume that $\mathcal{C}_i \perp\!\!\!\perp T_i^w \mid (X_i, W_i = w)$.

Assumption 5 (Censoring positivity) There exists a constant $c > 0$ such that $G_0^w(x; t) \geq c$, for all (x, t) in the support of (X_i, T_i) and all $w \in \{0, 1\}$.

Assumption 6 (Uniform convergence rate) For a fixed $t \in [t_{\min}, t_{\max}]$, assume the estimates, corresponding to $S_0^w(x; t)$, $G_0^w(x; t)$ and $e_0^w(x)$, satisfy

$$\sup_x |\hat{S}^w(x; t) - S_0^w(x; t)| = O_p(r_n^{(S)}), \quad \sup_x |\hat{G}^w(x; t) - G_0^w(x; t)| = O_p(r_n^{(G)}),$$

$$\sup_x |\hat{e}^w(x) - e_0^w(x)| = O_p(r_n^{(e)}),$$

for some positive sequences $r_n^{(S)}, r_n^{(G)}, r_n^{(e)}$ that converge to 0.

Assumptions 1–3 are standard in causal inference and ensure identification of the potential outcome distributions (Imbens & Rubin, 2015; Rosenbaum & Rubin, 1983). Assumptions 4–5 are commonly imposed in survival analysis to guarantee identifiability of the survival and censoring mechanisms (Andersen & Gill, 1982; Tsiatis, 2006). Assumption 6 requires uniform convergence of the estimated nuisance functions, which is standard in semiparametric and machine learning-based inference and determines the convergence rate of the proposed estimator (Chernozhukov et al., 2018b; van der Vaart, 1998).

Theorem 1 (Identification) Under Assumptions 1-5, if all nuisance functions in (2) equal their true counterparts, it holds that

$$\mathbb{E}\{\varphi_i(t) \mid X_i = x\} = \tau_0(x; t).$$

Theorem 1 establishes identification of $\tau_0(x; t)$ via the conditional expectation of $\varphi_i(t)$ given X_i , thereby reducing the problem of estimating heterogeneous treatment effects under censoring to a conditional mean regression problem.

Theorem 2 (Double robustness) Under Assumptions 1-5, suppose the censoring model is correctly specified, i.e., $G^w(x; t) = P(\mathcal{C}_i > t \mid X_i = x, W_i = w)$. If either the survival model is correctly specified, i.e., $S^w(x; t) = P(T_i > t \mid X_i = x, W_i = w)$, or the propensity score model is correctly specified, i.e., $e^w(x) = P(W_i = w \mid X_i = x)$, it holds that

$$\mathbb{E}\{\varphi_i(t) \mid X_i = x\} = \tau_0(x; t).$$

The double robustness of the proposed pseudo-outcome, established in Theorem 2, is especially appealing in observational survival studies, where it is often difficult to correctly specify all nuisance components.

The consistency of the final estimator further relies on the approximation and estimation properties of the DNN. We establish consistency for the scalar network estimator at a fixed time point t (i.e., $m = 1$ and $B_1(t) = 1$). We then extend the analysis to joint estimation at multiple time points with the temporal basis representation and a multi-output DNN. Recall that $\tau_0(x; t) = \mathbb{E}[\varphi_i^0(t) \mid X_i = x]$, where $\varphi_i^0(t)$ is the oracle pseudo-outcome and $\hat{\varphi}_i(t)$ is its cross-fitted counterpart (Section 2.3). The DNN estimator is obtained via

$$\hat{f} = \arg \min_{f \in \mathcal{F}_{n,m}} \mathbb{P}_n \left\{ \frac{1}{J} \|\hat{\varphi}_i - \mathbf{B}f(X_i)\|_2^2 \right\}.$$

where $\mathcal{F}_{n,m}$ is the DNN class. We begin with the decomposition

$$\hat{\tau}(x; t) - \tau_0(x; t) = \underbrace{\hat{\tau}(x; t) - \mathbb{E}[\hat{\varphi}_i(t) \mid X_i = x]}_{\text{regression error}} + \underbrace{\mathbb{E}[\hat{\varphi}_i(t) - \varphi_i^0(t) \mid X_i = x]}_{\text{nuisance-induced error}}.$$

The first term reflects the error from nonparametric regression, while the second term captures the impact of estimating nuisance functions. We bound the second term using the following results.

Theorem 3 (Bias decomposition with estimated nuisance functions)

Suppose Assumptions 3, 5, and 6 hold. Let $\hat{\varphi}_i(t)$ and $\varphi_i^0(t)$ denote the estimated and population pseudo-outcomes, respectively, obtained by plugging the cross-fitted and true nuisance functions into (2). Then, it holds that

$$\begin{aligned} \mathbb{E}[\hat{\varphi}_i(t) - \varphi_i^0(t) \mid X_i = x] &= \sum_{w=0}^1 (-1)^{w+1} \left[\frac{\hat{e}^w(x) - e_0^w(x)}{\hat{e}^w(x)} \{ \hat{S}^w(x; t) - S_0^w(x; t) \} \right. \\ &\quad \left. + \frac{e_0^w(x) S_0^w(x; t)}{\hat{e}^w(x) \hat{G}^w(x; t)} \{ G_0^w(x; t) - \hat{G}^w(x; t) \} \right]. \end{aligned} \quad (5)$$

Consequently, it follows that

$$\sup_x |\mathbb{E}[\hat{\varphi}_i(t) - \varphi_i^0(t) \mid X_i = x]| = O_p \left(r_n^{(e)} r_n^{(S)} + r_n^{(G)} \right).$$

In particular, if $\hat{G}^w(x; t) = G_0^w(x; t)$, then we have

$$\sup_x |\mathbb{E}[\hat{\varphi}_i(t) - \varphi_i^0(t) \mid X_i = x]| = O_p \left(r_n^{(e)} r_n^{(S)} \right).$$

We impose smoothness conditions on the CATE function in x and additional regularity conditions on the cross-fitted pseudo-outcomes and the network class, which are required for the consistency and error rate results at a fixed time point.

Definition 1 (Hölder class) Let $\mathcal{X} = [0, 1]^d$ and let $\beta > 0$. Write

$$\beta = q + \alpha, \quad q = \lfloor \beta \rfloor, \quad \alpha \in (0, 1],$$

where $q = \lfloor \beta \rfloor$ is the integer part of β . Let $\nu = (\nu_1, \dots, \nu_d)$, where each ν_j is a nonnegative integer. For $x = (x_1, \dots, x_d)$, define $D_x^\nu = \frac{\partial^{|\nu|_1}}{\partial x_1^{\nu_1} \dots \partial x_d^{\nu_d}}$. For $M > 0$, define $\mathcal{H}^\beta(M)$ as the class of functions $f : \mathcal{X} \rightarrow \mathbb{R}$ such that, for all mixed partial derivatives $D_x^\nu f(x)$ with $\|\nu\|_1 \leq q$, these derivatives exist and are uniformly bounded by M , that is,

$$\max_{\|\nu\|_1 \leq q} \sup_{x \in \mathcal{X}} |D_x^\nu f(x)| \leq M,$$

and moreover,

$$\sup_{x \neq y} \frac{|D_x^\nu f(x) - D_x^\nu f(y)|}{\|x - y\|_2^\alpha} \leq M, \quad \|\nu\|_1 = q.$$

Assumption 7 (Bounded second moments) The cross-fitted $\hat{\varphi}_i$ satisfies $\sup_{1 \leq i \leq n} \mathbb{E}[\hat{\varphi}_i^2(t)] < \infty$.

Assumption 8 (Network approximation capacity) For any fixed $m \geq 1$, smoothness parameter $\beta > 0$, and $g_1, \dots, g_m \in \mathcal{H}^\beta(M)$, the network class $\mathcal{F}_{n,m}$ is rich such that there exists $f^* = (f_1^*, \dots, f_m^*)^\top \in \mathcal{F}_{n,m}$ satisfying

$$\max_{1 \leq k \leq m} \|f_k^* - g_k\|_{L_\infty(P_X)} \leq Cn^{-\beta/(2\beta+d)}.$$

Assumption 9 (Uniform convergence over network) The empirical squared-loss process for the scalar network class $\mathcal{F}_{n,1}$ satisfies

$$\sup_{f \in \mathcal{F}_{n,1}} |\mathbb{P}_n[\{\hat{\varphi}_i(t) - f(X_i)\}^2] - \mathbb{E}[\{\hat{\varphi}_i(t) - f(X_i)\}^2]| = O_p(\rho_n),$$

for some $\rho_n \rightarrow 0$.

Assumption 10 (Controlled optimization error) For fixed $J \geq 1$ and $m \leq J$, let $\hat{f} = (\hat{f}_1, \dots, \hat{f}_m)^\top \in \mathcal{F}_{n,m}$ denote the fitted m -output DNN. The optimization error is controlled in the sense that

$$\mathbb{P}_n \left[\frac{1}{J} \sum_{j=1}^J \left\{ \hat{\varphi}_i(t_j) - \sum_{k=1}^m \hat{f}_k(X_i) B_k(t_j) \right\}^2 \right] \leq \inf_{f \in \mathcal{F}_{n,m}} \mathbb{P}_n \left[\frac{1}{J} \sum_{j=1}^J \left\{ \hat{\varphi}_i(t_j) - \sum_{k=1}^m f_k(X_i) B_k(t_j) \right\}^2 \right] + \delta_{n,J,m}^2$$

where $\delta_{n,J,m}^2 = o_p(1)$

The Hölder class in Definition 1 is a standard smoothness condition in nonparametric regression used to characterize approximation error and minimax rates (Györfi et al., 2006; Tsybakov, 2009). Assumption 7 imposes a bounded second moment condition on the pseudo-outcomes, ensuring well-behaved risk and concentration properties (Chernozhukov et al., 2018a). Assumption 8 requires that the network class $\mathcal{F}_n$ can approximate $\tau_0(\cdot; t)$ at the optimal rate, consistent with approximation results for ReLU networks under Hölder smoothness at $m = 1$ (Schmidt-Hieber, 2020). For $m > 1$, the scalar approximating networks can be combined in parallel into a single m -output ReLU network without changing the approximation rate in n . Assumption 9 ensures uniform convergence of the empirical loss to its population counterpart over $\mathcal{F}_n$, providing control of the estimation error in empirical risk minimization (Bartlett et al., 2017; van der Vaart & Wellner, 1996). Finally, Assumption 10 controls the optimization error arising from approximate minimization of the empirical loss (Schmidt-Hieber, 2020).

The consistency of the final estimator relies on the approximation and estimation properties of DNNs. We first consider the scalar network estimator at a fixed time t . This corresponds to the special case of general basis representation with $J = m = 1$ and $B_1(t) = 1$. We denote $\tau(x; t) = f_1(x)$, $f_1 \in \mathcal{F}_{n,1}$. We now establish consistency and convergence rate for this fixed time estimator. For the fixed-time analysis, we write $\delta_n = \delta_{n,1,1}$ for simplicity.

Theorem 4 (Consistency of the constructed estimator at a fixed t)

Suppose Assumptions 1–9 hold and that the convergence rates of nuisance functions in Assumption 6 satisfy that $r_n^{(e)} r_n^{(S)} + r_n^{(G)} \rightarrow 0$. It follows that

$$\|\hat{\tau}(\cdot; t) - \tau_0(\cdot; t)\|_{L_2(P_X)} \xrightarrow{P} 0.$$

.

Theorem 5 (Rate for the constructed CATE estimator at a fixed t)

Suppose $\tau_0(\cdot; t) \in \mathcal{H}^\beta(M)$ for some $\beta > 0$ and $M > 0$, and that Assumptions 1–10 hold. It follows that

$$\|\hat{\tau}(\cdot; t) - \tau_0(\cdot; t)\|_{L_2(P_X)} = O_p\left(n^{-\beta/(2\beta+d)} + \rho_n^{1/2} + \delta_n + r_n^{(e)} r_n^{(S)} + r_n^{(G)}\right).$$

In particular, if $\rho_n \lesssim n^{-2\beta/(2\beta+d)}$, it holds that

$$\|\hat{\tau}(\cdot; t) - \tau_0(\cdot; t)\|_{L_2(P_X)} = O_p\left(n^{-\beta/(2\beta+d)} + \delta_n + r_n^{(e)} r_n^{(S)} + r_n^{(G)}\right).$$

Remark 1 Cross-fitting ensures that the pseudo-outcome $\hat{\varphi}_i(t)$ is constructed using nuisance estimators that are independent of observation i . This allows us to treat the estimated nuisance functions as fixed when taking conditional expectations, which is critical for establishing the bias decomposition in Theorem 3. In particular, it justifies the expansion $\mathbb{E}\{\hat{\varphi}_i(t) \mid X_i = x\}$, leading to the second-order term $r_n^{(e)} r_n^{(S)}$ arising from the product of estimation errors in the propensity score and outcome model. Without cross-fitting, additional

first-order bias terms of order $r_n^{(e)} + r_n^{(S)}$ may appear due to overfitting. We note that estimation of the censoring model introduces an additional first-order term $r_n^{(G)}$, as reflected in the bias decomposition.

