



Causal survival analysis

Estimation of the Average Treatment Effect (ATE): Practical Recommendations

ISSN 2824-7795

Charlotte Voinot¹ CMEI, Sanofi R&D
Premedical, INRIA
INSERM

Université de Montpellier
Clément Berenfeld² Universität Potsdam, Potsdam, Germany
Imke Mayer³ Charité – Universität Berlin, Berlin, Germany

Bernard Sebastien⁴ CMEI, Sanofi R&D
Julie Josse⁵ Premedical, INRIA
INSERM
Université de Montpellier

Date published: 2025-02-20 Last modified: 2025-02-20

Abstract

Causal survival analysis combines survival analysis and causal inference to evaluate the effect of a treatment or intervention on a time-to-event outcome, such as survival time. It offers an alternative to relying solely on Cox models for assessing these effects. In this paper, we present a comprehensive review of estimators for the average treatment effect measured with the restricted mean survival time, including regression-based methods, weighting approaches, and hybrid techniques. We investigate their theoretical properties and compare their performance through extensive numerical experiments. Our analysis focuses on the finite-sample behavior of these estimators, the influence of nuisance parameter selection, and their robustness and stability under model misspecification. By bridging theoretical insights with practical evaluation, we aim to equip practitioners with both state-of-the-art implementations of these methods and practical guidelines for selecting appropriate estimators for treatment effect estimation. Among the approaches considered, G-formula two-learners, AIPCW-AIPTW, Buckley-James estimators, and causal survival forests emerge as particularly promising.

Keywords: Restricted Mean Survival Time, Randomized Control Trial, Observational Study, Censoring

Contents

1 Introduction	2
1.1 Context and motivations	2
1.2 Definition of the estimand: the RMST	3
1.3 Organisation of the paper	5
1.4 Notations	5

¹Corresponding author: charlotte.voinot@sanofi.com

²Corresponding author: clement.berenfeld@uni-potsdam.de

³Corresponding author: imke.mayer@owkin.com, now affiliated with Owkin

⁴Corresponding author: bernard.sebastien@sanofi.com

⁵Corresponding author: julie.josse@inria.fr

2	Causal survival analysis in Randomized Controlled Trials	6
2.1	Independent censoring: the Kaplan-Meier estimator	7
2.2	Conditionally independent censoring	9
2.2.1	The G-formula and the Cox Model	10
2.2.2	Censoring unbiased transformations	12
2.2.3	Augmented corrections	16
3	Causal survival analysis in observational studies	16
3.1	Independent censoring	17
3.2	Conditional independent censoring	19
3.2.1	IPTW-IPCW transformations	19
3.2.2	IPTW-BJ transformations	20
3.2.3	Double augmented corrections	21
4	Implementation	21
4.1	Summary of the estimators	22
4.2	Implementation of the estimators	22
4.3	Available packages	24
5	Simulations	25
5.1	RCT	25
5.2	Observational data	28
5.3	Mispecification of nuisance components	30
6	Conclusion	31
7	Disclosure	34
	References	34
8	Appendix A: Proofs	36
8.1	Proofs of Section 2.1	36
8.2	Proofs of Section 2.2	38
8.3	Proofs of Section 3.1	39
8.4	Proofs of Section 3.2	39
9	Appendix B: Descriptive statistics	41
9.1	RCT	41
9.2	Observational study with linear relationship	43
9.3	Observational study with interaction	46

1 Introduction

1.1 Context and motivations

Causal survival analysis is a growing field that integrates causal inference (D. B. Rubin 1974; Hernán and Robins 2010) with survival analysis (Kalbfleisch and Prentice 2002) to evaluate the impact of treatments on time-to-event outcomes, while accounting for censoring situations where only partial information about an event’s occurrence is available.

Being a relatively new domain, the existing literature, though vast, remains fragmented. As a result, a clear understanding of the theoretical properties of various estimators is challenging to obtain.

Moreover, the implementation of proposed methods is limited, leaving researchers confronted with a range of available estimators and the need to make numerous methodological decisions. There is a pressing need for a comprehensive survey that organizes the available methods, outlines the underlying assumptions, and provides an evaluation of estimator performance — particularly in finite sample settings. Such a survey also has the potential to help identify remaining methodological challenges that need to be addressed. This need becomes increasingly urgent as causal survival analysis gains traction in both theoretical and applied domains. For instance, its applications to external control arm analyses are particularly relevant in the context of single-arm clinical trials, where traditional comparator arms are unavailable. Regulatory guidelines have begun to acknowledge and support such semi-experimental approaches, reflecting the broader evolution of trial design and therapeutic innovation in precision medicine, see for instance (European Medicines Agency 2024).

By synthesizing the theoretical foundations, assumptions, and performance of various estimators, a survey on existing causal survival analysis methods would provide researchers and practitioners with the necessary tools to make informed methodological choices. This is crucial for fostering robust and reliable applications of causal survival analysis in both academic research and practical settings, where precise and valid results are paramount.

In this paper, we focus our attention to the estimation of the Restricted Mean Survival Time (RMST), a popular causal measure in survival analysis which offers an intuitive interpretation of the average survival time over a specified period. In particular, we decided to not cover the estimation of Hazard Ratio (HR), which has been prominently used but often questioned due to its potential non-causal nature (Martinussen, Vansteelandt, and Andersen 2020). Additionally, unlike the Hazard Ratio, the RMST has the desirable property of being a collapsible measure, meaning that the population effect can be expressed as a weighted average of subgroup effects, with positive weights that sum to one (Huitfeldt, Stensrud, and Suzuki 2019).

1.2 Definition of the estimand: the RMST

We set the analysis in the potential outcome framework, where a patient, described by a vector of covariates $X \in \mathbb{R}^p$, either receives a treatment ($A = 1$) or is in the control group ($A = 0$). The patient will then survive up to a certain time $T(0) \in \mathbb{R}^+$ in the control group, or up to a time $T(1) \in \mathbb{R}^+$ in the treatment group. In practice, we cannot simultaneously have access to $T(0)$ and $T(1)$, as one patient is either treated or control, but only to T defined as follows:

Assumption. (Stable Unit Treatment Value Assumption: SUTVA)

$$T = AT(1) + (1 - A)T(0). \quad (1)$$

Due to potential censoring, the outcome T is not completely observed. The most common form of censoring is right-censoring (also known as type II censoring), which occurs when the event of interest has not taken place by the end of the observation period, indicating that it may have occurred later if the observation had continued (Turkson, Ayiah-Mensah, and Nimoh 2021). We focus in this study on this type of censoring only and we assume that we observe $\tilde{T} = T \wedge C = \min(T, C)$ for some censoring time $C \in \mathbb{R}^+$. When an observation is censored, the observed time is equal to the censoring time.

We also assume that we know whether an outcome is censored or not. In other words, we observe the censoring status variable $\Delta = \mathbb{I}\{T < C\}$, where $\mathbb{I}\{\cdot\}$ is the indicator function. Δ is 1 if the true outcome is observed, and 0 if it is censored.

We assume observing a n -sample of variables $(X, A, \tilde{T}, \Delta)$ stemming from an n -sample of the partially unobservable variables $(X, A, T(0), T(1), C)$. A toy example of such data is given in Table 1.

Table 1: Example of a survival dataset. In practice, only X , A , \tilde{T} and Δ are observed.

ID	Co-variates			Treatment	Censoring	Status	Potential outcomes		True outcome	Observed outcome
	X_1	X_2	X_3				$T(0)$	$T(1)$		
i	X_1	X_2	X_3	A	C	Δ	$T(0)$	$T(1)$	T	\tilde{T}
1	1	1.5	4	1	?	1	?	200	200	200
2	5	1	2	0	?	1	100	?	100	100
3	9	0.5	3	1	200	0	?	?	?	200

Our aim is to estimate the Average Treatment Effect (ATE) defined as the difference between the Restricted Mean Survival Time of the treated and controls (Royston and Parmar 2013).

Definition 1.1. (Causal effect: Difference between Restricted Mean Survival Time)

$$\theta_{\text{RMST}} = \mathbb{E} [T(1) \wedge \tau - T(0) \wedge \tau],$$

where $a \wedge b := \min(a, b)$ for $a, b \in \mathbb{R}$.

We define the survival functions $S^{(a)}(t) := \mathbb{P}(T(a) > t)$ for $a \in \{0, 1\}$, i.e., the probability that a treated or non-treated individual will survive beyond a given time t . Likewise, we let $S(t) := \mathbb{P}(T > t)$, and $S_C(t) := \mathbb{P}(C > t)$. We also let $G(t) := \mathbb{P}(C \leq t)$ be the left-limit of the survival function S_C . Because $T(a) \wedge \tau$ are non-negative random variables, one can easily express the restricted mean survival time using the survival functions:

$$\mathbb{E}(T(a) \wedge \tau) = \int_0^\infty \mathbb{P}(T(a) \wedge \tau > t) dt = \int_0^\tau S^{(a)}(t) dt. \quad (2)$$

Consequently, θ_{RMST} can be interpreted as the mean difference between the survival function of treated and control until a fixed time horizon τ . A difference in RMST $\theta_{\text{RMST}} = 10$ days with $\tau = 200$ means that, on average, the treatment increases the survival time by 10 days at 200 days. We give a visual interpretation of RMST in Figure 1.