To quantify the gain from joint estimation of $\{\tau_0(x; t_j)\}_{j=1}^J$, we impose additional regularity on the CATE surface in the time dimension. In contrast to the single-time-point setting, where only smoothness in x is required, joint estimation exploits cross-time dependence by modeling $\tau_0(x; t)$ as a function that is jointly smooth in (x, t) . We formalize this through a bivariate Hölder class and prove that it induces a low-complexity representation in the time direction and enables control of the joint approximation error. The following lemma establishes this representation, and the subsequent theorem uses it to characterize the advantage of shared network estimation across time. Without loss of generality, we take $[t_{\min}, t_{\max}] = [0, 1]$ throughout the development.

Definition 2 (Hölder smoothness in (x, t)) Let $\mathcal{X} = [0, 1]^d$ and $\mathcal{T} = [0, 1]$. For $\beta_x > 0$ and $\beta_t > 0$, write

$$\beta_x = q_x + \alpha_x, \quad q_x = \lfloor \beta_x \rfloor, \quad \alpha_x \in (0, 1],$$

and

$$\beta_t = q_t + \alpha_t, \quad q_t = \lfloor \beta_t \rfloor, \quad \alpha_t \in (0, 1].$$

Let $\nu = (\nu_1, \dots, \nu_d)$, where ν_j is a nonnegative integer. For $x = (x_1, \dots, x_d)$, define $D_x^\nu = \frac{\partial^{\|\nu\|_1}}{\partial x_1^{\nu_1} \dots \partial x_d^{\nu_d}}$. For $M > 0$, define $\mathcal{H}^{\beta_x, \beta_t}(M)$ as the class of functions $f : \mathcal{X} \times \mathcal{T} \rightarrow \mathbb{R}$ such that, for all mixed derivatives

$$D_x^\nu \partial_t^m f(x, t), \quad \|\nu\|_1 \leq q_x, \quad m = 0, 1, \dots, q_t,$$

these derivatives exist and are uniformly bounded by M , and moreover satisfy

$$\sup_{t \in \mathcal{T}} \sup_{x \neq y} \frac{|D_x^\nu \partial_t^m f(x, t) - D_x^\nu \partial_t^m f(y, t)|}{\|x - y\|_2^{\alpha_x}} \leq M, \quad \|\nu\|_1 = q_x, \quad m = 0, 1, \dots, q_t,$$

and

$$\sup_{x \in \mathcal{X}} \sup_{s \neq t} \frac{|D_x^\nu \partial_t^{q_t} f(x, t) - D_x^\nu \partial_t^{q_t} f(x, s)|}{|t - s|^{\alpha_t}} \leq M, \quad \|\nu\|_1 \leq q_x.$$

Lemma 1 (Basis expansion under Hölder smoothness) Suppose $\tau_0(\cdot; \cdot) \in \mathcal{H}^{\beta_x, \beta_t}(M)$ in the sense of Definition 2. Let $q_t = \lfloor \beta_t \rfloor$. For any integer $K \geq 1$, there exist $m = (q_t + 1)K$ basis functions $B_1, \dots, B_m$ on $[0, 1]$ and coefficient functions $b_1, \dots, b_m \in \mathcal{H}^{\beta_x}(C_M) : \mathcal{X} \rightarrow \mathbb{R}$ such that

$$\sup_{1 \leq j \leq J} \left\| \tau_0(\cdot; t_j) - \sum_{k=1}^m b_k(\cdot) B_k(t_j) \right\|_{L_2(P_X)}^2 \leq C m^{-2\beta_t},$$

where $C > 0$ and $C_M > 0$ depend only on the Hölder smoothness parameters. Consequently, it holds that

$$\frac{1}{J} \sum_{j=1}^J \left\| \tau_0(\cdot; t_j) - \sum_{k=1}^m b_k(\cdot) B_k(t_j) \right\|_{L_2(P_X)}^2 \leq C m^{-2\beta_t}.$$

Theorem 6 (Error bound for joint estimation over time) *For fixed $J > 1$, let $B_1, \dots, B_m$, with $m \leq J$, be fixed basis functions satisfying the approximation property in Lemma 1, and let $\mathbf{B} = \{B_k(t_j)\}_{1 \leq j \leq J, 1 \leq k \leq m}$ denote the corresponding basis evaluation matrix. The joint estimator $\hat{\tau}(x; t) = \sum_{k=1}^m \hat{f}_k(x) B_k(t)$ is defined by*

$$\hat{f} = \arg \min_{f \in \mathcal{F}_{n,m}} \mathbb{P}_n \left\{ \frac{1}{J} \|\hat{\varphi}_i - \mathbf{B}f(X_i)\|_2^2 \right\}$$

where $\mathcal{F}_{n,m}$ is the class of m -output neural networks with shared hidden layers. Suppose $\tau_0 \in \mathcal{H}^{\beta_x, \beta_t}(M)$ and that the basis stability condition holds:

$$\lambda_{\max} \left(\frac{m}{J} \mathbf{B}^\top \mathbf{B} \right) \leq C_B.$$

Under Assumptions 1–10, then

$$\begin{aligned} & \frac{1}{J} \sum_{j=1}^J \|\hat{\tau}(\cdot; t_j) - \tau_0(\cdot; t_j)\|_{L_2(P_X)}^2 \\ &= O_p \left(m^{-2\beta_t} + n^{-2\beta_x/(2\beta_x+d)} + \sqrt{m} \rho_n + \delta_{n,J,m}^2 + \left\{ r_n^{(e)} r_n^{(S)} + r_n^{(G)} \right\}^2 \right). \end{aligned}$$

Remark 2 If one instead fits J separate DNNs, one for each time point t_j without using smoothness in t , then the empirical process fluctuation is of order $\sqrt{J} \rho_n$. By contrast, the joint estimator exploits the temporal smoothness through a lower-dimensional representation with $m < J$ basis functions and achieves a smaller empirical process fluctuation of order $\sqrt{m} \rho_n$ at the cost of temporal approximation error of order $m^{-2\beta_t}$. Therefore, the joint estimator provides a complexity advantage over fitting J time-specific DNNs separately when the CATE varies smoothly over time and can be well represented using m basis functions. This advantage need not hold when such temporal structure is weak or misspecified.

4 Simulation Study

To evaluate the finite-sample performance of DSL and to empirically verify its double robustness property, we conducted simulation studies under settings where either the survival model or the propensity score model is misspecified. We compared DSL with several commonly used causal machine learning methods, including the M-learner (Bo et al., 2025; Tian et al., 2014), X-learner (Künzel et al., 2019), and R-learner (Nie & Wager, 2021). We also included two doubly robust learners based on nonparametric failure time Bayesian additive regression trees (NFT-BART), namely DRL-NFT-BART and PSDRL-NFT-BART (Yang et al., 2025), and the causal survival forest (Cui et al., 2023).

Across all scenarios, we generated independent datasets consisting of 800 training samples and 400 testing samples. For each subject i , a 30-dimensional

covariate vector X_i was generated independently from a uniform distribution on $[-3, 3]^{30}$.

Case 1: Correct survival and propensity score models In this setting, both the survival model and the propensity score model were correctly specified. The treatment assignment W_i was generated from a logistic regression model,

$$\text{logit}\{P(W_i = 1 \mid X_i)\} = \alpha^\top X_i,$$

where $\text{logit}(p) = \log\{p/(1-p)\}$. The survival time followed a Weibull-Cox proportional hazards model,

$$\lambda_i(t) = \lambda \nu t^{\nu-1} \exp(\beta^\top X_i + W_i \theta^\top X_i),$$

and the censoring time followed an Exponential-Cox model,

$$\lambda_i^C(t) = \lambda^C \exp(\gamma^\top X_i).$$

We set $\lambda = 0.1$, $\nu = 1.2$, and $\lambda^C = 0.2$, so the average censoring rate is approximately 35%. The coefficient vectors $\alpha, \beta, \theta, \gamma \in \mathbb{R}^{30}$ were generated as evenly spaced values on the intervals $[-1, 1]$, $[0.4, -0.2]$, $[0.3, -0.1]$, and $[0.3, -0.3]$, respectively. Additional simulation results under a simplified setting with lower-dimensional covariates and a larger sample size are provided in the appendix for comparison.

Case 2: Correct survival model, misspecified propensity score This setting evaluates robustness to misspecification of the propensity score model. The survival and censoring mechanisms were identical to those in Case 1, while the treatment assignment followed a nonlinear model,

$$\text{logit}\{P(W_i = 1 \mid X_i)\} = 0.5 + 0.45 \tanh(\eta_i/1.2) + 0.25\mathbb{I}(\eta_i > 0) - 0.15\mathbb{I}(\eta_i < 1) + 0.1 \sin(\eta_i^2/3),$$

where $\eta_i = \alpha^\top X_i$ and α is defined as in Case 1. The resulting censoring rate was approximately 40%.

Case 3: Misspecified survival model, correct propensity score This setting evaluates robustness to misspecification of the survival model. The propensity score was generated as in Case 1, while the survival time followed a non-proportional hazards model with time-varying treatment effects,

$$\lambda_i(t) = \lambda_0 \exp\{g(X_i) + W_i \cdot h(X_i) \cdot \sin(v(X_i) t)\},$$

where

$$\begin{aligned} g(X_i) &= \beta^\top X_i, \quad u(X_i) = \gamma^\top X_i, \quad v(X_i) = \theta^\top X_i, \\ h(X_i) &= 0.5 \cos(u(X_i)) + 0.2u(X_i)^2. \end{aligned}$$

The coefficient vectors β, γ, θ were generated as evenly spaced values on the intervals $[-0.3, 0.5]$, $[-0.2, 0.3]$, and $[-0.2, 0.4]$, respectively. We set $\lambda_0 = 0.05$. The censoring mechanism was the same as in Case 1, except that $\lambda^C = 0.05$, yielding a censoring rate of approximately 50%.

In each simulation setting, we evaluated CATE over a grid of J time points, with $J \in \{1, 50, 100\}$. The time grid consisted of J equally spaced points between the 20th and 80th percentiles of the observed survival time distribution.

When $J = 1$, the evaluation time corresponded to the median observed survival time.

We generated 200 independent replicates. For each replicate, DSL was trained on the training sample and evaluated on the testing sample. For replicate $r = 1, \dots, 200$, we computed the bias and mean squared error (MSE) averaged over test subjects and time points,

$$\text{Bias}^{(r)} = \frac{1}{NJ} \sum_{i=1}^n \sum_{j=1}^J \left\{ \hat{\tau}^{(r)}(X_i^{(r)}; t_j^{(r)}) - \tau_0(X_i^{(r)}; t_j^{(r)}) \right\},$$

$$\text{MSE}^{(r)} = \frac{1}{NJ} \sum_{i=1}^n \sum_{j=1}^J \left\{ \hat{\tau}^{(r)}(X_i^{(r)}; t_j^{(r)}) - \tau_0(X_i^{(r)}; t_j^{(r)}) \right\}^2,$$

where $N = 400$ denotes the number of test samples. We report the average bias and MSE across replicates, along with their corresponding 95% Monte Carlo confidence intervals, in Tables 1 and 2.

We employed five-fold cross-fitting to construct the estimated pseudo-outcomes. The nuisance functions for the survival and censoring time was estimated using Cox proportional hazards models with linear predictors, and the treatment assignment was modeled using a logistic regression with linear covariates. The DSL architecture uses the basis-expansion representation described in Section 2.2. Specifically, we use the indicator basis ($B_\ell(t_j) = I(\ell = j)$), $\ell, j = 1, \dots, J$, so that the basis evaluation matrix is ($\mathbf{B} = \mathbf{I}_J$). Under this representation, the coefficient functions $f_1(x), \dots, f_J(x)$ are jointly estimated by a J -output DNN with shared hidden layers, resulting in $\tau(x; t_j) = f_j(x)$ at each evaluation time (t_j). The DNN consists of three fully connected hidden layers with 128, 64, and 32 neurons, respectively, each using ReLU activation functions, and is trained using the Adam optimizer with a learning rate of 0.003. Regularization is achieved through the architectural design of the network, including its depth and width, which implicitly control model complexity. All neural network models are implemented in **PyTorch**. For competing methods, the M-Learner used logistic regression (via **glm**) for propensity score estimation and generalized boosted regression models (GBM; **gbm** package) for outcome regression, with inverse probability of censoring weights estimated using random survival forests. The X-Learner employed survival random forests implemented in **randomForestSRC**, with tuning over terminal node sizes from 1 to 100 and candidate splitting variables starting at 15. The outcome regressions for the X-Learner were also fitted using GBM. The Causal Survival Forest was implemented using the **grf** package with 2000 trees and a minimum node size of 5. Finally, we implemented two doubly robust learners, DRL-NFT-BART and SDRL-NFT-BART (Yang et al., 2025) via **nftbart**. DRL-NFT-BART combines inverse probability of censoring weighting with outcome regression, while PSDRL-NFT-BART uses a propensity-score-weighted pseudo-outcome. Propensity scores were estimated via Super Learner, censoring probabilities via random survival forests, and NFT-BART

was fitted with 1000 burn-in and 2000 posterior samples; CATE was estimated using a regression forest with 2000 trees.