Although the present work focuses on the estimation of the difference in RMST, we would like to stress that the causal effect can be assessed through other measures, such as for instance the difference of the survival functions

$$\theta_{\text{SP}} := S^{(1)}(\tau) - S^{(0)}(\tau)$$

for some time τ , see for instance (Ozenne et al. 2020). As mentioned in Section 1.1, another widely used measure (though not necessarily causal) is the hazards ratio, defined as

$$\theta_{\text{HR}} := \frac{\lambda^{(1)}(\tau)}{\lambda^{(0)}(\tau)},$$

where the hazard function $\lambda^{(a)}$ is defined as

$$\lambda^{(a)}(t) := \lim_{h \rightarrow 0^+} \frac{\mathbb{P}(T(a) \in [t, t + h) | T(a) \geq t)}{h}.$$

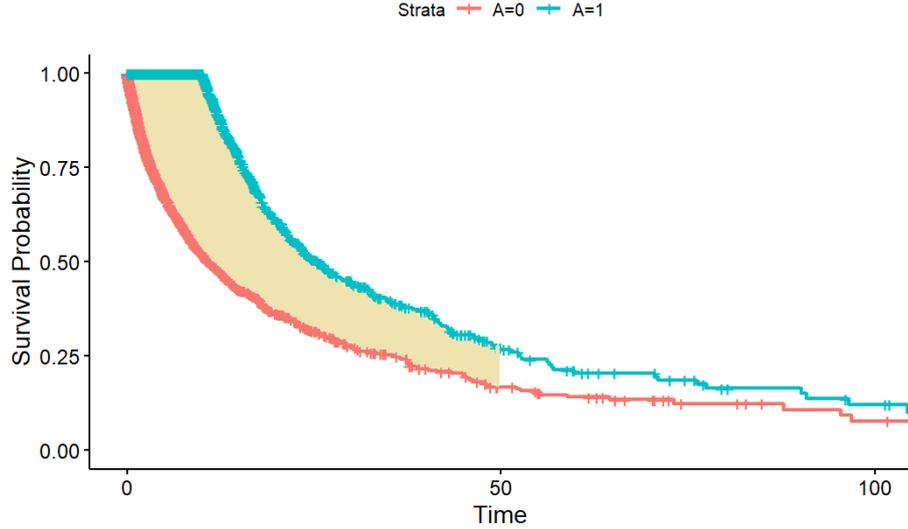


Figure 1: Plot of the estimated survival curves on synthetic toy-data. The θ_{RMST} at $\tau = 50$ corresponds to the yellow shaded area between the two survival curves. The curves have been estimated using Kaplan-Meier estimator, see Section 2.1.

in a continuous setting, or as $\lambda^{(a)}(t) := \mathbb{P}(T(a) = t | T(a) \geq t)$ when the survival times are discrete. The hazard functions and the survival functions are linked through the identities

$$S^{(a)}(t) = \exp(-\Lambda^{(a)}(t)) \quad \text{where} \quad \Lambda^{(a)}(t) := \int_0^t \lambda^{(a)}(s) ds, \quad (3)$$

in the continuous case. The functions $\Lambda^{(a)}$ are called the cumulative hazard functions. In the discrete case, we have in turn

$$S^{(a)}(t) = \prod_{t_k \leq t} (1 - \lambda^{(a)}(t_k)), \quad (4)$$

where $\{t_1, \dots, t_K\}$ are the atoms of $T^{(a)}$. These hazard functions are classically used to model the survival times and the censoring times, see Section 2.2.1.

1.3 Organisation of the paper

In this paper, we detail the minimal theoretical framework that allows the analysis of established RMST estimators in the context of both Randomized Controlled Trials (Section 2) and observational data (Section 3). We give their statistical properties (consistency, asymptotic normality) along with proofs when possible. We then conduct in Section 5 a numerical study of these estimators through simulations under various experimental conditions, including independent and conditionally independent censoring and correct and incorrect model specifications. We conclude in Section 6 with practical recommendations on estimator selection, based on criteria such as asymptotic behavior, computational complexity, and efficiency.

1.4 Notations

We provide in Table 2 a summary of the notation used throughout the paper.

Table 2: Summary of the notations.

Symbol	Description
X	Covariates
A	Treatment indicator ($A = 1$ for treatment, $A = 0$ for control)
T	Survival time
$T(a), a \in \{0, 1\}$	Potential survival time respectively with and without treatment
$S^{(a)}, a \in \{0, 1\}$	Survival function $S^{(a)}(t) = \mathbb{P}(T(a) > t)$ of the potential survival times
$\lambda^{(a)}, a \in \{0, 1\}$	Hazard function $\lambda^{(a)}(t) = \lim_{h \rightarrow 0^+} \mathbb{P}(T(a) \in [t, t+h] T(a) > t) / h$ of the potential survival times
$\Lambda^{(a)}, a \in \{0, 1\}$	Cumulative hazard function of the potential survival times
C	Censoring time
S_C	Survival function $S_C(t) = \mathbb{P}(C > t)$ of the censoring time
G	Left-limit of the survival function $G(t) = \mathbb{P}(C > t)$ of the censoring time
\tilde{T}	Observed time ($T \wedge C$)
Δ	Censoring status $\mathbb{I}\{T < C\}$
Δ^τ	Censoring status of the restricted time $\Delta^\tau = \max\{\Delta, \mathbb{I}\{\tilde{T} > \tau\}\}$
$\{t_1, t_2, \dots, t_K\}$	Discrete times
$e(x)$	Propensity score $\mathbb{E}[A X = x]$
$\mu(x, a), a \in \{0, 1\}$	$\mathbb{E}[T \wedge \tau X = x, A = a]$
$S(t x, a), a \in \{0, 1\}$	Conditional survival function, $\mathbb{P}(T > t X = x, A = a)$.
$\lambda^{(a)}(t x), a \in \{0, 1\}$	Conditional hazard functions of the potential survival times
$G(t x, a), a \in \{0, 1\}$	left-limit of the conditional survival function of the censoring $\mathbb{P}(C > t X = x, A = a)$
$Q_S(t x, a), a \in \{0, 1\}$	$\mathbb{E}[T \wedge \tau X = x, A = a, T \wedge \tau > t]$

2 Causal survival analysis in Randomized Controlled Trials

Randomized Controlled Trials (RCTs) are the gold standard for establishing the effect of a treatment on an outcome, because treatment allocation is controlled through randomization, which ensures (asymptotically) the balance of covariates between treated and controls, and thus avoids problems of confounding between treatment groups. The core assumption in a classical RCT is the random assignment of the treatment (D. B. Rubin 1974).

Assumption. (Random treatment assignment) There holds:

$$A \perp\!\!\!\perp T(0), T(1), X \quad (5)$$

We also assume that there is a positive probability of receiving the treatment, which we rephrase under the following assumption.

Assumption. (Trial positivity)

$$0 < \mathbb{P}(A = 1) < 1 \quad (6)$$

Under Assumptions 5 and 6, classical causal identifiability equations can be written to express θ_{RMST} without potential outcomes.

$$\begin{aligned} \theta_{\text{RMST}} &= \mathbb{E}[T(1) \wedge \tau - T(0) \wedge \tau] \\ &= \mathbb{E}[T(1) \wedge \tau | A = 1] - \mathbb{E}[T(0) \wedge \tau | A = 0] \quad (\text{Random treatment assignment}) \\ &= \mathbb{E}[T \wedge \tau | A = 1] - \mathbb{E}[T \wedge \tau | A = 0]. \quad (\text{SUTVA}) \end{aligned} \quad (7)$$

However, Equation 7 still depends on T , which remains only partially observed due to censoring. To ensure that censoring does not compromise the identifiability of treatment effects, we must impose certain assumptions on the censoring process, standards in survival analysis, namely, independent censoring and conditionally independent censoring. These assumptions lead to different estimation approaches. We focus on two strategies: those that aim to directly estimate $E[T \wedge \tau | A = a]$ directly (e.g., through censoring-unbiased transformations, see Section 2.2), and those that first estimate the survival curves to derive RMST via Equation 2 (such as the Kaplan-Meier estimator and all its variants, see the next Section).

2.1 Independent censoring: the Kaplan-Meier estimator

In a first approach, one might assume that the censoring times are independent from the rest of the variables.

Assumption. (Independent censoring)

$$C \perp\!\!\!\perp T(0), T(1), X, A. \quad (8)$$

Under Equation 8, subjects censored at time t are representative of all subjects who remain at risk at time t . Figure 2 represents the causal graph when the study is randomized and outcomes are observed under independent censoring.

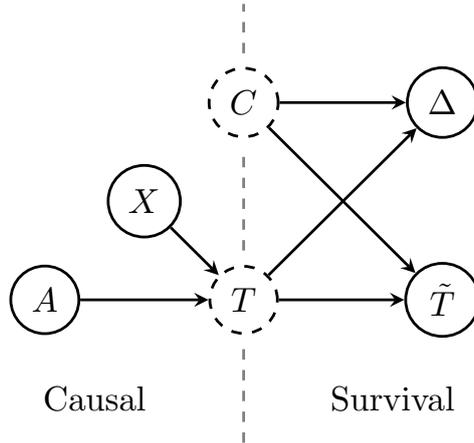


Figure 2: Causal graph in RCT survival data with independent censoring.

We also assume that there is no almost-sure upper bound on the censoring time before τ , which we rephrase under the following assumption.

Assumption. (Positivity of the censoring process) There exists $\varepsilon > 0$ such that

$$G(t) \geq \varepsilon \quad \text{for all } t \in [0, \tau). \quad (9)$$

If indeed it was the case that $\mathbb{P}(C < t) = 1$ for some $t < \tau$, then we would not be able to infer anything on the survival function on the interval $[t, \tau]$ as all observation times \tilde{T}_i would be in $[0, t]$ almost surely. In practice, adjusting the threshold time τ can help satisfy the positivity assumption. For instance, in a clinical study, if a subgroup of patients has zero probability of remaining uncensored at a given time, τ can be modified to ensure that participants have a feasible chance of remaining uncensored up to the revised threshold.

The two Assumptions 8 and 9 together allow the distributions of $T(a)$ to be identifiable, in the sense that there exists an identity that expresses $S^{(a)}$ as a function of the joint distribution of $(\tilde{T}, \Delta, A = a)$,

see for instance Ebrahimi, Molefe, and Ying (2003) for such a formula in a non-causal framework. This enables several estimation strategies, the most well-known being the Kaplan-Meier product-limit estimator.

To motivate the definition of the latter and explicit the identifiability identity, we set the analysis in the discrete case. We let $\{t_k\}_{k=1}$ be a set of positive and increasing times and assume that $T \in \{t_k\}_{k=1}$ almost surely. Then for any $t \in [0, \tau]$, it holds, letting $t_0 = 0$ by convention, thanks to Equation 4,

$$\begin{aligned} S(t|A = a) &= \mathbb{P}(T > t|A = a) = \prod_{t_k \leq t} (1 - \mathbb{P}(T = t_k|T > t_{k-1}, A = a)) \\ &= \prod_{t_k \leq t} \left(1 - \frac{\mathbb{P}(T = t_k, A = a)}{\mathbb{P}(T \geq t_k, A = a)} \right). \end{aligned}$$

Using Assumptions 8 and 9, we find that

$$\frac{\mathbb{P}(T = t_k, A = a)}{\mathbb{P}(T \geq t_k, A = a)} = \frac{\mathbb{P}(T = t_k, C = t_k, A = a)}{\mathbb{P}(T \geq t_k, C = t_k, A = a)} = \frac{\mathbb{P}(\tilde{T} = t_k, \Delta = 1, A = a)}{\mathbb{P}(\tilde{T} \geq t_k, A = a)}, \quad (10)$$

yielding the final identity

$$S(t|A = a) = \prod_{t_k \leq t} \left(1 - \frac{\mathbb{P}(\tilde{T} = t_k, \Delta = 1, A = a)}{\mathbb{P}(\tilde{T} \geq t_k, A = a)} \right). \quad (11)$$

Notice that the right hand side only depends on the distribution of the observed tuple (A, \tilde{T}, Δ) . This last equation suggests in turn to introduce the quantities

$$D_k(a) := \sum_{i=1}^n \mathbb{I}(\tilde{T}_i = t_k, \Delta_i = 1, A = a) \quad \text{and} \quad N_k(a) := \sum_{i=1}^n \mathbb{I}(\tilde{T}_i \geq t_k, A = a), \quad (12)$$

which correspond respectively to the number of deaths $D_k(a)$ and of individuals at risk $N_k(a)$ at time t_k in the treated group ($a=1$) or in the control group ($a=0$).