Table 1 summarizes the mean bias and corresponding 95% Monte Carlo confidence intervals across all simulation settings. Overall, the proposed DSL consistently exhibits the smallest bias among competing approaches. When both the survival and treatment models are correctly specified, the confidence intervals for DSL include zero at $J = 1$ and $J = 50$, indicating negligible bias; similar behavior persists when only the survival model is correctly specified, with coverage maintained across all J . In contrast, competing learners, including M-, X-, and R-learners, exhibit substantial and persistent bias that does not diminish with increasing J . Among the doubly robust methods, DRL-NFT-BART shows partial robustness, with bias substantially reduced under treatment model misspecification but remaining non-negligible when the survival model is misspecified. PSDRL-NFT-BART, which relies on propensity-score-weighted pseudo-outcomes, exhibits consistently large bias across all scenarios, indicating limited robustness in finite samples. When only the treatment model is correctly specified, DSL continues to outperform all competitors, achieving markedly smaller bias. For example, at $J = 100$, DSL attains a bias of 0.011 (95% CI: 0.005, 0.016), compared with 0.219 (95% CI: 0.203, 0.234) for the R-learner and 0.236 (95% CI: 0.217, 0.255) for DRL-NFT-BART.

Table 2 reports the corresponding mean squared error (MSE). Across all scenarios, DSL achieves competitive performance, with only a mild increase in MSE as the number of time points J grows. In the fully correctly specified setting, DRL-NFT-BART and the X-learner attain slightly lower MSE than DSL, although the differences are modest. Under nuisance misspecification, DSL demonstrates a clear advantage. When only the survival model is correctly specified, DSL substantially outperforms all competing methods, with MSE increasing only slightly from 0.070 (95% CI: 0.068, 0.072) at $J = 1$ to 0.076 (95% CI: 0.074, 0.077) at $J = 100$, compared to higher MSEs for alternatives such as the X-learner. When the survival model is misspecified, DRL-NFT-BART achieves lower MSE at small J , but this advantage diminishes as J increases, and DSL remains competitive. In contrast, PSDRL-NFT-BART exhibits consistently higher MSE across all settings, consistent with its large bias.

In summary, DSL remains competitive under correct specification, incurring minimal efficiency loss, while offering clear advantages under nuisance misspecification through reduced bias and enhanced robustness, particularly as model complexity increases.

5 Boston Lung Cancer Study

We apply the proposed method to investigate heterogeneous treatment effects of perioperative chemotherapy using data from the Boston Lung Cancer Study (BLCS), a large prospective longitudinal cancer cohort initiated in 1992 (Alvarez et al., 2026; Wang et al., 2023; Zhai et al., 2022). The BLCS has enrolled

Table 1 Bias of conditional average treatment effect estimators across time under nuisance model misspecification

		Case 1: Correct Survival Correct Treatment	Case 2: Correct Survival Wrong Treatment	Case 3: Wrong Survival Correct Treatment
$J = 1$	Deep Survival Learner	0.002(-0.003, 0.007)	0.003(-0.003, 0.010)	0.012(0.005, 0.018)
	M-Learner	0.199(0.193, 0.204)	0.221(0.213, 0.229)	0.305(0.300, 0.310)
	R-Learner	0.068(0.054, 0.082)	0.030(0.011, 0.049)	0.223(0.207, 0.239)
	X-Learner	0.204(0.198, 0.211)	0.221(0.210, 0.232)	0.297(0.289, 0.305)
	DRL-NFT-BART	0.177(0.161, 0.194)	-0.033(-0.067, 0.001)	0.209(0.188, 0.231)
	PSDRL-NFT-BART	0.360(0.354, 0.367)	0.372(0.364, 0.380)	0.386(0.380, 0.391)
	Causal Survival Forest	0.378(0.371, 0.384)	0.392(0.382, 0.402)	0.475(0.468, 0.481)
$J = 50$	Deep Survival Learner	-0.001(-0.006, 0.004)	0.001(-0.006, 0.007)	0.012(0.006, 0.018)
	M-Learner	0.199(0.194, 0.204)	0.221(0.213, 0.229)	0.305(0.299, 0.310)
	R-Learner	0.068(0.054, 0.081)	0.022(0.003, 0.041)	0.220(0.205, 0.234)
	X-Learner	0.206(0.200, 0.213)	0.226(0.216, 0.236)	0.300(0.292, 0.308)
	DRL-NFT-BART	0.167(0.147, 0.187)	-0.024(-0.056, 0.009)	0.238(0.221, 0.255)
	PSDRL-NFT-BART	0.360(0.354, 0.366)	0.373(0.364, 0.381)	0.386(0.380, 0.392)
	Causal Survival Forest	0.378(0.372, 0.384)	0.393(0.383, 0.403)	0.474(0.468, 0.481)
$J = 100$	Deep Survival Learner	-0.007(-0.012, -0.002)	-0.001(-0.008, 0.005)	0.011(0.005, 0.016)
	M-Learner	0.198(0.193, 0.204)	0.221(0.213, 0.229)	0.305(0.299, 0.310)
	R-Learner	0.073(0.061, 0.086)	0.028(0.010, 0.047)	0.219(0.203, 0.234)
	X-Learner	0.205(0.198, 0.212)	0.229(0.219, 0.240)	0.295(0.286, 0.303)
	DRL-NFT-BART	0.158(0.136, 0.180)	-0.037(-0.071, -0.002)	0.236(0.217, 0.255)
	PSDRL-NFT-BART	0.360(0.354, 0.367)	0.373(0.364, 0.381)	0.384(0.378, 0.389)
	Causal Survival Forest	0.378(0.372, 0.384)	0.393(0.384, 0.403)	0.474(0.468, 0.481)

¹ Mean bias and 95% Monte Carlo confidence intervals are computed based on 200 simulation replicates.

² Conditional average treatment effects are evaluated at $J \in \{1, 50, 100\}$ time points, equally spaced between the 20th and 80th percentiles of the observed survival time distribution.

³ Columns correspond to combinations of correctly and incorrectly specified survival and treatment assignment models; the censoring model is correctly specified in all scenarios.

over 12,000 patients treated at Massachusetts General Hospital and includes detailed clinical, pathological, and follow-up information, including treatment history and survival outcomes.

We evaluate the effect of perioperative chemotherapy on overall survival among patients with non-small cell lung cancer (NSCLC). The primary outcome is the time from diagnosis to death or censoring. Because perioperative chemotherapy is primarily administered to patients with early-stage disease, we restrict the analysis to patients with stage II or III NSCLC who underwent surgical resection. Within this subset, we compare patients who received perioperative chemotherapy with those who did not. Patients with missing information on cancer histology, stage, treatment status, or diagnosis date are excluded. After applying these inclusion and exclusion criteria, a total of 784 patients from the BLCS cohort are included in the analysis.

Based on the descriptive statistics in Table 3, the study cohort consists of 784 patients, among whom 521 (66%) deaths were observed during follow-up. The median survival time for the overall population is 5.4 years (95% CI: 4.7–6.1). Patients who received surgery plus perioperative chemotherapy are defined as the treatment group, whereas those who underwent surgery alone constitute the control group. A total of 214 patients (27%) received

Table 2 Mean squared error of conditional average treatment estimators across time under nuisance model misspecification

		Case 1: Correct Survival Correct Treatment	Case 2: Correct Survival Wrong Treatment	Case 3: Wrong Survival Correct Treatment
$J = 1$	Deep Survival Learner	0.077(0.075, 0.080)	0.070(0.068, 0.072)	0.104(0.101, 0.106)
	M-Learner	0.114(0.110, 0.117)	0.122(0.117, 0.127)	0.190(0.185, 0.194)
	R-Learner	0.174(0.168, 0.180)	0.235(0.227, 0.244)	0.260(0.251, 0.269)
	X-Learner	0.081(0.078, 0.084)	0.093(0.089, 0.098)	0.154(0.150, 0.159)
	DRL-NFT-BART	0.070(0.064, 0.075)	0.133(0.113, 0.153)	0.087(0.077, 0.097)
	PSDRL-NFT-BART	0.142(0.138, 0.146)	0.152(0.147, 0.157)	0.152(0.148, 0.156)
	Causal Survival Forest	0.156(0.151, 0.160)	0.170(0.163, 0.176)	0.229(0.222, 0.235)
$J = 50$	Deep Survival Learner	0.084(0.082, 0.086)	0.074(0.072, 0.076)	0.112(0.109, 0.115)
	M-Learner	0.114(0.110, 0.117)	0.122(0.117, 0.127)	0.189(0.185, 0.194)
	R-Learner	0.174(0.168, 0.180)	0.237(0.229, 0.245)	0.261(0.251, 0.270)
	X-Learner	0.083(0.080, 0.085)	0.094(0.090, 0.098)	0.157(0.152, 0.161)
	DRL-NFT-BART	0.077(0.062, 0.091)	0.126(0.108, 0.143)	0.087(0.081, 0.094)
	PSDRL-NFT-BART	0.141(0.137, 0.146)	0.152(0.147, 0.158)	0.152(0.148, 0.156)
	Causal Survival Forest	0.156(0.152, 0.160)	0.170(0.164, 0.177)	0.228(0.222, 0.235)
$J = 100$	Deep Survival Learner	0.084(0.082, 0.086)	0.076(0.074, 0.077)	0.114(0.111, 0.117)
	M-Learner	0.114(0.110, 0.117)	0.123(0.118, 0.128)	0.190(0.185, 0.194)
	R-Learner	0.173(0.167, 0.179)	0.233(0.225, 0.241)	0.260(0.252, 0.269)
	X-Learner	0.082(0.080, 0.085)	0.096(0.091, 0.100)	0.154(0.150, 0.158)
	DRL-NFT-BART	0.082(0.066, 0.098)	0.139(0.117, 0.161)	0.088(0.082, 0.094)
	PSDRL-NFT-BART	0.142(0.138, 0.146)	0.152(0.147, 0.158)	0.150(0.146, 0.155)
	Causal Survival Forest	0.156(0.152, 0.161)	0.171(0.164, 0.177)	0.228(0.222, 0.234)

¹ Mean squared error and 95% Monte Carlo confidence intervals are computed based on 200 simulation replicates.

² Conditional average treatment effects are evaluated at $J \in \{1, 50, 100\}$ time points, equally spaced between the 20th and 80th percentiles of the observed survival time distribution.

³ Columns correspond to combinations of correctly and incorrectly specified survival and treatment assignment models; the censoring model is correctly specified in all scenarios.

perioperative chemotherapy. Baseline characteristics differ between the two groups. Patients in the treatment group are younger (median age 64 years, IQR: 58–71) than those in the control group (median age 68 years, IQR: 60–75), and a higher proportion of treated patients are male (54% vs. 48%). Survival outcomes appear more favorable in the treatment group: 110 deaths (51%) were observed among treated patients compared with 411 deaths (72%) in the control group. The median survival time is 6.5 years (95% CI: 5.3–8.6) for the treatment group, versus 4.8 years (95% CI: 4.1–5.8) for the control group. Missing values in age, BMI, and smoking intensity are imputed using their respective sample means (mean imputation). Kaplan–Meier curves (Figure 1) further illustrate differences in survival probabilities between patients who received perioperative chemotherapy and those who did not.

We estimate CATEs at 10 prespecified time points ranging from 1 to 23 years after surgery. At each time point t , the CATE is defined as the difference in survival probabilities at time t between surgery plus perioperative chemotherapy and surgery alone, conditional on pre-treatment covariates. In our implementation, DSL employs an indicator basis over the 10 prespecified evaluation time points, such that $B_\ell(t_j) = I(\ell = j)$ and the corresponding basis evaluation matrix is $\mathbf{B} = \mathbf{I}_{10}$. Thus, the 10 coefficient functions are jointly

Table 3 Clinical Characteristics of Patients from the Boston Lung Cancer Cohort

Characteristic	Overall N = 784 ¹	Surgery N = 570 ¹	Surgery + Chemotherapy N = 214 ¹
Death	521 (66%)	411 (72%)	110 (51%)
Median Survival Time (years)	5.4 (4.7, 6.1)	4.8 (4.1, 5.8)	6.5 (5.3, 8.6)
Age at diagnosis (yrs)	67 (59, 74)	68 (60, 75)	64 (58, 71)
Unknown	19	12	7
BMI (kg/m ²)	25.8 (22.9, 29.5)	25.8 (22.8, 29.6)	26.3 (23.5, 29.5)
Unknown	60	53	7
Tumor Stage			
II	445 (57%)	333 (58%)	112 (52%)
III	339 (43%)	237 (42%)	102 (48%)
Gender			
Female	394 (50%)	296 (52%)	98 (46%)
Male	390 (50%)	274 (48%)	116 (54%)
Race			
White	746 (95%)	547 (96%)	199 (93%)
Asian	15 (1.9%)	7 (1.2%)	8 (3.7%)
Black	8 (1.0%)	5 (0.9%)	3 (1.4%)
Other	15 (1.9%)	11 (1.9%)	4 (1.9%)
Smoking Status			
Current Smoker	213 (27%)	166 (29%)	47 (22%)
Former Smoker	490 (62%)	353 (62%)	137 (64%)
Never Smoker	81 (10%)	51 (8.9%)	30 (14%)

¹ Median (IQR); n (%)

estimated by a 10-output DNN with shared hidden layers. The DNN consists of three-layer feedforward network with hidden layers of sizes 64, 32, and 16, respectively, and uses ReLU activation functions.

Figure 2 presents the estimated CATEs across key baseline covariates, including gender, tumor stage, age, body mass index (BMI), and smoking intensity. Across all subgroups, perioperative chemotherapy is associated with positive treatment effects over the time horizon considered, corresponding to higher survival probabilities relative to surgery alone. The magnitude of the estimated CATE increases during the first ten years following surgery and attenuates thereafter.

Gender-specific results (Figure 2(a)) indicate larger treatment effects for female patients. Holding other covariates at their mean or mode, the estimated increases in 5-, 10-, and 20-year survival probabilities are 3.9%, 9.7%, and 6.8% for female patients, compared with 2.7%, -2.8%, and -0.5% for male patients. Substantial heterogeneity is also observed by tumor stage (Figure 2(b)). For patients with stage II NSCLC, perioperative chemotherapy is associated with increases of 3.9%, 9.7%, and 6.8% in 5-, 10-, and 20-year survival probabilities,




Fig. 1 Kaplan–Meier survival curves by treatment group.

respectively. In contrast, the estimated effects for stage III patients are smaller, with corresponding increases of 3.3%, 3.7%, and 3.3%.

Age-specific effects (Figure 2(c)) are relatively stable. Holding other covariates at their mean or mode, the estimated increases in 5-, 10-, and 20-year survival probabilities are 3.9%, 10.1%, and 7.1% for patients aged 65 years, and 3.6%, 6.7%, and 5.1% for patients aged 75 years. Figure 2(d) suggests that patients with higher BMI experience larger treatment effects: patients with BMI = 32 show increases of 4.0%, 10.1%, and 7.1% at 5, 10, and 20 years, compared with 3.8%, 8.1%, and 5.9% for patients with BMI = 22. Smoking intensity also modifies treatment effects (Figure 2(e)). Patients smoking 10 cigarettes per day exhibit larger survival gains (4.3%, 13.2%, and 8.9%) than those smoking 30 cigarettes per day (3.7%, 7.6%, and 5.6%). Overall, perioperative chemotherapy is associated with meaningful survival benefits for patients with early-stage NSCLC, with effects concentrated in the first ten years following surgery and attenuating thereafter. This temporal pattern is consistent with long-term follow-up from the International Adjuvant Lung Cancer Trial, which reported a significant interaction between treatment and follow-up time, with benefits diminishing beyond five years (Arriagada et al., 2010). The observed heterogeneity by gender and tumor stage aligns with prior studies reporting greater benefits among female patients (Leiter et al., 2022; Sandler et al., 2016) and stronger effects in stage II disease relative to more advanced stages (Douillard et al., 2006). The findings on smoking intensity are also consistent with prior evidence linking lower smoking exposure to improved chemotherapy outcomes (Zhang et al., 2008). Finally, patients with higher BMI appear to benefit more from perioperative chemotherapy, a pattern consistent with the reported obesity paradox in lung cancer (Sakin et al., 2021). These results underscore the importance of accounting for treatment effect heterogeneity and equity in evaluating treatment effectiveness (Arriagada

et al., 2010; Douillard et al., 2006; Leiter et al., 2022; Sandler et al., 2016; Zhang et al., 2008).




Fig. 2 Estimated conditional average treatment effects across key baseline covariates. Each panel displays the estimated CATE as a function of one covariate of interest (gender, tumor stage, age, body mass index, or smoking intensity), with all remaining covariates fixed at their sample means or modal values.

6 Discussion

We propose the Deep Survival Learner (DSL), a causal deep learning framework for estimating heterogeneous treatment effects with right-censored survival outcomes. The key contribution is the integration of a doubly robust pseudo-outcome for time-specific conditional average treatment effects with joint learning across time via shared representations. This formulation directly targets treatment effect trajectories over clinically relevant follow-up intervals and improves efficiency relative to pointwise approaches that ignore temporal dependence.

The method relies on standard identification conditions, including unconfoundedness and appropriate handling of censoring. While the doubly robust construction mitigates misspecification of either the outcome or treatment model, correct specification of the censoring mechanism remains essential. This highlights the need for extensions that relax noninformative censoring, for example through joint modeling, sensitivity analysis, or alternative identification strategies.

From a statistical perspective, DSL performs joint estimation of a vector-valued target over a time grid, trading bias and variance through shared representations. This raises questions on optimal grid design, regularization, and

approximation error in high-dimensional settings. Extensions to competing risks would further broaden the framework to multivariate event processes and require redefining target functionals.

Finally, our current analysis focuses on point estimation. Developing valid uncertainty quantification for time-indexed CATEs under censoring remains an important open problem, requiring new theoretical tools that account for dependence across time and the complexity of deep learners.

Acknowledgements

This work was supported by NIH grants 2R01CA249096-05 and R01CA269398. We are grateful to the Editor, the Associate Editor, and the two referees for their insightful reviews, which substantially improved the article. We also thank our long-term collaborator, Dr. David Christiani, for providing the BLCS data and for his many insightful clinical suggestions.

References

- Alvarez, A. A. R., Sun, Y., Li, Y., & Christiani, D. C. (2026). Effects of sex on mortality in patients with lung cancer: A multiple mediation analysis of the Boston Lung Cancer Study. *Clinical Lung Cancer*, 27(2), 201–209.e3.
- Andersen, P. K., & Gill, R. D. (1982). Cox’s regression model for counting processes: A large sample study. *Annals of Statistics*, 10(4), 1100–1120.
- Arriagada, R., Dunant, A., Pignon, J.-P., Bergman, B., Chabowski, M., Grunewald, D., Kozłowski, M., Le Péchoux, C., Pirker, R., Pinel, M.-I. S., et al. (2010). Long-term results of the international adjuvant lung cancer trial evaluating adjuvant cisplatin-based chemotherapy in resected lung cancer. *Journal of Clinical Oncology*, 28(1), 35–42.
- Bartlett, P. L., Foster, D. J., & Telgarsky, M. (2017). Spectrally-normalized margin bounds for neural networks. *Advances in Neural Information Processing Systems*, 30.
- Bo, N., Jeong, J.-H., Forno, E., & Ding, Y. (2025). Evaluating meta-learners to analyze treatment heterogeneity in survival data: Application to electronic health records of pediatric asthma care in covid-19 pandemic. *Statistics in Medicine*, 44(3-4), e10333.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018a). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1), C1–C68. <https://doi.org/10.1111/ectj.12097>
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018b). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1), C1–C68.

- Cui, Y., Kosorok, M. R., Sverdrup, E., Wager, S., & Zhu, R. (2023). Estimating heterogeneous treatment effects with right-censored data via causal survival forests. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(2), 179–211.
- Curth, A., Lee, C., & van der Schaar, M. (2021). Survite: Learning heterogeneous treatment effects from time-to-event data. *Advances in Neural Information Processing Systems*, 34, 26740–26753.
- Dandl, S., Bender, A., & Hothorn, T. (2024). Heterogeneous treatment effect estimation for observational data using model-based forests. *Statistical Methods in Medical Research*, 33(3), 392–413.
- de Boor, C. (2001). *A practical guide to splines* (Revised, Vol. 27). Springer-Verlag. <https://doi.org/10.1007/978-1-4612-6333-3>
- Dharmarajan, S. H., Bragg-Gresham, J. L., Morgenstern, H., Gillespie, B. W., Li, Y., Powe, N. R., Tuot, D. S., Banerjee, T., Ríos Burrows, N., Rolka, D. B., Saydah, S. H., & Saran, R. (2017). State-level awareness of chronic kidney disease in the US. *American Journal of Preventive Medicine*, 53(3), 300–307. <https://doi.org/10.1016/j.amepre.2017.03.023>
- Douillard, J.-Y., Rosell, R., De Lena, M., Carpagnano, F., Ramlau, R., González-Larriba, J. L., Grodzki, T., Pereira, J. R., Le Groumellec, A., Lorusso, V., et al. (2006). Adjuvant vinorelbine plus cisplatin versus observation in patients with completely resected stage ib–iiia non-small-cell lung cancer (adjuvant navelbine international trialist association [anita]): A randomised controlled trial. *The Lancet Oncology*, 7(9), 719–727.
- Györfi, L., Kohler, M., Krzyżak, A., & Walk, H. (2006). *A distribution-free theory of nonparametric regression*. Springer.
- Haddad, R., Crum, C., Chen, Z., Krane, J., Posner, M., Li, Y., & Burk, R. (2008). HPV16 transmission between a couple with HPV-related head and neck cancer. *Oral Oncology*, 44(8), 812–815. <https://doi.org/10.1016/j.oraloncology.2007.09.004>
- Horvitz, D. G., & Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260), 663–685.
- Hu, L. (2024). A new method for clustered survival data: Estimation of treatment effect heterogeneity and variable selection. *Biometrical Journal*, 66(1), 2200178.
- Imbens, G. W., & Rubin, D. B. (2015). *Causal inference for statistics, social, and biomedical sciences: An introduction*. Cambridge University Press.
- Imbens, G. W., & Rubin, D. B. (2016). *Causal inference for statistics, social, and biomedical sciences: An introduction*. Taylor & Francis.
- Kennedy, E. H. (2020). Towards optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497*.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., & Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National academy of Sciences*, 116(10), 4156–4165.

- Leiter, A., Kong, C. Y., Gould, M. K., Kale, M. S., Veluswamy, R. R., Smith, C. B., Mhango, G., Huang, B. Z., Wisnivesky, J. P., & Sigel, K. (2022). The benefits and harms of adjuvant chemotherapy for non-small cell lung cancer in patients with major comorbidities: A simulation study. *Plos one*, 17(11), e0263911.
- Nie, X., & Wager, S. (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2), 299–319.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55.
- Sakin, A., Sahin, S., Mustafa Atci, M., Yasar, N., Demir, C., Geredeli, C., Sakin, A., & Cihan, S. (2021). The effect of body mass index on treatment outcomes in patients with metastatic non-small cell lung cancer treated with platinum-based therapy. *Nutrition and Cancer*, 73(8), 1411–1418.
- Sandler, B. J., Wang, Z., Hancock, J. G., Boffa, D. J., Detterbeck, F. C., & Kim, A. W. (2016). Gender, age, and comorbidity status predict improved survival with adjuvant chemotherapy following lobectomy for non-small cell lung cancers larger than 4 cm. *Annals of surgical oncology*, 23, 638–645.
- Schmidt-Hieber, J. (2020). Deep relu network approximation of functions on a manifold. *Journal of Machine Learning Research*, 21(52), 1–26.
- Schumaker, L. L. (2007). *Spline functions: Basic theory* (3rd). Cambridge University Press. <https://doi.org/10.1017/CBO9780511618994>
- Tian, L., Alizadeh, A. A., Gentles, A. J., & Tibshirani, R. (2014). A simple method for estimating interactions between a treatment and a large number of covariates. *Journal of the American Statistical Association*, 109(508), 1517–1532.
- Tsiatis, A. A. (2006). *Semiparametric theory and missing data*. Springer.
- Tsybakov, A. B. (2009). *Introduction to nonparametric estimation*. Springer. <https://doi.org/10.1007/b13794>
- van der Vaart, A. W. (1998). *Asymptotic statistics*. Cambridge University Press.
- van der Vaart, A. W., & Wellner, J. A. (1996). *Weak convergence and empirical processes*. Springer.
- Wang, X., Romero-Gutierrez, C. W., Kothari, J., Shafer, A., Li, Y., & Christiani, D. C. (2023). Prediagnosis smoking cessation and overall survival among patients with non-small cell lung cancer. *JAMA Network Open*, 6(5), e2311966.
- Xu, Y., Ignatiadis, N., Sverdrup, E., Fleming, S., Wager, S., & Shah, N. (2023). Treatment heterogeneity with survival outcomes. In *Handbook of matching and weighting adjustments for causal inference* (pp. 445–482). Chapman; Hall/CRC.
- Yang, Y., Hu, W., Liao, Z., & Chen, F. (2025). Doubly robust estimators for heterogeneous treatment effects in heteroskedastic survival data. *Statistics in Medicine*, 44(23-24), e70301.

-
- Zhai, T., Li, Y., Brown, R., Lanuti, M., Gainor, J. F., & Christiani, D. C. (2022). Spirometry at diagnosis and overall survival in non-small cell lung cancer patients. *Cancer Medicine*, 11(24), 4796–4805.
- Zhang, Z., Xu, F., Wang, S., Li, N., & Wang, C. (2008). Influence of smoking on histologic type and the efficacy of adjuvant chemotherapy in resected non-small cell lung cancer. *Lung Cancer*, 60(3), 434–440.