Definition 2.1. (Kaplan-Meier estimator, Kaplan and Meier (1958)) With $D_k(a)$ and $N_k(a)$ defined in Equation 12, we let

$$\hat{S}_{KM}(t|A = a) := \prod_{t_k \leq t} \left(1 - \frac{D_k(a)}{N_k(a)} \right). \quad (13)$$

The associated RMST estimator is then simply defined as

$$\hat{\theta}_{KM} = \int_0^\tau \hat{S}_{KM}(t|A = 1) - \hat{S}_{KM}(t|A = 0) dt. \quad (14)$$

The Kaplan-Meier estimator is the Maximum Likelihood Estimator (MLE) of the survival functions, see for instance Kaplan and Meier (1958). Furthermore, because $D_k(a)$ and $N_k(a)$ are sums of i.i.d. random variables, the Kaplan-Meier estimator inherits some convenient statistical properties.

Proposition 2.1. Under Assumptions 1, 5, 6, 8 and 9, and for all $t \in [0, \tau]$, the estimator $\hat{S}_{KM}(t|A = a)$ of $S^{(a)}(t)$ is strongly consistent and admits the following bounds for its bias:

$$0 \leq S^{(a)}(t) - \mathbb{E}[\hat{S}_{KM}(t|A = a)] \leq O(\mathbb{P}(N_k(a) = 0)),$$

where k is the greatest time t_k such that $t \leq t_k$.

Gill (1983) gives a more precise lower-bound on the bias in the case of continuous distributions, which was subsequently refined by Zhou (1988). The bound we give, although slightly looser, still exhibits the same asymptotic regime. In particular, as soon as $S^{(a)}(t) > 0$ (and Assumption 9 holds), then the bias decays exponentially fast towards 0. We give in Section 8.1 a simple proof of our bound in our context.

Proposition 2.2. *Under Assumptions 1, 5, 6, 8 and 9, and for all $t \in [0, \tau]$, $\hat{S}_{\text{KM}}(t|A = a)$ is asymptotically normal and $\sqrt{n}(\hat{S}_{\text{KM}}(t|A = a) - S^{(a)}(t))$ converges in distribution towards a centered Gaussian of variance*

$$V_{\text{KM}}(t|A = a) := S^{(a)}(t)^2 \sum_{t_k \leq t} \frac{1 - s_k(a)}{s_k(a)r_k(a)},$$

where $s_k(a) = S^{(a)}(t_k)/S^{(a)}(t_{k-1})$ and $r_k(a) = \mathbb{P}(\tilde{T} \geq t_k, A = a)$.

The proof of Proposition 2.2 can be found in Section 8.1. Because $D_k(a)/N_k(a)$ is a natural estimator of $1 - s_k(a)$ and, $\frac{1}{n}N_k(a)$ a natural estimator for $r_k(a)$, the asymptotic variance of the Kaplan-Meier estimator can be estimated with the so-called Greenwood formula, as already derived heuristically in Kaplan and Meier (1958):

$$\widehat{\text{Var}}(\hat{S}_{\text{KM}}(t|A = a)) := \hat{S}_{\text{KM}}(t|A = a)^2 \sum_{t_k \leq t} \frac{D_k(a)}{N_k(a)(N_k(a) - D_k(a))}. \quad (15)$$

We finally mention that the KM estimator as defined in Definition 2.1 still makes sense in a non-discrete setting, and one only needs to replace the fixed grid $\{t_k\}$ by the values at which we observed an event ($\tilde{T}_i = t_k, \Delta_i = 1$). We refer to Breslow and Crowley (1974) for a study of this estimator in the continuous case and to Aalen, Borgan, and Gjessing (2008), Sec 3.2 for a general study of the KM estimator through the prism of point processes.

2.2 Conditionally independent censoring

An alternative hypothesis in survival analysis that relaxes the assumption of independent censoring is conditionally independent censoring, also referred sometimes as *informative censoring*. It allows to model more realistic censoring processes, in particular in situations where there are reasons to believe that C may be dependent from A and X (for instance, if patient is more like to leave the study when treated because of side-effects of the treatment).

Assumption. (Conditionally independent censoring)

$$C \perp\!\!\!\perp T(0), T(1) \mid X, A \quad (16)$$

Under Equation 16, subjects within a same stratum defined by $X = x$ and $A = a$ have equal probability of censoring at time t , for all t . In case of conditionally independent censoring, we also need to assume that all subjects have a positive probability to remain uncensored at their time-to-event.

Assumption. (Positivity / Overlap for censoring) There exists $\varepsilon > 0$ such that for all $t \in [0, \tau]$, it almost surely holds

$$G(t|A, X) \geq \varepsilon. \quad (17)$$

Figure 3 represents the causal graph when the study is randomized with conditionally independent censoring.

Under dependent censoring, the Kaplan-Meier estimator as defined in Definition 2.1 can fail to estimate survival, in particular because Equation 10 does not hold anymore. Alternatives include plug-in G-formula estimators (Section 2.2.1) and unbiased transformations (Section 2.2.2).

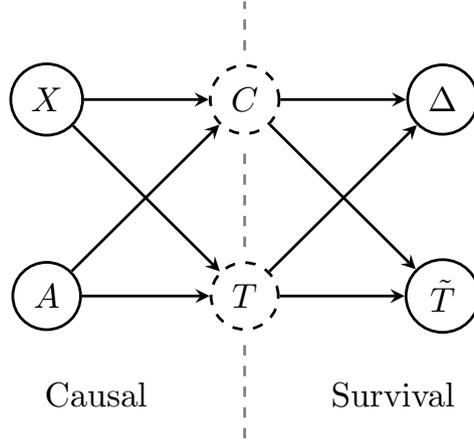


Figure 3: Causal graph in RCT survival data with dependent censoring.

2.2.1 The G-formula and the Cox Model

Because the censoring is now independent from the potential outcome conditionally to the covariates, it would seem natural to model the response of the survival time conditionally to these covariates too:

$$\mu(x, a) := \mathbb{E}[T \wedge \tau | X = x, A = a].$$

Building on Equation 7, one can express the RMST as a function of μ :

$$\theta_{\text{RMST}} = \mathbb{E} [\mathbb{E}[T \wedge \tau | X, A = 1]] - \mathbb{E} [T \wedge \tau | X, A = 0] = \mathbb{E}[\mu(X, 1) - \mu(X, 0)].$$

An estimator $\hat{\mu}$ of μ would then straightforwardly yield an estimator of the difference in RMST through the so-called *G-formula* plug-in estimator:

$$\hat{\theta}_{\text{G-formula}} = \frac{1}{n} \sum_{i=1}^n \hat{\mu}(X_i, 1) - \hat{\mu}(X_i, 0). \quad (18)$$

We would like to stress that a G-formula approach works also in an observational context as the one introduced in Section 3.2. However, because the estimation strategies presented in the next sections relies on estimating nuisance parameters, and that this latter task is often done in the same way as we estimate the conditional response μ , we decided to not delay the introduction of the G-formula any further, and we present below a few estimation methods for μ . These methods are sub-divided in two categories: *T-learners*, where $\mu(\cdot, 1)$ is estimated separately from $\mu(\cdot, 0)$, and *_S-learners*, where $\hat{\mu}$ is obtained by fitting a single model based on covariates (X, A) .

Cox's Model. There are many ways to model μ in a survival context, the most notorious of which being the Cox proportional hazards model (Cox 1972). It relies on a semi-parametric modelling the conditional hazard functions $\lambda^{(a)}(t|X)$ as

$$\lambda^{(a)}(t|X) = \lambda_0^{(a)}(t) \exp(X^\top \beta^{(a)}),$$

where $\lambda_0^{(a)}$ is a baseline hazard function and $\beta^{(a)}$ has the same dimension as the vector of covariate X . The conditional survival function then take the simple form (in the continuous case)

$$S^{(a)}(t|X) = S_0^{(a)}(t) \exp(-X^\top \beta^{(a)}),$$

where $S_0^{(a)}(t)$ is the survival function associated with $\lambda_0^{(a)}$. The vector $\beta^{(a)}$ is classically estimated by maximizing the so-called *partial likelihood* function as introduced in the original paper of Cox (1972):

$$\mathcal{L}(\beta) := \prod_{\Delta_i=1} \frac{\exp(X_i^\top \beta)}{\sum_{\tilde{T}_j \leq \tilde{T}_i} \exp(X_j^\top \beta)},$$

while the cumulative baseline hazard function can be estimated through the Breslow's estimator (Breslow 1974):

$$\hat{\Lambda}_0^{(a)}(t) = \sum_{\Delta_i=1, \tilde{T}_i \leq t} \frac{1}{\sum_{\tilde{T}_j \leq \tilde{T}_i} \exp(X_j^\top \hat{\beta}^{(a)})}$$

where $\hat{\beta}^{(a)}$ is a partial likelihood maximizer. One can show that $(\hat{\beta}^{(a)}, \hat{\Lambda}_0^{(a)})$ is the MLE of the true likelihood, when $\hat{\Lambda}_0^{(a)}$ is optimized over all step functions of the form

$$\Lambda_0(t) := \sum_{\Delta_i=1} h_i, \quad h_i \in \mathbb{R}^+.$$

This fact was already hinted in the original paper by Cox and made rigorous in many subsequent papers, see for instance Fan, Feng, and Wu (2010). Furthermore, if the true distribution follows a Cox model, then both $\hat{\beta}^{(a)}$ and $\hat{\Lambda}_0^{(a)}$ are strongly consistent and asymptotically normal estimator of the true parameters $\beta^{(a)}$ and $\Lambda^{(a)}$, see Kalbfleisch and Prentice (2002), Sec 5.7. When using a *T*-learner approach, one simply finds $(\hat{\beta}^{(a)}, \hat{\Lambda}_0^{(a)})$ for $a \in \{0, 1\}$ by considering the control group and the treated group separately. When using a *S*-learner approach, the treatment status A becomes a covariate and the model becomes

$$\lambda(t|X, A) = \lambda_0(t) \exp(X^\top \beta + \alpha A). \quad (19)$$

for some $\alpha \in \mathbb{R}$. One main advantage of Cox's model is that it makes it very easy to interpret the effect of a covariate on the survival time. If indeed $\alpha > 0$, then the treatment has a negative effect on the survival times. Likewise, if $\beta_i > 0$, then the i -th coordinate of X has a negative effect as well. We finally mention that the hazard ratio takes a particularly simple form under the later model since

$$\theta_{\text{HR}} = e^\alpha.$$

In particular, it does not depend on the time horizon τ , and is thus sometimes referred to as *proportional hazard*. Figure 4 illustrates the estimation of the difference in Restricted Mean Survival Time using G-formula with Cox models.

Weibull Model. A popular parametric model for survival is the Weibull Model, which amounts to assume that

$$\lambda^{(a)}(t|X) = \lambda_0^{(a)}(t) \exp(X^\top \beta)$$

where $\lambda_0^{(a)}(t)$ is the instant hazard function of a Weibull distribution, that is to say a function proportional to t^γ for some shape parameter $\gamma > 0$. We refer to Zhang (2016) for a study of this model.

Survival Forests. On the non-parametric front, we mention the existence of survival forests (Ishwaran et al. 2008).

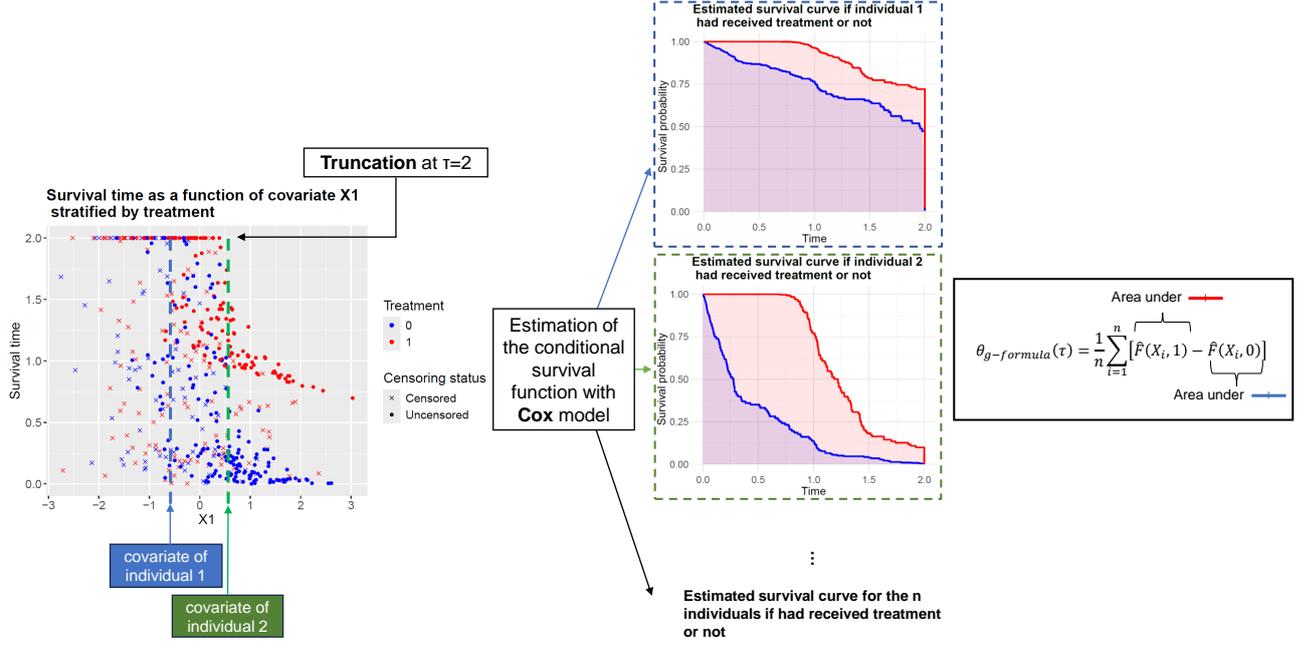


Figure 4: Illustration of the G-formula for estimating θ_{RMST} in an RCT when only one covariate X_1 influences the outcome.

2.2.2 Censoring unbiased transformations

Censoring unbiased transformations involve applying a transformation to T . Specifically, we compute a new time T^* of the form

$$T^* := T^*(\tilde{T}, X, A, \Delta) = \begin{cases} \phi_0(\tilde{T} \wedge \tau, X, A) & \text{if } \Delta^\tau = 0, \\ \phi_1(\tilde{T} \wedge \tau, X, A) & \text{if } \Delta^\tau = 1. \end{cases} \quad (20)$$

for two wisely chosen transformations ϕ_0 and ϕ_1 , and where

$$\Delta^\tau := \mathbb{I}\{T \wedge \tau \leq C\} = \Delta + (1 - \Delta)\mathbb{I}\{\tilde{T} > \tau\} \quad (21)$$

is the indicator of the event where the individual is either uncensored or censored after time τ . The idea behind the introduction of Δ^τ is that because we are only interested in computing the expectation of the survival time thresholded by τ , any censored observation coming after time τ can in fact be considered as uncensored ($\Delta^\tau = 1$).

A censoring unbiased transformation T^* shall satisfy: for $a \in \{0, 1\}$, it holds

$$\mathbb{E}[T^* | A = a, X] = \mathbb{E}[T(a) \wedge \tau | X] \quad \text{almost surely.} \quad (22)$$

A notable advantage of this approach is that it enables the use of the full transformed dataset (X_i, A_i, T_i^*) as if no censoring occurred. Because it holds

$$\mathbb{E}[\mathbb{E}[T^* | A = a, X]] = \mathbb{E} \left[\frac{\mathbb{I}\{A = a\}}{\mathbb{P}(A = a)} T^* \right], \quad (23)$$

there is a very natural way to derive an estimator of the difference in RMST from any censoring unbiased transformation T^* as:

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i}{\pi} - \frac{1 - A_i}{1 - \pi} \right) T_i^* \quad (24)$$

where $\pi = P(A = 1) \in (0, 1)$ by Assumption 6 and where $T_i^* = T^*(\tilde{T}_i, X_i, A_i, \Delta_i)$. We easily get the following result.

Proposition 2.3. *Under Assumptions 5 and 6, the estimator $\hat{\theta}$ derived as in Equation 24 from a square integrable censoring unbiased transformations satisfying Equation 22 is an unbiased, strongly consistent, and asymptotically normal estimator of the difference in RMST.*

Square integrability will be ensured any time the transformation is bounded, which will always be the case of the ones considered in this work. It is natural in a RCT setting to assume that probability of being treated π is known. If not, it is usual to replace π by its empirical counterpart $\hat{\pi} = n_1/n$ where $n_a = \sum_i 1\{A = a\}$. The resulting estimator takes the form

$$\hat{\theta} = \frac{1}{n_1} \sum_{A_i=1} T_i^* - \frac{1}{n_0} \sum_{A_i=0} T_i^*. \quad (25)$$

Note however that this estimator is slightly biased due to the estimation of π (see for instance Colnet et al. (2022), Lemma 2), but it is still strongly consistent and asymptotically normal, and its biased is exponentially small in n .

Proposition 2.4. *Under Assumptions 5 and 6, the estimator $\hat{\theta}$ derived as in Equation 25 from a square integrable censoring unbiased transformations satisfying Equation 22 is a strongly consistent, and asymptotically normal estimator of the difference in RMST.*

The two most popular transformations are Inverse-Probability-of-Censoring Weighting (Koul, Susarla, and Ryzin (1981)) and Buckley-James (Buckley and James (1979)), both illustrated in Figure 5 and detailed below. In the former, only non-censored observations are considered and they are weighted while in the latter, censored observations are imputed with an estimated survival time.

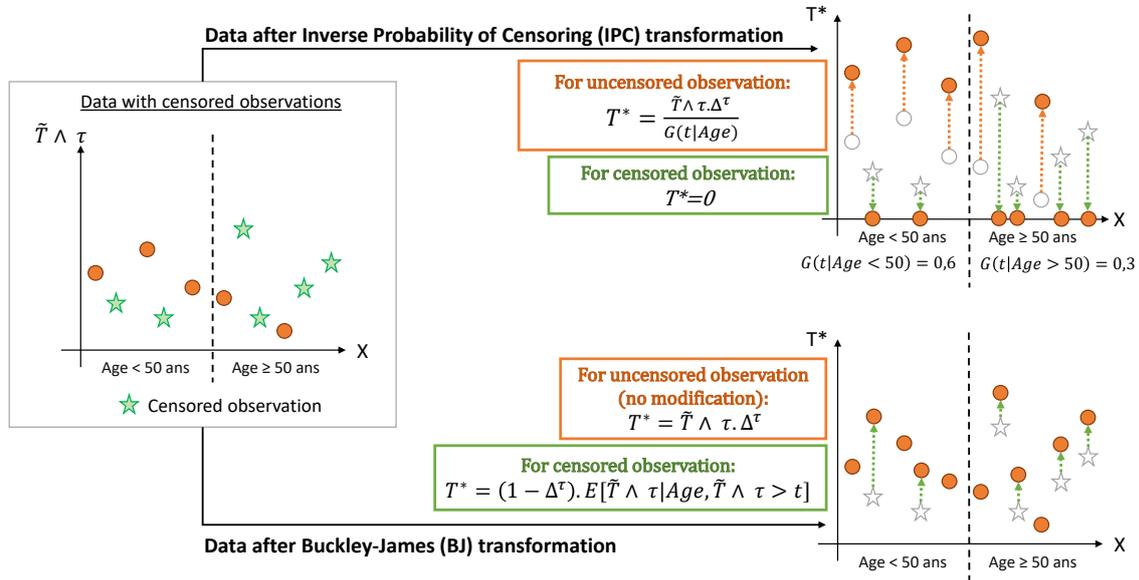


Figure 5: Illustration of Inverse-Probability-of-Censoring and Buckley-James transformations.