**Supplementary Materials for
“estimating heterogeneous treatment effects with survival outcomes
via a deep survival learner”**

Proofs of Theorems

Proof of Theorem 1: It suffices to show that, for each $w \in \{0, 1\}$,

$$\mathbb{E}\{\pi_i^w(t) \mid X_i = x\} = P(T_i^w > t \mid X_i = x),$$

since this immediately implies

$$\mathbb{E}\{\varphi_i(t) \mid X_i = x\} = P(T_i^1 > t \mid X_i = x) - P(T_i^0 > t \mid X_i = x) = \tau_0(x; t).$$

We prove the result for $w = 1$; the case $w = 0$ follows by symmetry.

Recall that

$$\pi_i^1(t) = S^1(x; t) + \frac{\mathbb{I}(W_i = 1)}{e^1(x)} \left\{ \frac{\mathbb{I}(U_i > t)}{G^1(x; t)} - S^1(x; t) \right\},$$

where $S^1(x; t) = P(T_i > t \mid X_i = x, W_i = 1)$, $e^1(x) = P(W_i = 1 \mid X_i = x)$, and $G^1(x; t) = P(\mathcal{C}_i > t \mid X_i = x, W_i = 1)$.

Taking conditional expectation given $X_i = x$ yields

$$\mathbb{E}\{\pi_i^1(t) \mid X_i = x\} = S^1(x; t) + \mathbb{E}\left[\frac{\mathbb{I}(W_i = 1)}{e^1(x)} \left\{ \frac{\mathbb{I}(U_i > t)}{G^1(x; t)} - S^1(x; t) \right\} \middle| X_i = x \right].$$

By the law of iterated expectations, the second term can be written as

$$\mathbb{E}\left[\mathbb{E}\left\{ \frac{\mathbb{I}(W_i = 1)}{e^1(x)} \left(\frac{\mathbb{I}(U_i > t)}{G^1(x; t)} - S^1(x; t) \right) \middle| X_i = x, W_i \right\} \middle| X_i = x \right].$$

Conditioning on W_i , and noting that $\mathbb{I}(W_i = 1)$ is deterministic given W_i , we obtain

$$= \mathbb{E}\left[\mathbb{I}(W_i = 1) \frac{1}{e^1(x)} \left\{ \mathbb{E}\left[\frac{\mathbb{I}(U_i > t)}{G^1(x; t)} \middle| X_i = x, W_i = 1 \right] - S^1(x; t) \right\} \middle| X_i = x \right].$$

Using the definition $U_i = \min(T_i, \mathcal{C}_i)$, we have $\mathbb{I}(U_i > t) = \mathbb{I}(T_i > t, \mathcal{C}_i > t)$. Under Assumption 4, $\mathcal{C}_i \perp T_i^w \mid (X_i, W_i = w)$ for $w \in \{0, 1\}$. Under Assumption 1, when conditioning on $W_i = w$ we have $T_i = T_i^w$, which implies $\mathcal{C}_i \perp T_i \mid (X_i, W_i = w)$, and in particular, $T_i \perp \mathcal{C}_i \mid (X_i, W_i = 1)$. This implies

$$\mathbb{E}[\mathbb{I}(T_i > t, \mathcal{C}_i > t) \mid X_i = x, W_i = 1] = P(T_i > t \mid X_i = x, W_i = 1)P(\mathcal{C}_i > t \mid X_i = x, W_i = 1).$$

Therefore,

$$\mathbb{E}\left[\frac{\mathbb{I}(U_i > t)}{G^1(x; t)} \middle| X_i = x, W_i = 1 \right] = \frac{P(T_i > t \mid X_i = x, W_i = 1)P(\mathcal{C}_i > t \mid X_i = x, W_i = 1)}{P(\mathcal{C}_i > t \mid X_i = x, W_i = 1)} = S^1(x; t).$$

Substituting back yields

$$\begin{aligned}\mathbb{E}\{\pi_i^1(t) \mid X_i = x\} &= S^1(x; t) + \mathbb{E}\left[\mathbb{I}(W_i = 1) \frac{1}{e^1(x)} \{S^1(x; t) - S^1(x; t)\} \mid X_i = x\right] \\ &= S^1(x; t).\end{aligned}$$

Finally, by Assumptions 1 and 2,

$$S^1(x; t) = P(T_i > t \mid X_i = x, W_i = 1) = P(T_i^1 > t \mid X_i = x).$$

This completes the proof.

Proof of Theorem 2: It suffices to show that for each $w \in \{0, 1\}$,

$$\mathbb{E}\{\pi_i^w(t) \mid X_i = x\} = P(T_i^w > t \mid X_i = x),$$

because then $\mathbb{E}\{\varphi_i(t) \mid X_i = x\} = P(T_i^1 > t \mid X_i = x) - P(T_i^0 > t \mid X_i = x) = \tau_0(x; t)$. We prove the claim for $w = 1$; the case $w = 0$ follows by symmetry.

Recall the definition

$$\pi_i^1(t) = S^1(X_i; t) + \frac{\mathbb{I}(W_i = 1)}{e^1(X_i)} \left\{ \frac{Y_i(t)}{G^1(X_i; t)} - S^1(X_i; t) \right\},$$

where $Y_i(t) = \mathbb{I}(U_i > t) = \mathbb{I}(T_i > t, C_i > t)$, and where the nuisance functions $S^1(x; t)$, $G^1(x; t)$ and $e^1(x)$ are intended to approximate

$$S_0^1(x; t) = P(T_i > t \mid X_i = x, W_i = 1), \quad G_0^1(x; t) = P(C_i > t \mid X_i = x, W_i = 1),$$

$$e_0^1(x) = P(W_i = 1 \mid X_i = x).$$

Throughout, Assumptions 1–5 ensure that these conditional probabilities are well-defined and that $G_0^1(x; t) > 0$.

Case A: the survival model is correctly specified. Assume $S^1(x; t) = S_0^1(x; t)$ and the censoring model is correctly specified, i.e., $G^1(x; t) = G_0^1(x; t)$. Then, regardless of whether the propensity score model $e^1(\cdot)$ is correct, the argument in Theorem 1 applies verbatim and yields

$$\mathbb{E}\{\pi_i^1(t) \mid X_i = x\} = S_0^1(x; t) = P(T_i^1 > t \mid X_i = x),$$

where the last equality follows from Assumptions 1 and 2. For completeness, we provide the key step:

$$\mathbb{E}\left[\frac{Y_i(t)}{G_0^1(x; t)} \mid X_i = x, W_i = 1\right] = \frac{P(T_i > t, C_i > t \mid X_i = x, W_i = 1)}{P(C_i > t \mid X_i = x, W_i = 1)} = P(T_i > t \mid X_i = x, W_i = 1),$$

where the second equality uses Assumption 4 (noninformative censoring), i.e., $T_i \perp C_i \mid (X_i, W_i = 1)$. This establishes unbiasedness in Case A.

Case B: the propensity score model is correctly specified. Now assume the censoring model is correctly specified, $G^1(x; t) = G_0^1(x; t)$, and the propensity score is correctly specified, $e^1(x) = e_0^1(x) = P(W_i = 1 \mid X_i = x)$,

while allowing the survival model $S^1(x; t)$ to be misspecified. We show that $\mathbb{E}\{\pi_i^1(t) \mid X_i = x\} = S_0^1(x; t)$.

Condition on $X_i = x$. Using the definition of $\pi_i^1(t)$,

$$\mathbb{E}\{\pi_i^1(t) \mid X_i = x\} = S^1(x; t) + \mathbb{E}\left[\frac{\mathbb{I}(W_i = 1)}{e_0^1(x)} \left\{ \frac{Y_i(t)}{G_0^1(x; t)} - S^1(x; t) \right\} \mid X_i = x\right].$$

Apply the law of iterated expectations conditioning on (X_i, W_i) :

$$\begin{aligned} & \mathbb{E}\left[\frac{\mathbb{I}(W_i = 1)}{e_0^1(x)} \left\{ \frac{Y_i(t)}{G_0^1(x; t)} - S^1(x; t) \right\} \mid X_i = x\right] \\ &= \mathbb{E}\left[\mathbb{E}\left\{\frac{\mathbb{I}(W_i = 1)}{e_0^1(x)} \left(\frac{Y_i(t)}{G_0^1(x; t)} - S^1(x; t) \right) \mid X_i = x, W_i\right\} \mid X_i = x\right]. \end{aligned}$$

Inside the inner expectation, $\mathbb{I}(W_i = 1)$ is fixed given W_i . Hence

$$= \mathbb{E}\left[\frac{\mathbb{I}(W_i = 1)}{e_0^1(x)} \left\{ \mathbb{E}\left[\frac{Y_i(t)}{G_0^1(x; t)} \mid X_i = x, W_i = 1\right] - S^1(x; t) \right\} \mid X_i = x\right].$$

We next compute the conditional mean term. Since $Y_i(t) = \mathbb{I}(T_i > t, \mathcal{C}_i > t)$ and $G_0^1(x; t) = P(\mathcal{C}_i > t \mid X_i = x, W_i = 1)$,

$$\mathbb{E}\left[\frac{Y_i(t)}{G_0^1(x; t)} \mid X_i = x, W_i = 1\right] = \frac{\mathbb{E}[\mathbb{I}(T_i > t, \mathcal{C}_i > t) \mid X_i = x, W_i = 1]}{P(\mathcal{C}_i > t \mid X_i = x, W_i = 1)}.$$

By Assumption 4 and a similar argument in the Proof of Theorem 1, $T_i \perp \mathcal{C}_i \mid (X_i, W_i = 1)$, so

$$P(T_i > t, \mathcal{C}_i > t \mid X_i = x, W_i = 1) = P(T_i > t \mid X_i = x, W_i = 1) P(\mathcal{C}_i > t \mid X_i = x, W_i = 1).$$

Therefore,

$$\mathbb{E}\left[\frac{Y_i(t)}{G_0^1(x; t)} \mid X_i = x, W_i = 1\right] = P(T_i > t \mid X_i = x, W_i = 1) = S_0^1(x; t).$$

Substituting this back,

$$\begin{aligned} \mathbb{E}\{\pi_i^1(t) \mid X_i = x\} &= S^1(x; t) + \mathbb{E}\left[\frac{\mathbb{I}(W_i = 1)}{e_0^1(x)} \{S_0^1(x; t) - S^1(x; t)\} \mid X_i = x\right] \\ &= S^1(x; t) + \{S_0^1(x; t) - S^1(x; t)\} \mathbb{E}\left[\frac{\mathbb{I}(W_i = 1)}{e_0^1(x)} \mid X_i = x\right]. \end{aligned}$$

Because $e_0^1(x) = P(W_i = 1 \mid X_i = x)$, we have

$$\mathbb{E}\left[\frac{\mathbb{I}(W_i = 1)}{e_0^1(x)} \mid X_i = x\right] = \frac{\mathbb{E}\{\mathbb{I}(W_i = 1) \mid X_i = x\}}{e_0^1(x)} = \frac{P(W_i = 1 \mid X_i = x)}{P(W_i = 1 \mid X_i = x)} = 1.$$

This identity follows from correct specification of the propensity score and Assumption 3, which guarantees $0 < e_0^1(x) < 1$. Hence,

$$\mathbb{E}\{\pi_i^1(t) \mid X_i = x\} = S_0^1(x; t) = P(T_i > t \mid X_i = x, W_i = 1).$$

Finally, by Assumptions 1 and 2,

$$P(T_i > t \mid X_i = x, W_i = 1) = P(T_i^1 > t \mid X_i = x),$$

so $\mathbb{E}\{\pi_i^1(t) \mid X_i = x\} = P(T_i^1 > t \mid X_i = x)$, completing the proof for $w = 1$.

The proof for $w = 0$ is identical after replacing 1 by 0 throughout.