The Inverse-Probability-of-Censoring Weighted transformation

The Inverse-Probability-of-Censoring Weighted (IPCW) transformation, introduced by (Koul, Susarla, and Ryzin (1981)) in the context of censored linear regression, involves discarding censored observa-

tions and applying weights to uncensored data. More precisely, we let

$$T_{\text{IPCW}}^* = \frac{\Delta^\tau}{G(\tilde{T} \wedge \tau | X, A)} \tilde{T} \wedge \tau, \quad (26)$$

where we recall that $G(t|X, A) := \mathbb{P}(C \leq t | X, A)$ is the left limit of the conditional survival function of the censoring. This estimator assigns higher weights to uncensored subjects within a covariate group that is highly prone to censoring, thereby correcting for conditionally independent censoring and reducing selection bias (Howe et al. 2016).

Proposition 2.5. *Under Assumptions 1, 5, 6, 16 and 17, the IPCW transform 26 is a censoring unbiased transformation in the sense of Equation 22.*

The proof of Proposition 2.5 is in Section 8.2. The IPCW depends on the unknown conditional survival function of the censoring $G(\cdot | X, A)$, which thus needs to be estimated. Estimating conditional censoring or the conditional survival function can be approached similarly, as both involve estimating a time—whether for survival or censoring. Consequently, we can use semi-parametric methods, such as the Cox model, or non-parametric approaches like survival forests, and we refer to Section 2.2.1 for a development on these approaches. Once an estimator $\hat{G}(\cdot | A, X)$ of the later is provided, one can construct an estimator of the difference in RMST based on Equation 24 or Equation 25

$$\hat{\theta}_{\text{IPCW}} = \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i}{\pi} - \frac{1 - A_i}{1 - \pi} \right) T_{\text{IPCW},i}^* \quad (27)$$

or

$$\hat{\theta}_{\text{IPCW}} = \frac{1}{n_1} \sum_{A_i=1} T_{\text{IPCW},i}^* - \frac{1}{n_0} \sum_{A_i=0} T_{\text{IPCW},i}^* \quad (28)$$

where we recall that $n_a := \#\{i \in [n] \mid A_i = a\}$. By Proposition 2.3, Proposition 2.4 and Proposition 2.5, we easily deduce that $\hat{\theta}_{\text{IPCW}}$ is asymptotically consistent as soon as \hat{G} is.

Corollary 2.1. *Under Assumptions 1, 5, 6, 16 and 17, if \hat{G} is uniformly weakly (resp. strongly) consistent then so is $\hat{\theta}_{\text{IPCW}}$, either as in defined in Equation 27 or in Equation 28.*

This result simply comes from the fact that $\hat{\theta}_{\text{IPCW}}$ depends continuously on \hat{G} and that G is lower-bounded (Assumption 17). Surprisingly, we found limited use of this estimator in the literature, beside its first introduction in Koul, Susarla, and Ryzin (1981). This could potentially be explained by the fact that, empirically, we observed that this estimator is highly variable. Consequently, we do not explore its properties further and will not include it in the numerical experiments. A related and more popular estimator is the IPCW-Kaplan-Meier, defined as follows.

Definition 2.2. (IPCW-Kaplan-Meier) We let again $\hat{G}(\cdot | X, A)$ be an estimator of the (left limit of) the conditional censoring survival function and we introduce

$$D_k^{\text{IPCW}}(a) := \sum_{i=1}^n \frac{\Delta_i^\tau}{\hat{G}(\tilde{T}_i \wedge \tau | X_i, A = a)} \mathbb{I}(\tilde{T}_i = t_k, A_i = a)$$

and

$$N_k^{\text{IPCW}}(a) := \sum_{i=1}^n \frac{\Delta_i^\tau}{\hat{G}(\tilde{T}_i \wedge \tau | X_i, A = a)} \mathbb{I}(\tilde{T}_i \geq t_k, A_i = a),$$

be the weight-corrected numbers of deaths and of individual at risk at time t_k . The IPCW version of the KM estimator is then defined as:

$$\hat{S}_{\text{IPCW}}(t | A = a) = \prod_{t_k \leq t} \left(1 - \frac{D_k^{\text{IPCW}}(a)}{N_k^{\text{IPCW}}(a)} \right).$$

Note that the quantity π is not present in the definition of $D_k^{\text{IPCW}}(a)$ and $N_k^{\text{IPCW}}(a)$ because it would simply disappear in the ratio $D_k^{\text{IPCW}}(a)/N_k^{\text{IPCW}}(a)$. The subsequent RMST estimator then take the form

$$\hat{\theta}_{\text{IPCW-KM}} = \int_0^\tau \hat{S}_{\text{IPCW}}(t|A=1) - \hat{S}_{\text{IPCW}}(t|A=0) dt. \quad (29)$$

Like before for the classical KM estimator, this new reweighted KM estimator enjoys good statistical properties.

Proposition 2.6. *Under Assumptions 1, 5, 6, 16 and 17, and for all $t \in [0, \tau]$, the oracle estimator $S_{\text{IPCW}}^*(t|A=a)$ defined as in Definition 2.2 with $\hat{G} = G$ is a strongly consistent and asymptotically normal estimator of $S^{(a)}(t)$.*

The proof of Proposition 2.6 can be found in Section 8.2. Because the evaluation of $N_k^{\text{IPCW}}(a)$ now depends on information gathered after time t_k (through the computation of the weights), the previous proofs on the absence of bias and on the derivation of the asymptotic variance unfortunately do not carry over in this case. Whether its bias is exponentially small and whether the asymptotic variance can be derived in a closed form are questions left open. We are also not aware of any estimation schemes for the asymptotic variance in this context. In the case where we do not have access to the oracle survival function G , we can again still achieve consistency if the estimator $\hat{G}(\cdot|A, X)$ that we provide is consistent.

Corollary 2.2. *Under Assumptions 1, 5, 16 and 17, if \hat{G} is uniformly weakly (resp. strongly) consistent then so is $\hat{S}_{\text{IPCW}}(t|A=a)$.*

This corollary ensures that the corresponding RMST estimator defined in Equation 29 will be consistent as well.

The Buckley-James transformation

One weakness of the IPCW transformation is that it discards all censored data. The Buckley-James (BJ) transformation, introduced by (Buckley and James (1979)), takes a different path by leaving all uncensored values untouched, while replacing the censored ones by an extrapolated value. Formally, it is defined as follows:

$$T_{\text{BJ}}^* = \Delta^\tau(\tilde{T} \wedge \tau) + (1 - \Delta^\tau)Q_S(\tilde{T} \wedge \tau|X, A), \quad (30)$$

where, for $t \leq \tau$,

$$Q_S(t|X, A) := \mathbb{E}[T \wedge \tau|X, A, T \wedge \tau > t] = \frac{1}{S(t|X, A)} \int_t^\tau S(u|X, A) du$$

where $S(t|X, A = a) := \mathbb{P}(T(a) > t|X)$ is the conditional survival function. We refer again to Figure 5 for a diagram of this transformation.

Proposition 2.7. *Under Assumptions 1, 5, 16 and 17, the BJ transform 30 is a censoring unbiased transformation in the sense of Equation 22.*

The proof of Proposition 2.7 can be found in Section 8.2. Again, the BJ transformation depends on a nuisance parameter (here $Q_S(\cdot|X, A)$) that needs to be estimated. We again refer to Section 2.2.1 for a brief overview of possible estimation strategies for Q_S . Once provided with an estimator $\hat{Q}_S(\cdot|A, X)$, a very natural estimator of the RMST based on the BJ transformation and on Equation 24 or Equation 25 would be

$$\hat{\theta}_{\text{BJ}} = \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i}{\pi} - \frac{1 - A_i}{1 - \pi} \right) T_{\text{BJ},i}^* \quad (31)$$

or

$$\hat{\theta}_{\text{BJ}} = \frac{1}{n_1} \sum_{A_i=1} T_{\text{BJ},i}^* - \frac{1}{n_0} \sum_{A_0=1} T_{\text{BJ},i}^* \quad (32)$$

Like for the IPCW transformation, the BJ transformation yields a consistent estimate of the RMST as soon as the model is well-specified.

Corollary 2.3. *Under Assumptions 1, 5, 16 and 17, if \hat{Q}_S is uniformly weakly (resp. strongly) consistent then so is $\hat{\theta}_{\text{BJ}}$ defined as in Equation 31 or Equation 32.*

The proof is again a mere application of Propositions 2.3, 2.4 and 2.7, and relies on the continuity of $S \mapsto Q_S$. The BJ transformation is considered as the best censoring transformation of the original response in the following sense.

Theorem 2.1. *For any transformation T^* of the form 20, it holds*

$$\mathbb{E}[(T_{\text{BJ}}^* - T \wedge \tau)^2] \leq \mathbb{E}[(T^* - T \wedge \tau)^2].$$

This result is stated in Fan and Gijbels (1994) but without a proof. We detail it in Section 8.2 for completeness.

2.2.3 Augmented corrections

The main disadvantage of the two previous transformations is that they heavily rely on the specification of good estimator \hat{G} (for IPCW) or \hat{S} (for BJ). In order to circumvent this limitation, D. Rubin and Laan (2007) proposed the following transformations, whose expression is based on theory of semi-parametric estimation developed in Laan and Robins (2003),

$$T_{\text{DR}}^* = \frac{\Delta^\tau \tilde{T} \wedge \tau}{G(\tilde{T} \wedge \tau | X, A)} + \frac{(1 - \Delta^\tau) Q_S(\tilde{T} \wedge \tau | X, A)}{G(\tilde{T} \wedge \tau | X, A)} - \int_0^{\tilde{T} \wedge \tau} \frac{Q_S(t | X, A)}{G(t | X, A)^2} d\mathbb{P}_C(t | X, A), \quad (33)$$

where $d\mathbb{P}_C(t | X, A)$ is the distribution of C conditional on the covariates X and A . We stress that this distribution is entirely determined by the $G(\cdot | X, A)$, so that this transformation only depends on the knowledge of both conditional survival functions G and S , will be thus sometimes denoted $T_{\text{DR}}^*(G, S)$ to underline this dependency. This transformation is not only a censoring unbiased transform in the sense of Equation 22, but is also doubly robust in the following sense.

Proposition 2.8. *We let F, R be two conditional survival functions. Under Assumptions 1, 5, 6, 16 and 17, if F also satisfies Assumption 17, and if $F(\cdot | X, A)$ is absolutely continuous wrt $G(\cdot | X, A)$, then the transformation $T_{\text{DR}}^* = T_{\text{DR}}^*(F, R)$ is a censoring unbiased transformation in the sense of Equation 22 whenever $F = G$ or $R = S$.*

The statement and proof of this results is found in D. Rubin and Laan (2007) in the case where C and T are continuous. A careful examination of the proofs show that the proof translates straight away to our discrete setting.

3 Causal survival analysis in observational studies

Unlike RCT, observational data — such as from registries, electronic health records, or national healthcare systems — are collected without controlled randomized treatment allocation. In such cases, treated and control groups are likely unbalanced due to the non-randomized design, which results in the treatment effect being confounded by variables influencing both the survival outcome T and the treatment allocation A . To enable identifiability of the causal effect, additional standard assumptions are needed.

Assumption. (Conditional exchangeability / Unconfoundedness) It holds

$$A \perp\!\!\!\perp T(0), T(1)|X \quad (34)$$

Under Equation 34, the treatment assignment is randomly assigned conditionally on the covariates X . This assumption ensures that there are no unmeasured confounder as the latter would make it impossible to distinguish correlation from causality.

Assumption. (Positivity / Overlap for treatment) Letting $e(X) := \mathbb{P}(A = 1|X)$ be the *propensity score*, there holds

$$0 < e(X) < 1 \quad \text{almost surely.} \quad (35)$$

The Equation 35 assumption requires adequate overlap in covariate distributions between treatment groups, meaning every observation must have a non-zero probability of being treated. Because Assumption 5 does not hold anymore, neither does the previous identifiability Equation 7. In this new context, we can write

$$\begin{aligned} \theta_{\text{RMST}} &= \mathbb{E}[T(1) \wedge \tau - T(0) \wedge \tau] \\ &= \mathbb{E} \left[\mathbb{E}[T(1) \wedge \tau | X] - \mathbb{E}[T(0) \wedge \tau | X] \right] \\ &= \mathbb{E} \left[\mathbb{E}[T(1) \wedge \tau | X, A = 1] - \mathbb{E}[T(0) \wedge \tau | X, A = 0] \right] \quad (\text{unconfoundedness}) \\ &= \mathbb{E} \left[\mathbb{E}[T \wedge \tau | X, A = 1] - \mathbb{E}[T \wedge \tau | X, A = 0] \right]. \quad (\text{SUTVA}) \end{aligned} \quad (36)$$

In another direction, one could wish to identify the treatment effect through the survival curve as in Equation 2:

$$S^{(a)}(t) = \mathbb{P}(T(a) > t) = \mathbb{E} \left[\mathbb{P}(T > t | X, A = a) \right]. \quad (37)$$

Again, both identities still depend on the unknown quantity T and suggest two different estimation strategies. These strategies differ according to the censoring assumptions and are detailed below.

3.1 Independent censoring

Figure 6 illustrates a causal graph in observational survival data with independent censoring (Assumption 8).

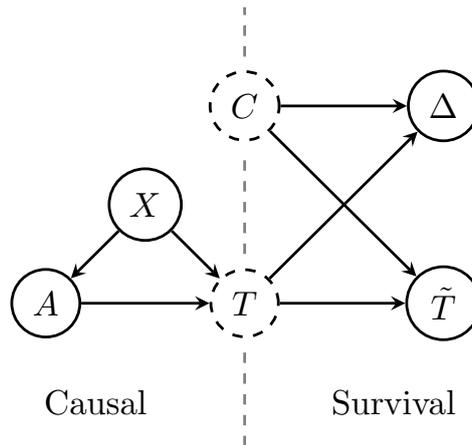


Figure 6: Causal graph in observational survival data with independent censoring.

Under Assumption 8, we saw in Section 2.1 that the Kaplan-Meier estimator could conveniently handle censoring. Building on Equation 37, we can write

$$S^{(1)}(t) = \mathbb{E} \left[\frac{\mathbb{E}[\mathbb{I}\{A = 1, T > t\}|X]}{\mathbb{E}[\mathbb{I}\{A = 1\}|X]} \right] = \mathbb{E} \left[\frac{\mathbb{A}\mathbb{I}\{T > t\}}{e(X)} \right],$$

which suggests to adapt the classical KM estimator by reweighting it by the propensity score. The use of propensity score in causal inference has been initially introduced by Rosenbaum and Rubin (1983) and further developed in Hirano, Imbens, and Ridder (2003). It was extended to survival analysis by Xie and Liu (2005) through the adjusted Kaplan-Meier estimator as defined below.

Definition 3.1. (IPTW Kaplan-Meier estimator) We let $\hat{e}(\cdot)$ be an estimator of the propensity score $e(\cdot)$. We introduce

$$D_k^{\text{IPTW}}(a) := \sum_{i=1}^n \left(\frac{a}{\hat{e}(X_i)} + \frac{1-a}{1-\hat{e}(X_i)} \right) \mathbb{I}(\tilde{T}_i = t_k, \Delta_i = 1, A_i = a)$$

and $N_k^{\text{IPTW}}(a) := \sum_{i=1}^n \left(\frac{a}{\hat{e}(X_i)} + \frac{1-a}{1-\hat{e}(X_i)} \right) \mathbb{I}(\tilde{T}_i \leq t_k, A_i = a),$

be the numbers of deaths and of individual at risk at time t_k , reweighted by the propensity score. The Inverse-Probability-of-Treatment Weighting (IPTW) version of the KM estimator is then defined as:

$$\hat{S}_{\text{IPTW}}(t|A = a) = \prod_{t_k \leq t} \left(1 - \frac{D_k^{\text{IPTW}}(a)}{N_k^{\text{IPTW}}(a)} \right). \quad (38)$$

We let $S_{\text{IPTW}}^*(t|A = a)$ be the oracle KM-estimator defined as above with $\hat{e}(\cdot) = e(\cdot)$. Although the reweighting slightly changes the analysis, the good properties of the classical KM carry on to this setting.

Proposition 3.1. *Under Assumptions 1, 34, 35, 8 and 9 The oracle IPTW Kaplan-Meier estimator $S_{\text{IPTW}}^*(t|A = a)$ is a strongly consistent and asymptotically normal estimator of $S^{(a)}(t)$.*

The proof of this result simply relies again on the law of large number and the δ -method and can be found in Section 8.3. Because now S_{IPTW}^* is a continuous function of $e(\cdot)$, and because e and $1 - e$ are lower-bounded as per Assumptions 35, we easily derive the following corollary.

Corollary 3.1. *Under the same assumptions as Proposition 3.1, if $\hat{e}(\cdot)$ satisfies $\|\hat{e} - e\|_\infty \rightarrow 0$ almost surely (resp. in probability), then the IPTW Kaplan-Meier estimator $\hat{S}_{\text{IPTW}}(t|A = a)$ is a strongly (resp. weakly) consistent estimator of $S^{(a)}(t)$.*

The resulting RMST estimator simply takes the form:

$$\hat{\theta}_{\text{IPTW-KM}} = \int_0^\tau \hat{S}_{\text{IPTW}}(t|A = 1) - \hat{S}_{\text{IPTW}}(t|A = 0) dt. \quad (39)$$

which will be consistent under the same Assumptions as the previous Corollary. Note that, we are not aware of any formal results concerning the bias and the asymptotic variance of the oracle estimator $S_{\text{IPTW}}^*(t|A = a)$, and we refer to Xie and Liu (2005) for heuristics concerning these questions.

3.2 Conditional independent censoring

Under Assumptions 34 (uncounfoundeness) and 16 (conditional independent censoring), the causal effect is affected both by confounding variables and by conditional censoring. The associated causal graph is depicted in Figure 7. In this setting, one can still use the G-formula exactly as in Section 2.2.1.

A natural alternative approach is to weight the IPCW and BJ transformations from Section 2.2.2 by the propensity score to disentangle both confounding effects and censoring at the same time.

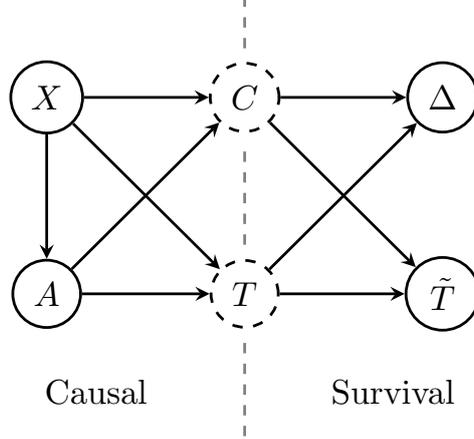


Figure 7: Causal graph in observational survival data with dependent censoring.

We mention that the G-formula approach developed in Section 2.2.1 does work in that context. In particular, Chen and Tsiatis (2001) prove consistency and asymptotic normality results for Cox estimators in an observational study, and they give an explicit formulation of the asymptotic variance as a function of the parameters of the Cox model. In the non-parametric world, Foster, Taylor, and Ruberg (2011) and Künzel et al. (2019) empirically study this estimator using survival forests, with the former employing it as a T-learner (referred to as *Virtual Twins*) and the latter as an S-learner.

3.2.1 IPTW-IPCW transformations

One can check that the IPCW transformation as introduced in Equation 26 is also a censoring unbiased transformation in that context.

Proposition 3.2. *Under Assumptions 1, 34, 35, 16 and 17, the IPTW-IPCW transform 26 is a censoring unbiased transformation in the sense of Equation 22.*

The proof of Proposition 3.2 can be found in Section 8.4. Deriving an estimator of the difference in RMST is however slightly different in that context. In particular, Equation 23 rewrites

$$\mathbb{E}[\mathbb{E}[T^*|X, A = 1]] = \mathbb{E}\left[\frac{A}{e(X)}T^*\right],$$

Which in turn suggests to define

$$\hat{\theta}_{\text{IPTW-IPCW}} = \frac{1}{n} \sum_{i=1}^n \left(\frac{A}{e(X)} - \frac{1-A}{1-e(X)} \right) T_{\text{IPCW},i}^* \quad (40)$$

This transformation now depends on two nuisance parameters, namely the conditional survival function of the censoring (through T_{IPCW}^*) and the propensity score. Once estimators of these quantities are provided, one could look at the corresponding quantity computed with these estimators.

Proposition 3.3. *Under Assumptions 1, 34, 35, 16 and 17, and if $\widehat{G}(\cdot|X, A)$ and $\widehat{e}(\cdot)$ are uniformly weakly (resp. strongly) consistent estimators, then estimator 40 defined with \widehat{e} and \widehat{G} is a weakly (resp. strongly) consistent estimator of the difference in RMST.*

The proof of Proposition 3.3 can be found in Section 8.4. We can also use the same strategy as for the IPCW transformation and incorporate the new weights into a Kaplan-Meier estimator.