Proof of Theorem 3: We first analyze the treatment-specific component $\hat{\pi}_i^w(t)$ for fixed $w \in \{0, 1\}$. As defined in the notation section, conditioning on $X_i = x$, and using cross-fitting so that $\hat{S}^w(x; t)$, $\hat{e}^w(x)$, and $\hat{G}^w(x; t)$ are fixed, we take expectation with respect to $T_i, \mathcal{C}_i$:

$$\mathbb{E}\{\hat{\pi}_i^w(t) \mid X_i = x\} = \hat{S}^w(x; t) + \frac{1}{\hat{e}^w(x)} \mathbb{E} \left[\mathbb{I}(W_i = w) \left\{ \frac{Y_i(t)}{\hat{G}^w(x; t)} - \hat{S}^w(x; t) \right\} \middle| X_i = x \right].$$

By iterated expectation,

$$\begin{aligned} & \mathbb{E} \left[\mathbb{I}(W_i = w) \left\{ \frac{Y_i(t)}{\hat{G}^w(x; t)} - \hat{S}^w(x; t) \right\} \middle| X_i = x \right] \\ &= e_0^w(x) \left\{ \mathbb{E} \left[\frac{Y_i(t)}{\hat{G}^w(x; t)} \middle| X_i = x, W_i = w \right] - \hat{S}^w(x; t) \right\}. \end{aligned}$$

Since $Y_i(t) = \mathbb{I}(T_i > t, \mathcal{C}_i > t)$, by noninformative censoring,

$$\begin{aligned} \mathbb{E} \left[\frac{Y_i(t)}{\hat{G}^w(x; t)} \middle| X_i = x, W_i = w \right] &= \frac{P(T_i > t, \mathcal{C}_i > t \mid X_i = x, W_i = w)}{\hat{G}^w(x; t)} \\ &= \frac{P(T_i > t \mid X_i = x, W_i = w) P(\mathcal{C}_i > t \mid X_i = x, W_i = w)}{\hat{G}^w(x; t)} \\ &= \frac{S_0^w(x; t) G_0^w(x; t)}{\hat{G}^w(x; t)}. \end{aligned}$$

Hence,

$$\mathbb{E}\{\hat{\pi}_i^w(t) \mid X_i = x\} = \hat{S}^w(x; t) + \frac{e_0^w(x)}{\hat{e}^w(x)} \left\{ \frac{S_0^w(x; t) G_0^w(x; t)}{\hat{G}^w(x; t)} - \hat{S}^w(x; t) \right\}.$$

Subtracting $S_0^w(x; t)$ from both sides gives

$$\begin{aligned} & \mathbb{E}\{\hat{\pi}_i^w(t) \mid X_i = x\} - S_0^w(x; t) \\ &= \hat{S}^w(x; t) - S_0^w(x; t) + \frac{e_0^w(x)}{\hat{e}^w(x)} \left\{ \frac{S_0^w(x; t) G_0^w(x; t)}{\hat{G}^w(x; t)} - \hat{S}^w(x; t) \right\}. \end{aligned}$$

Add and subtract $\frac{e_0^w(x)}{\hat{e}^w(x)} S_0^w(x; t)$ to obtain

$$\begin{aligned} & \mathbb{E}\{\hat{\pi}_i^w(t) \mid X_i = x\} - S_0^w(x; t) \\ &= \left\{1 - \frac{e_0^w(x)}{\hat{e}^w(x)}\right\} \{\hat{S}^w(x; t) - S_0^w(x; t)\} + \frac{e_0^w(x)}{\hat{e}^w(x)} \left\{\frac{S_0^w(x; t) G_0^w(x; t)}{\hat{G}^w(x; t)} - S_0^w(x; t)\right\} \\ &= \frac{\hat{e}^w(x) - e_0^w(x)}{\hat{e}^w(x)} \{\hat{S}^w(x; t) - S_0^w(x; t)\} + \frac{e_0^w(x)}{\hat{e}^w(x)} \frac{S_0^w(x; t)}{\hat{G}^w(x; t)} \{G_0^w(x; t) - \hat{G}^w(x; t)\}. \end{aligned}$$

Now use

$$\hat{\varphi}_i(t) = \hat{\pi}_i^1(t) - \hat{\pi}_i^0(t), \quad \varphi_i^0(t) = \pi_{i,0}^1(t) - \pi_{i,0}^0(t),$$

together with

$$\mathbb{E}\{\pi_{i,0}^w(t) \mid X_i = x\} = S_0^w(x; t), \quad w = 0, 1,$$

to obtain

$$\begin{aligned} \mathbb{E}[\hat{\varphi}_i(t) - \varphi_i^0(t) \mid X_i = x] &= \left(\mathbb{E}\{\hat{\pi}_i^1(t) \mid X_i = x\} - S_0^1(x; t)\right) \\ &\quad - \left(\mathbb{E}\{\hat{\pi}_i^0(t) \mid X_i = x\} - S_0^0(x; t)\right), \end{aligned}$$

which yields (5).

It remains to bound the right-hand side uniformly in x . By positivity and uniform consistency of $\hat{e}^w$ and $\hat{G}^w$ (Assumption 6), with probability tending to one,

$$\inf_x \hat{e}^w(x) \geq c/2, \quad \inf_x \hat{G}^w(x; t) \geq c/2.$$

Also $0 \leq S_0^w(x; t) \leq 1$. Therefore,

$$\begin{aligned} \sup_x \left| \frac{\hat{e}^w(x) - e_0^w(x)}{\hat{e}^w(x)} \{\hat{S}^w(x; t) - S_0^w(x; t)\} \right| &= O_p\left(r_n^{(e)} r_n^{(S)}\right), \\ \sup_x \left| \frac{e_0^w(x)}{\hat{e}^w(x)} \frac{S_0^w(x; t)}{\hat{G}^w(x; t)} \{G_0^w(x; t) - \hat{G}^w(x; t)\} \right| &= O_p\left(r_n^{(G)}\right). \end{aligned}$$

Summing over $w = 0, 1$ yields

$$\sup_x |\mathbb{E}[\hat{\varphi}_i(t) - \varphi_i^0(t) \mid X_i = x]| = O_p\left(r_n^{(e)} r_n^{(S)} + r_n^{(G)}\right).$$

If $\hat{G}^w = G_0^w$, the second term vanishes and the bound reduces to

$$O_p\left(r_n^{(e)} r_n^{(S)}\right).$$

This completes the proof.

Proof of Theorem 4: For notational simplicity, write

$$\hat{\varphi}_i = \hat{\varphi}_i(t), \quad \varphi_i^0 = \varphi_i^0(t), \quad \tau_0(x) = \tau_0(x; t).$$

Define the conditional mean of the estimated pseudo-outcome by

$$m_n(x) := \mathbb{E}[\hat{\varphi}_i \mid X_i = x].$$

Then

$$\hat{\tau}(x) - \tau_0(x) = \{\hat{\tau}(x) - m_n(x)\} + \{m_n(x) - \tau_0(x)\},$$

and therefore

$$\|\hat{\tau} - \tau_0\|_{L_2(P_X)} \leq \|\hat{\tau} - m_n\|_{L_2(P_X)} + \|m_n - \tau_0\|_{L_2(P_X)}. \quad (6)$$

We study the two terms separately.

By definition of τ_0 ,

$$\tau_0(x) = \mathbb{E}[\varphi_i^0 \mid X_i = x].$$

Hence

$$m_n(x) - \tau_0(x) = \mathbb{E}[\hat{\varphi}_i - \varphi_i^0 \mid X_i = x].$$

By Theorem 3,

$$\sup_x |m_n(x) - \tau_0(x)| = \sup_x |\mathbb{E}[\hat{\varphi}_i - \varphi_i^0 \mid X_i = x]| = O_p(r_n^{(e)} r_n^{(S)} + r_n^{(G)}).$$

Since $r_n^{(e)} r_n^{(S)} + r_n^{(G)} \rightarrow 0$,

$$\sup_x |m_n(x) - \tau_0(x)| = o_p(1).$$

Therefore,

$$\|m_n - \tau_0\|_{L_2(P_X)} \leq \sup_x |m_n(x) - \tau_0(x)| = o_p(1). \quad (7)$$

For any measurable function f ,

$$R_n(f) := \mathbb{E}[(\hat{\varphi}_i - f(X_i))^2] < \infty$$

by Assumption 7

Expanding the square conditional on X_i ,

$$\begin{aligned} R_n(f) &= \mathbb{E}\left(\mathbb{E}[(\hat{\varphi}_i - f(X_i))^2 \mid X_i]\right) \\ &= \mathbb{E}\left(\text{Var}(\hat{\varphi}_i \mid X_i) + \{m_n(X_i) - f(X_i)\}^2\right). \end{aligned}$$

Thus,

$$R_n(f) = \mathbb{E}[\text{Var}(\hat{\varphi}_i \mid X_i)] + \|f - m_n\|_{L_2(P_X)}^2.$$

Therefore m_n is the unique population minimizer of $R_n(f)$ over all square-integrable functions, and over $\mathcal{F}_n$ the excess risk satisfies

$$R_n(f) - R_n(m_n) = \|f - m_n\|_{L_2(P_X)}^2. \quad (8)$$

Let

$$f_n^* \in \arg \min_{f \in \mathcal{F}_n} R_n(f)$$

be a population risk minimizer over $\mathcal{F}_n$. Since $\hat{\tau}$ minimizes the empirical risk,

$$\mathbb{P}_n[(\hat{\varphi}_i - \hat{\tau}(X_i))^2] \leq \mathbb{P}_n[(\hat{\varphi}_i - f_n^*(X_i))^2].$$

Subtract $R_n(\hat{\tau})$ from the left-hand side and $R_n(f_n^*)$ from the right-hand side:

$$\begin{aligned} R_n(\hat{\tau}) - R_n(f_n^*) &\leq \left\{ R_n(\hat{\tau}) - \mathbb{P}_n[(\hat{\varphi}_i - \hat{\tau}(X_i))^2] \right\} \\ &\quad + \left\{ \mathbb{P}_n[(\hat{\varphi}_i - f_n^*(X_i))^2] - R_n(f_n^*) \right\}. \end{aligned}$$

Hence

$$R_n(\hat{\tau}) - R_n(f_n^*) \leq 2 \sup_{f \in \mathcal{F}_n} |\mathbb{P}_n[(\hat{\varphi}_i - f(X_i))^2] - R_n(f)| = o_p(1)$$

by Assumption 9. Using (8),

$$\|\hat{\tau} - m_n\|_{L_2(P_X)}^2 \leq \|f_n^* - m_n\|_{L_2(P_X)}^2 + o_p(1). \quad (9)$$

We now relate m_n to τ_0 . For any $f \in \mathcal{F}_n$,

$$\|f - m_n\|_{L_2(P_X)} \leq \|f - \tau_0\|_{L_2(P_X)} + \|\tau_0 - m_n\|_{L_2(P_X)}.$$

Taking infima over $f \in \mathcal{F}_n$ and using Assumption 8 together with (7), we obtain

$$\inf_{f \in \mathcal{F}_n} \|f - m_n\|_{L_2(P_X)} \leq \inf_{f \in \mathcal{F}_n} \|f - \tau_0\|_{L_2(P_X)} + \|\tau_0 - m_n\|_{L_2(P_X)} =$$

$$\inf_{f \in \mathcal{F}_n} \|f - \tau_0\|_{L_\infty(P_X)} + o_p(1) = o_p(1).$$

Since f_n^* minimizes $R_n(f)$ over $\mathcal{F}_n$, equivalently it minimizes $\|f - m_n\|_{L_2(P_X)}^2$ over $\mathcal{F}_n$. Therefore,

$$\|f_n^* - m_n\|_{L_2(P_X)} = o_p(1). \quad (10)$$

Combining (9) and (10) yields

$$\|\hat{\tau} - m_n\|_{L_2(P_X)} = o_p(1).$$

Finally, substitute this and (7) into (6):

$$\|\hat{\tau} - \tau_0\|_{L_2(P_X)} \leq \|\hat{\tau} - m_n\|_{L_2(P_X)} + \|m_n - \tau_0\|_{L_2(P_X)} = o_p(1).$$

Hence

$$\|\hat{\tau}(\cdot; t) - \tau_0(\cdot; t)\|_{L_2(P_X)} \xrightarrow{p} 0.$$

This completes the proof.

Proof of Theorem 5: Write

$$\hat{\varphi}_i = \hat{\varphi}_i(t), \quad \varphi_i^0 = \varphi_i^0(t), \quad \tau_0(x) = \tau_0(x; t),$$

and define

$$m_n(x) := \mathbb{E}(\hat{\varphi}_i \mid X_i = x).$$

Because $\hat{\tau}(\cdot; t)$ is obtained by regressing the cross-fitted pseudo-outcome on X_i , the relevant population regression target is m_n , not τ_0 directly. We therefore decompose

$$\|\hat{\tau} - \tau_0\|_{L_2(P_X)} \leq \|\hat{\tau} - m_n\|_{L_2(P_X)} + \|m_n - \tau_0\|_{L_2(P_X)}. \quad (11)$$

We first control the second term using the explicit construction of the pseudo-outcome. By definition,

$$m_n(x) - \tau_0(x) = \mathbb{E}\{\hat{\varphi}_i - \varphi_i^0 \mid X_i = x\},$$

since

$$\tau_0(x) = \mathbb{E}\{\varphi_i^0 \mid X_i = x\}.$$

By Theorem 3,

$$\sup_x |m_n(x) - \tau_0(x)| = O_p(r_n^{(e)} r_n^{(S)} + r_n^{(G)}).$$

Hence

$$\|m_n - \tau_0\|_{L_2(P_X)} \leq \sup_x |m_n - \tau_0| = O_p(r_n^{(e)} r_n^{(S)} + r_n^{(G)}), \quad (12)$$

which accommodates censoring estimation error with $r_n^{(G)}$.