Definition 3.2. (IPTW-IPCW-Kaplan-Meier) We let again $\widehat{G}(\cdot|X, A)$ and $\widehat{e}(\cdot)$ be estimators of the conditional censoring survival function and of the propensity score. We introduce

$$D_k^{\text{IPTW-IPCW}}(a) := \sum_{i=1}^n \left(\frac{A_i}{\widehat{e}(X_i)} + \frac{1 - A_i}{1 - \widehat{e}(X_i)} \right) \frac{\Delta_i^\tau}{\widehat{G}(\tilde{T}_i \wedge \tau | X_i, A = a)} \mathbb{I}(\tilde{T}_i = t_k, A_i = a)$$

and

$$N_k^{\text{IPTW-IPCW}}(a) := \sum_{i=1}^n \left(\frac{A_i}{\widehat{e}(X_i)} + \frac{1 - A_i}{1 - \widehat{e}(X_i)} \right) \frac{\Delta_i^\tau}{\widehat{G}(\tilde{T}_i \wedge \tau | X_i, A = a)} \mathbb{I}(\tilde{T}_i \geq t_k, A_i = a),$$

be the weight-corrected numbers of deaths and of individual at risk at time t_k . The IPTW-IPCW version of the KM estimator is then defined as:

$$\widehat{S}_{\text{IPTW-IPCW}}(t|A = a) = \prod_{t_k \leq t} \left(1 - \frac{D_k^{\text{IPTW-IPCW}}(a)}{N_k^{\text{IPTW-IPCW}}(a)} \right).$$

The difference in RMST estimated with IPTW-IPCW-Kaplan-Meier survival curves is then simply as

$$\widehat{\theta}_{\text{IPTW-IPCW-KM}} = \int_0^\tau \widehat{S}_{\text{IPTW-IPCW}}(t|A = 1) - \widehat{S}_{\text{IPTW-IPCW}}(t|A = 0) dt. \quad (41)$$

Proposition 3.4. *Under Assumptions 1, 34, 35, 16 and 17, and for all $t \in [0, \tau]$, if the oracle estimator $S_{\text{IPTW-IPCW}}^*(t|A = a)$ defined as in Definition 3.2 with $\widehat{G}(\cdot|A, X) = G(\cdot|A, X)$ and $\widehat{e} = e$ is a strongly consistent and asymptotically normal estimator of $S^{(a)}(t)$.*

The proof of Proposition 3.4 can be found in Section 8.4. Under consistency of the estimators of the nuisance parameters, the previous proposition implies that this reweighted Kaplan-Meier is a consistent estimator of the survival curve, which in turn implies consistency of $\widehat{\theta}_{\text{IPTW-IPCW-KM}}$.

Corollary 3.2. *Under Assumptions 1, 34, 35, 16 and 17, and for all $t \in [0, \tau]$, if the nuisance estimators $\widehat{G}(\cdot|A, X)$ and \widehat{e} are weakly (resp. strongly) uniformly consistent, then $\widehat{S}_{\text{IPTW-IPCW}}(t|A = a)$ is a weakly (resp. strongly) consistent estimator of $S^{(a)}(t)$.*

We are not aware of general formula for the asymptotic variances in this context. We mention nonetheless that Schaubel and Wei (2011) have been able to derive asymptotic laws in this framework in the particular case of Cox-models.

3.2.2 IPTW-BJ transformations

Like IPCW transformation, BJ transformation is again a censoring unbiased transformation in an observational context.

Proposition 3.5. *Under Assumptions 1, 34, 35, 16 and 17, the IPTW-BJ transform 30 is a censoring unbiased transformation in the sense of Equation 22.*

The proof of Proposition 3.5 can be found in Section 8.4. The corresponding estimator of the difference in RMST is

$$\hat{\theta}_{\text{IPTW-BJ}} = \frac{1}{n} \sum_{i=1}^n \left(\frac{A}{e(X)} - \frac{1-A}{1-e(X)} \right) T_{\text{BJ},i}^* \quad (42)$$

This transformation depends on the conditional survival function S (through T_{BJ}^*) and the propensity score. Consistency of the nuisance parameter estimators implies again consistency of the RMST estimator.

Proposition 3.6. *Under Assumptions 1, 34, 35, 16 and 17, and if $\hat{S}(\cdot|X, A)$ and $\hat{e}(\cdot)$ are uniformly weakly (resp. strongly) consistent estimators, then $\hat{\theta}_{\text{IPTW-BJ}}$ defined with \hat{S} and \hat{e} is a weakly (resp. strongly) consistent estimator of the RMST.*

The proof of Proposition 3.6 can be found in Section 8.4.

3.2.3 Double augmented corrections

Building on the classical doubly-robust AIPTW estimator from causal inference (Robins, Rotnitzky, and Zhao 1994), we could incorporate the doubly-robust transformations of Section 2.2.3 to obtain a *quadruply robust* transformation

$$\Delta_{\text{QR}}^* = \Delta_{\text{QR}}^*(G, S, \mu, e) := \left(\frac{A}{e(X)} - \frac{1-A}{1-e(X)} \right) (T_{\text{DR}}^*(G, S) - \mu(X, A)) + \mu(X, 1) - \mu(X, 0),$$

where we recall that T_{DR}^* is defined in Section 2.2.3. This transformation depends on four nuisance parameters: G and S through T_{DR}^* , and now the propensity score e and the conditional response μ . This transformation doesn't really fall into the scope of censoring unbiased transform, but it is easy to show that Δ_{QR}^* is quadruply robust in the following sense.

Proposition 3.7. *Let F, R be two conditional survival function, p be a propensity score, and v be a conditional response. Then, under the same assumption on F, R as in Proposition 2.8, and under Assumptions 1, 34, 35, 16 and 17, the transformations $\Delta_{\text{QR}}^* = \Delta_{\text{QR}}^*(F, R, p, v)$ satisfies, for $a \in \{0, 1\}$,*

$$\mathbb{E}[\Delta_{\text{QR}}^*|X] = \mathbb{E}[T(1) \wedge \tau - T(0) \wedge \tau|X] \quad \text{if} \quad \begin{cases} F = G & \text{or} & R = S & \text{and} \\ p = e & \text{or} & v = \mu. \end{cases}$$

This result is similar to Ozenne et al. (2020), Thm 1, and its proof can be found in Section 8.4. Based on estimators $(\hat{G}, \hat{S}, \hat{\mu}, \hat{e})$ of (G, S, μ, e) , one can then propose the following estimator of the RMST, coined the AIPTW-AIPCW estimator in Ozenne et al. (2020):

$$\begin{aligned} \hat{\theta}_{\text{AIPTW-AIPCW}} &:= \frac{1}{n} \sum_{i=1}^n \Delta_{\text{QR},i}^*(\hat{G}, \hat{S}, \hat{\mu}, \hat{e}) \\ &= \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i}{\hat{e}(X_i)} - \frac{1-A_i}{1-\hat{e}(X_i)} \right) (T_{\text{DR}}^*(\hat{G}, \hat{S})_i - \hat{\mu}(X_i, A_i)) + \hat{\mu}(X_i, 1) - \hat{\mu}(X_i, 0). \end{aligned} \quad (43)$$

This estimator enjoys good asymptotic properties under parametric models, as detailed in Ozenne et al. (2020).

4 Implementation

In this section, we first summarize the various estimators and their properties. We then provide custom implementations for all estimators, even those already available in existing packages. These

manual implementations serve two purposes: first, to make the methods accessible to the community when no existing implementation is available; and second, to facilitate a deeper understanding of the methods by detailing each step, even when a package solution exists. Finally, we present the packages available for directly computing θ_{RMST} .

4.1 Summary of the estimators

Table 3 provides an overview of the estimators introduced in this paper, along with the corresponding nuisance parameters needed for their estimation and an overview of their statistical properties in particular regarding their sensitivity to misspecification of the nuisance parameters.

Table 3: Estimators of the difference in RMST and nuisance parameters needed to compute each estimator. Empty boxes indicate that the nuisance parameter is not needed in the estimator thus misspecification has no impact. Estimators in *italic* are those that are already implemented in available packages.

Estimator	RCT	Obs	Ind Cens	Dep Cens	Outcome model	Censoring model	Treatment model	Robustness
<i>Unadjusted KM</i>	X		X					
IPCW-KM	X		X	X		G		
BJ	X		X	X	S			
<i>IPTW-KM</i>	X	X	X				e	
IPCW-IPTW-KM	X	X	X	X		G	e	
IPTW-BJ	X	X	X	X	S		e	
<i>G-formula</i>	X	X	X	X	μ			
AIPTW-AIPCW	X	X	X	X	S, μ	G	e	X (Prp 3.7)

4.2 Implementation of the estimators

Across different implementations, we use the following shared functions provided in the `utility.R` file.

- `estimate_propensity_score`: function to estimate propensity scores $e(X)$ using either parametric (i.e. logistic regression with the argument `type_of_model = "glm"`) or non-parametric methods (i.e. probability forest with the argument `type_of_model = "probability forest"` based on the `probability_forest` function from the [grf](#) (Tibshirani et al. 2017) package). This latter can include cross-fitting (`n.folds 1`).
- `estimate_survival_function`: function to estimate the conditional survival model, which supports either Cox models (argument `type_of_model = "cox"`) or survival forests (argument `type_of_model = "survival forest"`) which uses the function `survival_forest` from the [grf](#) (Tibshirani et al. 2017) package. This latter can include cross-fitting (`n.folds 1`). The estimation can be done with a single learner (argument `learner = "S-learner"`) or two learners (argument `learner = "T-learner"`).
- `estimate_hazard_function`: function to estimate the instantaneous hazard function by deriving the cumulative hazard function at each time point. This cumulative hazard function is estimated from the negative logarithm of the survival function.

- `Q_t_hat`: function to estimate the remaining survival function at all time points and for all individuals which uses the former `estimate_survival_function`.
- `Q_Y`: function to select the value of the remaining survival function from `Q_t_hat` at the specific time-to-event.
- `integral_rectangles`: function to estimate the integral of a decreasing step function using the rectangle method.
- `expected_survival`: function to estimate the integral with x, y coordinate (estimated survival function) using the trapezoidal method.
- `integrate`: function to estimate the integral at specific time points $Y.grid$ of a given integrand function which takes initially its values on `times` using the trapezoidal method.

Unadjusted Kaplan-Meier

Although Kaplan-Meier is implemented in the `survival` package (Therneau 2001), we provide a custom implementation, `Kaplan_meier_handmade`, for completeness. The difference in Restricted Mean Survival Time, estimated using Kaplan-Meier as in Equation 14 can then be calculated with the `RMST_1` function.