Next, consider the population risk associated with the constructed pseudo-outcome,

$$R_n(f) := \mathbb{E}[(\hat{\varphi}_i - f(X_i))^2].$$

Since $m_n(x) = \mathbb{E}(\hat{\varphi}_i \mid X_i = x)$, conditioning on X_i gives

$$R_n(f) = R_n(m_n) + \|f - m_n\|_{L_2(P_X)}^2.$$

Thus m_n is the population least-squares target for the regression step, and

$$R_n(f) - R_n(m_n) = \|f - m_n\|_{L_2(P_X)}^2. \quad (13)$$

We now use the fact that τ_0 is Hölder smooth. By Assumption 8, there exists $f_n^* \in \mathcal{F}_n$ such that

$$\|f_n^* - \tau_0\|_{L_\infty(P_X)} \leq C n^{-\beta/(2\beta+d)}.$$

Therefore,

$$\|f_n^* - m_n\|_{L_2(P_X)} \leq \|f_n^* - \tau_0\|_{L_2(P_X)} + \|\tau_0 - m_n\|_{L_2(P_X)} = O_p\left(n^{-\beta/(2\beta+d)} + r_n^{(e)} r_n^{(S)} + r_n^{(G)}\right),$$

where we used (12) and the density bound in Assumption 3. Consequently,

$$R_n(f_n^*) - R_n(m_n) = \|f_n^* - m_n\|_{L_2(P_X)}^2 = O_p\left(n^{-2\beta/(2\beta+d)} + \{r_n^{(e)}r_n^{(S)} + r_n^{(G)}\}^2\right). \quad (14)$$

Because $\hat{\tau}$ is an approximate empirical minimizer, it follows that

$$\mathbb{P}_n[(\hat{\varphi}_i - \hat{\tau}(X_i))^2] \leq \mathbb{P}_n[(\hat{\varphi}_i - f_n^*(X_i))^2] + \delta_n^2,$$

by Assumption 10. Adding and subtracting population risks yields

$$\begin{aligned} R_n(\hat{\tau}) - R_n(f_n^*) &\leq \{R_n(\hat{\tau}) - \mathbb{P}_n[(\hat{\varphi}_i - \hat{\tau}(X_i))^2]\} \\ &\quad + \{\mathbb{P}_n[(\hat{\varphi}_i - f_n^*(X_i))^2] - R_n(f_n^*)\} + \delta_n^2 \\ &\leq 2 \sup_{f \in \mathcal{F}_n} |\mathbb{P}_n[(\hat{\varphi}_i - f(X_i))^2] - R_n(f)| + \delta_n^2 \\ &= O_p(\rho_n) + \delta_n^2, \end{aligned}$$

where the last equality is due to Assumption 9. Combining this with (14), we obtain

$$R_n(\hat{\tau}) - R_n(m_n) = O_p\left(n^{-2\beta/(2\beta+d)} + \rho_n + \delta_n^2 + \{r_n^{(e)}r_n^{(S)} + r_n^{(G)}\}^2\right).$$

Applying (13),

$$\|\hat{\tau} - m_n\|_{L_2(P_X)}^2 = O_p\left(n^{-2\beta/(2\beta+d)} + \rho_n + \delta_n^2 + \{r_n^{(e)}r_n^{(S)} + r_n^{(G)}\}^2\right),$$

and hence

$$\|\hat{\tau} - m_n\|_{L_2(P_X)} = O_p\left(n^{-\beta/(2\beta+d)} + \rho_n^{1/2} + \delta_n + r_n^{(e)}r_n^{(S)} + r_n^{(G)}\right). \quad (15)$$

Finally, combine (11), (12), and (15) to conclude

$$\|\hat{\tau} - \tau_0\|_{L_2(P_X)} = O_p\left(n^{-\beta/(2\beta+d)} + \rho_n^{1/2} + \delta_n + r_n^{(e)}r_n^{(S)} + r_n^{(G)}\right).$$

This completes the proof.

Proof of Lemma 1 We give a constructive proof based on a local polynomial approximation in the time direction. We aim to show that the coefficient functions in this expansion inherit the β_x -Hölder smoothness of $\tau_0(x, t)$ in the covariate direction.

Write

$$\beta_t = q_t + \alpha_t, \quad q_t = \lfloor \beta_t \rfloor, \quad \alpha_t \in (0, 1].$$

Fix an integer $K \geq 1$. Partition $[0, 1]$ into K equal subintervals

$$I_r = \left[\frac{r-1}{K}, \frac{r}{K}\right), \quad r = 1, \dots, K-1, \quad I_K = \left[\frac{K-1}{K}, 1\right].$$

Let u_r denote the midpoint of I_r . For each interval I_r , we approximate $\tau_0(x, t)$ in the time variable t by its Taylor polynomial of order q_t around u_r .

For $r = 1, \dots, K$ and $\ell = 0, \dots, q_t$, define

$$B_{r,\ell}(t) := (t - u_r)^\ell \mathbf{1}\{t \in I_r\},$$

and define the corresponding coefficient functions

$$b_{r,\ell}(x) := \frac{\partial_t^\ell \tau_0(x, u_r)}{\ell!}, \quad x \in \mathcal{X}.$$

Then, for $t \in I_r$,

$$\sum_{\ell=0}^{q_t} b_{r,\ell}(x) B_{r,\ell}(t) = \sum_{\ell=0}^{q_t} \frac{\partial_t^\ell \tau_0(x, u_r)}{\ell!} (t - u_r)^\ell,$$

which is the Taylor polynomial of $\tau_0(x, \cdot)$ at u_r .

We now verify that, uniformly over r and ℓ , the function $b_{r,\ell}$ belongs to a β_x -Hölder ball on $\mathcal{X}$.

Write

$$\beta_x = q_x + \alpha_x, \quad q_x = \lfloor \beta_x \rfloor, \quad \alpha_x \in (0, 1].$$

Fix $r \in \{1, \dots, K\}$ and $\ell \in \{0, \dots, q_t\}$. By definition,

$$b_{r,\ell}(x) = \frac{\partial_t^\ell \tau_0(x, u_r)}{\ell!}.$$

For any multi-index ν with $|\nu| \leq q_x$,

$$D_x^\nu b_{r,\ell}(x) = \frac{D_x^\nu \partial_t^\ell \tau_0(x, u_r)}{\ell!}.$$

Because $\tau_0 \in \mathcal{H}^{\beta_x, \beta_t}(M)$, Definition 2 implies that all such derivatives exist and are uniformly bounded:

$$\sup_{x \in \mathcal{X}} |D_x^\nu b_{r,\ell}(x)| \leq \frac{1}{\ell!} \sup_{x \in \mathcal{X}} |D_x^\nu \partial_t^\ell \tau_0(x, u_r)| \leq \frac{M}{\ell!} \leq M.$$

Next, for any multi-index ν with $|\nu| = q_x$, the Hölder condition in the x -direction yields

$$\frac{|D_x^\nu b_{r,\ell}(x) - D_x^\nu b_{r,\ell}(y)|}{\|x - y\|^{\alpha_x}} = \frac{|D_x^\nu \partial_t^\ell \tau_0(x, u_r) - D_x^\nu \partial_t^\ell \tau_0(y, u_r)|}{\ell! \|x - y\|^{\alpha_x}} \leq \frac{M}{\ell!} \leq M.$$

Thus each $b_{r,\ell}$ lies in a β_x -Hölder class on $\mathcal{X}$, with Hölder radius bounded by a constant depending only on M, β_x, β_t . In particular, there exists a constant $C_M > 0$, depending only on M, β_x, β_t , such that

$$b_{r,\ell} \in \mathcal{H}^{\beta_x}(C_M) \quad \text{for all } r = 1, \dots, K, \ell = 0, \dots, q_t.$$

Fix $j \in \{1, \dots, J\}$, and let $r(j)$ denote the unique index such that $t_j \in I_{r(j)}$. Define the local polynomial approximation at the grid point t_j by

$$P_j(x) := \sum_{\ell=0}^{q_t} b_{r(j),\ell}(x) B_{r(j),\ell}(t_j) = \sum_{\ell=0}^{q_t} \frac{\partial_t^\ell \tau_0(x, u_{r(j)})}{\ell!} (t_j - u_{r(j)})^\ell.$$

We claim that

$$\sup_{x \in \mathcal{X}} |\tau_0(x, t_j) - P_j(x)| \leq CK^{-\beta_t},$$

uniformly in j .

To see this, fix $x \in \mathcal{X}$. By Taylor's theorem in the one-dimensional variable t , applied to the function $t \mapsto \tau_0(x, t)$ around the point $u_{r(j)}$, there exists a point $\xi_{x,j}$ between $u_{r(j)}$ and t_j such that

$$\tau_0(x, t_j) - P_j(x) = \frac{\partial_t^{q_t} \tau_0(x, \xi_{x,j}) - \partial_t^{q_t} \tau_0(x, u_{r(j)})}{q_t!} (t_j - u_{r(j)})^{q_t}.$$

Because $\tau_0 \in \mathcal{H}^{\beta_x, \beta_t}(M)$, the derivative $\partial_t^{q_t} \tau_0(x, \cdot)$ is α_t -Hölder uniformly in x , so

$$|\partial_t^{q_t} \tau_0(x, \xi_{x,j}) - \partial_t^{q_t} \tau_0(x, u_{r(j)})| \leq M |\xi_{x,j} - u_{r(j)}|^{\alpha_t} \leq M |t_j - u_{r(j)}|^{\alpha_t}.$$

Therefore,

$$|\tau_0(x, t_j) - P_j(x)| \leq C |t_j - u_{r(j)}|^{q_t + \alpha_t} = C |t_j - u_{r(j)}|^{\beta_t}.$$

Since both t_j and $u_{m(j)}$ lie in the same interval $I_{m(j)}$, whose length is K^{-1} , we have

$$|t_j - u_{m(j)}| \leq K^{-1}.$$

Hence

$$\sup_{x \in \mathcal{X}} |\tau_0(x, t_j) - P_j(x)| \leq CK^{-\beta_t},$$

uniformly in j . Since P_X is a probability measure,

$$\|\tau_0(\cdot, t_j) - P_j\|_{L_2(P_X)}^2 \leq \sup_{x \in \mathcal{X}} |\tau_0(x, t_j) - P_j(x)|^2 \leq CK^{-2\beta_t}.$$

Taking the supremum over j gives

$$\sup_{1 \leq j \leq J} \|\tau_0(\cdot, t_j) - P_j\|_{L_2(P_X)}^2 \leq CK^{-2\beta_t}.$$

So far we have constructed $(q_t + 1)K$ basis functions $B_{r,\ell}$ and the same number of coefficient functions $b_{r,\ell}$. Each approximant P_j can be written as

$$P_j(x) = \sum_{r=1}^K \sum_{\ell=0}^{q_t} b_{r,\ell}(x) B_{r,\ell}(t_j).$$

Relabel the collection

$$\{B_{m,\ell} : 1 \leq r \leq K, 0 \leq \ell \leq q_t\}$$

as

$$\psi_1, \dots, \psi_{(q_t+1)K},$$

and similarly relabel the coefficient functions

$$\{b_{r,\ell} : 1 \leq m \leq K, 0 \leq \ell \leq q_t\}$$

as

$$b_1, \dots, b_{(q_t+1)K}.$$

Now define

$$m := (q_t + 1)K.$$

Since q_t depends only on β_t , this is only a constant rescaling of K , and thus

$$K^{-2\beta_t} \asymp m^{-2\beta_t}.$$

Absorbing constants into C , we obtain

$$\sup_{1 \leq j \leq J} \left\| \tau_0(\cdot; t_j) - \sum_{k=1}^m b_k(\cdot) B_k(t_j) \right\|_{L_2(P_X)}^2 \leq C m^{-2\beta_t}.$$

Finally,

$$\frac{1}{J} \sum_{j=1}^J \left\| \tau_0(\cdot; t_j) - \sum_{k=1}^m b_k(\cdot) B_k(t_j) \right\|_{L_2(P_X)}^2 \leq C m^{-2\beta_t}.$$

This completes the proof. This completes the proof.

Proof of Theorem 6: For $j = 1, \dots, J$, let

$$\bar{\varphi}_j(x) = \mathbb{E}\{\hat{\varphi}_i(t_j) \mid X_i = x\}$$

By Theorem 3

$$\sup_x |\bar{\varphi}_j(x) - \tau_0(x; t_j)| = O_p \left(r_n^{(e)} r_n^{(S)} + r_n^{(G)} \right)$$

Therefore

$$\frac{1}{J} \sum_{j=1}^J \|\bar{\varphi}_j(\cdot) - \tau_0(\cdot; t_j)\|_{L_2(P_X)}^2 = O_p \left(\left\{ r_n^{(e)} r_n^{(S)} + r_n^{(G)} \right\}^2 \right) \quad (16)$$

For $f \in \mathcal{F}_{n,m}$, let

$$\tau_f(x; t) = \sum_{k=1}^m f_k(x) B_k(t)$$

We then define the empirical and population risks as

$$R_{n,J}(f) = \mathbb{P}_n \left[\frac{1}{J} \sum_{j=1}^J \{\hat{\varphi}_i(t_j) - \tau_f(X_i; t_j)\}^2 \right]$$

and

$$R_J(f) = \mathbb{E} \left[\frac{1}{J} \sum_{j=1}^J \{\hat{\varphi}_i(t_j) - \tau_f(X_i; t_j)\}^2 \right]$$

We can write

$$\begin{aligned} & \mathbb{E} \left[\{ \hat{\varphi}_i(t_j) - \tau_f(X_i; t_j) \}^2 \right] \\ &= \mathbb{E} \left[\{ \hat{\varphi}_i(t_j) - \bar{\varphi}_j(X_i) \}^2 \right] + \|\tau_f(\cdot; t_j) - \bar{\varphi}_j(\cdot)\|_{L_2(P_X)}^2, \end{aligned}$$

where the interaction term is zero because $\bar{\varphi}_j(X_i)$ is the expectation of $\hat{\varphi}_i(t_j)$ conditional on X_i . Then,

$$R_J(f) = C_{n,J} + \frac{1}{J} \sum_{j=1}^J \|\tau_f(\cdot; t_j) - \bar{\varphi}_j(\cdot)\|_{L_2(P_X)}^2, \quad (17)$$

where

$$C_{n,J} = \frac{1}{J} \sum_{j=1}^J \mathbb{E} \left[\{ \hat{\varphi}_i(t_j) - \bar{\varphi}_j(X_i) \}^2 \right]$$

does not depend on f .

By Lemma 1, for fixed m , there exist basis functions $B_1(t), \dots, B_m(t)$ and coefficient functions $b_1, \dots, b_m \in \mathcal{H}^{\beta_x}(C)$ such that

$$\sum_{j=1}^J \left\| \tau_0(\cdot; t_j) - \sum_{k=1}^m b_k(\cdot) B_m(t_j) \right\|_{L_2(P_X)}^2 \lesssim m^{-2\beta_t}.$$

Since each b_k is β_x -Hölder, by Assumption 8, there exists $f^* = (f_1^*, \dots, f_m^*)^\top \in \mathcal{F}_{n,m}$ such that

$$\max_{k=1, \dots, m} \|f_k^* - b_k\|_{L_2(P_X)}^2 \lesssim n^{-2\beta_x/(2\beta_x+d)}.$$

We define the optimization error as

$$\delta_{n,J,m}^2 = R_{n,J}(\hat{f}) - \inf_{f \in \mathcal{F}_{n,m}} R_{n,J}(f)$$

then for f^* , we have

$$\begin{aligned} R_J(\hat{f}) - R_J(f^*) &\leq \left\{ R_J(\hat{f}) - R_{n,J}(\hat{f}) \right\} \\ &\quad + \{ R_{n,J}(f^*) - R_J(f^*) \} + \delta_{n,J,m}^2 \\ &\leq 2 \sup_{f \in \mathcal{F}_{n,m}} |R_{n,J}(f) - R_J(f)| + \delta_{n,J,m}^2. \end{aligned}$$

Using (17), we have

$$\begin{aligned} & \frac{1}{J} \sum_{j=1}^J \left\| \tau_{\hat{f}}(\cdot; t_j) - \bar{\varphi}_j(\cdot) \right\|_{L_2(P_X)}^2 \\ &\leq \frac{1}{J} \sum_{j=1}^J \|\tau_{f^*}(\cdot; t_j) - \bar{\varphi}_j(\cdot)\|_{L_2(P_X)}^2 + 2 \sup_{f \in \mathcal{F}_{n,m}} |R_{n,J}(f) - R_J(f)| + \delta_{n,J,m}^2. \end{aligned} \quad (18)$$

We first bound $\sum_{j=1}^J \left\| \tau_{\hat{f}}(\cdot; t_j) - \bar{\varphi}_j(\cdot) \right\|_{L_2(P_X)}^2$. For the basis functions $B_1(t), \dots, B_m(t)$ and coefficient functions $b_1, \dots, b_m$ satisfy the conditions in Lemma 1, let

$$\tau_b(x; t) = \sum_{k=1}^m b_k(x) B_k(t),$$

we have

$$\tau_{f^*}(x; t_j) - \tau_b(x; t_j) = \sum_{k=1}^m \{f_k^*(x) - b_k(x)\} B_k(t_j).$$

By the spline stability condition,

$$\lambda_{\max} \left(\frac{1}{J} \mathbf{B}^\top \mathbf{B} \right) \leq \frac{C_B}{m},$$

and therefore

$$\begin{aligned} & \frac{1}{J} \sum_{j=1}^J \left\| \tau_{f^*}(\cdot; t_j) - \tau_b(\cdot; t_j) \right\|_{L_2(P_X)}^2 \\ & \leq \frac{C_B}{m} \sum_{k=1}^m \|f_k^* - b_k\|_{L_2(P_X)}^2 \\ & \lesssim n^{-2\beta_x/(2\beta_x+d)}. \end{aligned} \quad (19)$$

Combine this with

$$\sum_{j=1}^J \left\| \tau_0(\cdot; t_j) - \tau_b(\cdot; t_j) \right\|_{L_2(P_X)}^2 \lesssim m^{-2\beta_t}.$$

we have

$$\frac{1}{J} \sum_{j=1}^J \left\| \tau_{f^*}(\cdot; t_j) - \tau_0(\cdot; t_j) \right\|_{L_2(P_X)}^2 \lesssim m^{-2\beta_t} + n^{-2\beta_x/(2\beta_x+d)}. \quad (20)$$

Next, by

$$\tau_{f^*}(\cdot; t_j) - \bar{\varphi}_j(\cdot) = \{\tau_{f^*}(\cdot; t_j) - \tau_0(\cdot; t_j)\} + \{\tau_0(\cdot; t_j) - \bar{\varphi}_j(\cdot)\},$$

together with (16) and (20),

$$\begin{aligned} & \frac{1}{J} \sum_{j=1}^J \left\| \tau_{f^*}(\cdot; t_j) - \bar{\varphi}_j(\cdot) \right\|_{L_2(P_X)}^2 \\ & = O_p \left(m^{-2\beta_t} + n^{-2\beta_x/(2\beta_x+d)} + \left\{ r_n^{(e)} r_n^{(S)} + r_n^{(G)} \right\}^2 \right). \end{aligned} \quad (21)$$

We then bound the empirical process fluctuation $\sup_{f \in \mathcal{F}_{n,m}} |R_{n,J}(f) - R_J(f)|$. For the scalar output DNN $\mathcal{F}_{n,1}$, let ρ_n denote the corresponding empirical

process fluctuation from Theorem 5. For $\tau, \tilde{\tau} \in \mathcal{T}_{n,m}$, define the aggregated metric

$$d_J^2(\tau, \tilde{\tau}) = \frac{1}{J} \sum_{j=1}^J \|\tau(\cdot; t_j) - \tilde{\tau}(\cdot; t_j)\|_{L_2(P_X)}^2.$$

Since $\mathcal{F}_{n,m} \subseteq \mathcal{F}_{n,1}^m$ and, by the basis stability condition,

$$d_J^2(\tau_f, \tau_{\tilde{f}}) \leq \frac{C_B}{m} \sum_{k=1}^m \|f_k - \tilde{f}_k\|_{L_2(P_X)}^2,$$

the metric entropy of the induced joint class satisfies

$$\log N(\varepsilon, \mathcal{T}_{n,m}, d_J) \leq m \log N(c\varepsilon, \mathcal{F}_{n,1}, L_2(P_X))$$

for some constant $c > 0$. Applying the standard symmetrization, contraction, and entropy-integral argument for the squared-loss class, together with the entropy condition on $\mathcal{F}_{n,1}$, gives

$$\sup_{f \in \mathcal{F}_{n,m}} |R_{n,J}(f) - R_J(f)| = O_p(\sqrt{m} \rho_n).$$

Combining this together with (21) and (18), we have

$$\begin{aligned} & \frac{1}{J} \sum_{j=1}^J \left\| \tau_{\tilde{f}}(\cdot; t_j) - \bar{\varphi}_j(\cdot) \right\|_{L_2(P_X)}^2 \\ &= O_p \left(m^{-2\beta_t} + n^{-2\beta_x/(2\beta_x+d)} + \sqrt{m} \rho_n + \delta_{n,J,m}^2 + \left\{ r_n^{(e)} r_n^{(S)} + r_n^{(G)} \right\}^2 \right). \end{aligned} \quad (22)$$

Finally, since

$$\hat{\tau}(\cdot; t_j) - \tau_0(\cdot; t_j) = \{\hat{\tau}(\cdot; t_j) \bar{\varphi}_j(\cdot)\} + \{\bar{\varphi}_j(\cdot) - \tau_0(\cdot; t_j)\},$$

with (16) and (22), we obtain

$$\begin{aligned} & \frac{1}{J} \sum_{j=1}^J \left\| \hat{\tau}(\cdot; t_j) - \tau_0(\cdot; t_j) \right\|_{L_2(P_X)}^2 \\ &= O_p \left(m^{-2\beta_t} + n^{-2\beta_x/(2\beta_x+d)} + \sqrt{m} \rho_n + \delta_{n,J,m}^2 + \left\{ r_n^{(e)} r_n^{(S)} + r_n^{(G)} \right\}^2 \right). \end{aligned}$$

This completes the proof.

Additional Simulation Results

To further investigate the finite-sample behavior observed in Case 1, we conducted an additional set of simulations under a simplified setting with lower-dimensional covariates ($d = 10$) and a larger training sample size ($N = 2000$). This setting is closer to those often considered in prior methodological work and therefore provides a useful benchmark for assessing whether the

performance patterns in Case 1 were driven by limited sample size and model complexity. The results are reported in Table A.1.

As shown in Table A.1, all methods worked better under this easier setting, and several competing approaches achieve noticeably smaller MSE than DSL. In particular, PSDRL-NFT-BART attains the lowest MSE overall, approximately 0.014 across all choices of J , while X-Learner and causal survival forest also perform strongly, with MSEs around 0.020 and 0.021, respectively. DRL-NFT-BART improves substantially as the number of evaluation time points increases, with MSE decreasing from 0.042 at $J = 1$ to 0.022 at $J = 100$. By comparison, DSL yields MSEs of 0.042, 0.046, and 0.047 for $J = 1, 50, 100$, respectively, which are somewhat larger but remain reasonably stable across grid sizes.

A notable feature of DSL is its very small bias. Its bias remains close to zero across all three values of J , around 0.005, and is clearly smaller than that of several competitors, including M-Learner, PSDRL-NFT-BART, and causal survival forest, all of which exhibit substantially larger systematic bias. R-Learner and DRL-NFT-BART also have relatively small bias, but their MSE is not uniformly better than that of the strongest competitors. Taken together, these results suggest that in a simpler and more data-rich setting, competing methods can achieve very strong predictive accuracy, sometimes surpassing DSL in MSE, while DSL continues to provide stable estimation with consistently low bias. This supports our interpretation that the larger performance gaps seen in Case 1 may be primarily attributable to the more challenging finite-sample and high-complexity setting.

Table A.1 Performance comparison under simplified setting

	$J = 1$	$J = 50$	$J = 100$
MSE			
Deep Survival Learner	0.042(0.041, 0.043)	0.046(0.045, 0.048)	0.047(0.045, 0.048)
M-Learner	0.049(0.047, 0.051)	0.049(0.047, 0.051)	0.049(0.047, 0.051)
R-Learner	0.049(0.047, 0.051)	0.051(0.049, 0.053)	0.048(0.046, 0.050)
X-Learner	0.020(0.019, 0.021)	0.020(0.019, 0.021)	0.020(0.019, 0.021)
DRL-NFT-BART	0.042(0.026, 0.057)	0.025(0.020, 0.030)	0.022(0.020, 0.024)
PSDRL-NFT-BART	0.014(0.013, 0.015)	0.014(0.014, 0.015)	0.014(0.014, 0.015)
Causal Survival Forest	0.021(0.020, 0.022)	0.021(0.020, 0.022)	0.021(0.020, 0.022)
Bias			
Deep Survival Learner	0.005(0.001, 0.010)	0.005(0.000, 0.009)	0.005(0.000, 0.009)
M-Learner	0.109(0.105, 0.113)	0.109(0.105, 0.113)	0.109(0.105, 0.113)
R-Learner	0.012(0.006, 0.017)	0.007(0.000, 0.014)	0.012(0.006, 0.018)
X-Learner	0.037(0.031, 0.043)	0.041(0.036, 0.046)	0.038(0.032, 0.043)
DRL-NFT-BART	0.013(0.003, 0.023)	0.012(0.005, 0.020)	0.013(0.007, 0.020)
PSDRL-NFT-BART	0.083(0.079, 0.087)	0.084(0.081, 0.088)	0.084(0.080, 0.088)
Causal Survival Forest	0.107(0.103, 0.111)	0.107(0.102, 0.111)	0.107(0.102, 0.111)

¹ The simulation setting uses a training sample size of $N = 2000$ and a testing sample size of 400, with covariate dimension $d = 10$ and approximately 35% censoring.

² Mean squared error and 95% Monte Carlo confidence intervals are computed based on 200 simulation replicates.

³ Conditional average treatment effects are evaluated at $J \in \{1, 50, 100\}$ time points, equally spaced between the 20th and 80th percentiles of the observed survival time distribution.