As an alternative, one can also use the `survfit` function in the `survival` package (Therneau 2001) for Kaplan-Meier and specify the `rmean` argument equal to τ in the corresponding summary function:

IPCW Kaplan-Meier

We first provide an `adjusted.KM` function which is then used in the `IPCW_Kaplan_meier` function to estimate the difference in RMST $\hat{\theta}_{IPCW}$ as in Equation 29. The survival censoring function $G(t|X)$ is computed with the `estimate_survival_function` utility function from the `utility.R` file.

One could also use the `survfit` function in the `survival` package (Therneau 2001) in adding IPCW weights for treated and control group and specify the `rmean` argument equal to τ in the corresponding summary function:

This alternative approach for IPCW Kaplan-Meier would also be valid for IPTW and IPTW-IPCW Kaplan-Meier.

Buckley-James based estimator

The function `BJ` estimates θ_{RMST} by implementing the Buckley-James estimator as in Equation 31. It uses two functions available in the `utility.R` file, namely `Q_t_hat` and `Q_Y`.

IPTW Kaplan-Meier

The function `IPTW_Kaplan_meier` implements the IPTW-KM estimator in Equation 39. It uses the `estimate_propensity_score` function from the `utility.R`.

G-formula

We implement two versions of the G-formula: `g_formula_T_learner` and `g_formula_S_learner`. In `g_formula_T_learner`, separate models estimate survival curves for treated and control groups, whereas `g_formula_S_learner` uses a single model incorporating both covariates and treatment status to estimate survival time. The latter approach is also available in the `RISCA` package but is limited to Cox models.

IPTW-IPCW Kaplan-Meier

The `IPTW_IPCW_Kaplan_meier` function implements the IPTW-IPCW Kaplan Meier estimator from Equation 41. It uses the utility functions from the `utility.R` file `estimate_propensity_score` and `estimate_survival_function` to estimate the nuisance parameters, and the function `adjusted.KM` which computes an adjusted Kaplan Meier estimator using the appropriate weight.

IPTW-BJ estimator

The `IPTW_BJ` implements the IPTW-BJ estimator in Equation 42. It uses the utility functions, from the `utility.R` file, `estimate_propensity_score`, `Q_t_hat` and `Q_Y` to estimate the nuisance parameters.

AIPTW-AIPCW

The `AIPTW_AIPCW` function implements the `AIPTW_AIPCW` estimator in Equation 43 using the utility function from the `utility.R` file `estimate_propensity_score`, `Q_t_hat`, `Q_Y`, and `estimate_survival_function` to estimate the nuisance parameters.

4.3 Available packages

Currently, there are few sustained implementations available for estimating RMST in the presence of right censoring. Notable exceptions include the packages `survRM2` (Hajime et al. 2015), `grf` (Tibshirani et al. 2017) and `RISCA` (Foucher, Le Borgne, and Chatton 2019). Those packages are implemented in the `utility.R` files.

SurvRM2

The difference in RMST with Unadjusted Kaplan-Meier $\hat{\theta}_{KM}$ (Equation 14) can be obtained using the function `rmst2` which takes as arguments the observed time-to-event, the status, the arm which corresponds to the treatment and τ .

RISCA

The `RISCA` package provides several methods for estimating θ_{RMST} . The difference in RMST with Unadjusted Kaplan-Meier $\hat{\theta}_{KM}$ (Equation 14) can be derived using the `survfit` function from the `survival` package (Therneau 2001) which estimates Kaplan-Meier survival curves for treated and control groups, and then the `rmst` function calculates the RMST by integrating these curves, applying the rectangle method (`type="s"`), which is well-suited for step functions.

The IPTW Kaplan-Meier (Equation 38) can be applied using the `ipw.survival` and `rmst` functions. The `ipw.survival` function requires user-specified weights (i.e. propensity scores). To streamline this process, we define the `RISCA_ipw` function, which combines these steps and utilizes the `estimate_propensity_score` from the `utility.R` file.

A single-learner version of the G-formula, as introduced in Section 2.2.1, can be implemented using the `gc.survival` function. This function requires as input the conditional survival function which should be estimated beforehand with a Cox model via the `coxph` function from the `survival` package (Therneau 2001). Specifically, the single-learner approach applies a single Cox model incorporating both covariates and treatment, rather than separate models for each treatment arm. We provide a function `RISCA_gf` that consolidates these steps.

grf

The `grf` package (Tibshirani et al. 2017) enables estimation of the difference between RMST using the Causal Survival Forest approach (Cui et al. 2023), which extends the non-parametric causal forest framework to survival data. The RMST can be estimated with the `causal_survival_forest` function, requiring covariates X , observed event times, event status, treatment assignment, and the

time horizon τ as inputs. The `average_treatment_effect` function then evaluates the treatment effect based on predictions from the fitted forest.

5 Simulations

We compare the behaviors and performances of the estimators through simulations conducted across various experimental contexts. These contexts include scenarios based on RCTs and observational data, with both independent and dependent censoring. We also use semi-parametric and non-parametric models to generate censoring and survival times, as well as logistic and nonlinear models to simulate treatment assignment.

5.1 RCT

Data Generating Process

We generate RCTs with independent censoring (Scenario 1) and conditionally independent censoring (Scenario 2). We sample n iid datapoints $(X_i, A_i, C, T_i(0), T_i(1))_{i \in [n]}$ where $T_i(0)$, $T_i(1)$ and C follows Cox's models. More specifically, we set

- $X \sim \mathcal{N}(\mu, \Sigma)$ where $\mu = (1, 1, -1, 1)$ and $\Sigma = \text{Id}_4$.
- The hazard function of $T(0)$ is

$$\lambda^{(0)}(t|X) = 0.01 \exp\{\beta_0^\top X\} \quad \text{where} \quad \beta_0 = (0.5, 0.5, -0.5, 0.5).$$

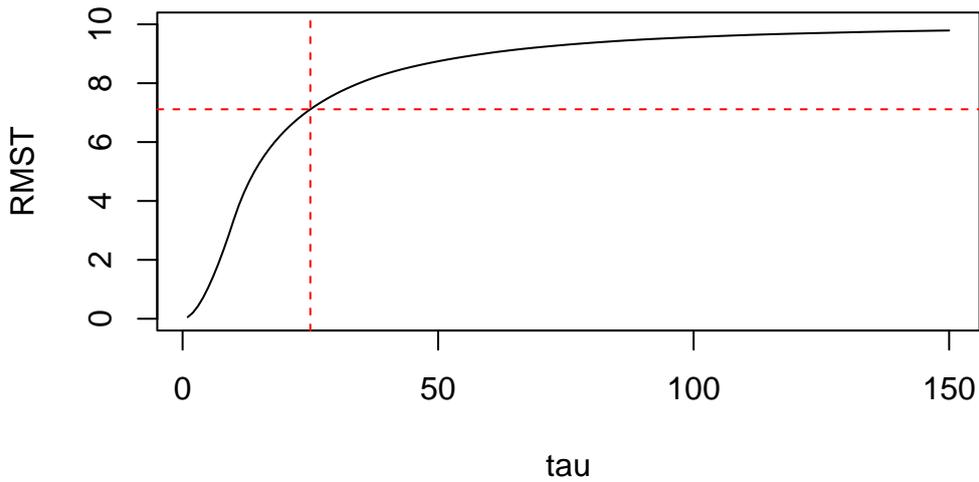
- The survival times in the treatment group are given by $T(1) = T(0) + 10$.
- The hazard function of the censoring time C is simply taken as $\lambda_C(t|X) = 0.03$ in Scenario 1, and in Scenario 2 as

$$\lambda_C(t|X) = 0.03 \cdot \exp\{\beta_C^\top X\} \quad \text{where} \quad \beta_C = (0.7, 0.7, -0.25, -0.1).$$

- The treatment allocation is independent of X : $e(X) = 0.5$.
- The threshold time τ is set to 25.

The descriptive statistics of the two datasets are displayed in Annex (Section 9.1). The graph of the difference in RMST as a function of τ for the two scenarii are displayed below; θ_{RMST} is the same in both setting.

True difference in RMST for RCT scenario 1



[1] "The ground truth for RCT scenario 1 and 2 at time 25 is 7.1"

Estimation of the RMST

For each Scenario, we estimate the difference in RMST using the methods summarized in Section 4.1. The methods used to estimate the nuisance components are indicated in brackets: either logistic regression or random forests for propensity scores and either cox models or survival random forests for survival and censoring models. A naive estimator where censored observations are simply removed and the survival time is averaged for treated and controls is also provided for a naive baseline.

Figure 8 shows the distribution of the difference in RMST for 100 simulations in Scenario 1 and different sample sizes: 500, 1000, 2000, 4000. The true value of θ_{RMST} is indicated by red dotted line.

In this setting, and in accordance with the theory, the simplest estimator (unadjusted KM) performs just as well as the others, and presents an extremely small bias (as derived in Section 2.1).

The naive estimator is biased, as expected, and the bias in both the G-formula (RISCA) and the manual G-formula S-learner arises because the treatment effect is additive $T(1) = T(0) + 10$ and violates the assumption that T would follow a Cox model in the variables (X, A) . However, $T|A = a$ is a Cox-model for $a \in \{0, 1\}$, which explain the remarkable performance of G-formula (Cox/T-learners) and some of the other models based on a Cox estimation of S .

Other estimators (IPTW KM (Reg.Log), IPCW KM (Cox), IPTW-IPCW KM (Cox & Log.Reg), IPTW-BJ (Cox & Log.Reg), AIPTW-AIPCW (Cox & Cox & Log.Reg)) involve unnecessary nuisance parameter estimates, such as propensity scores or censoring models. Despite this, their performance remains relatively stable in terms of variability, and there are roughly no differences between using (semi-)parametric or non-parametric estimation methods for nuisance parameters except for IPCW KM and IPTW-IPCW KM where there is a slight bias when using forest-based methods.

Figure 9 shows the results for the RCT simulation with conditionally independent censoring (Scenario 2). In this setting, the Naive estimator remains biased. Similarly, both the unadjusted Kaplan-Meier (KM) and its SurvRM2 equivalent, as well as the treatment-adjusted IPTW KM and its RISCA equivalent, are biased due to their failure to account for dependent censoring. As in Scenario 1, G-formula (Cox/ S-learner) and its RISCA equivalent also remain biased. The IPCW KM (Cox) is slightly biased up to 4,000 observations and quite variable due to extreme censoring probabilities. IPTW-IPCW KM (Cox & Log.Reg.) is not biased but shows high variance. In contrast, the Buckley-

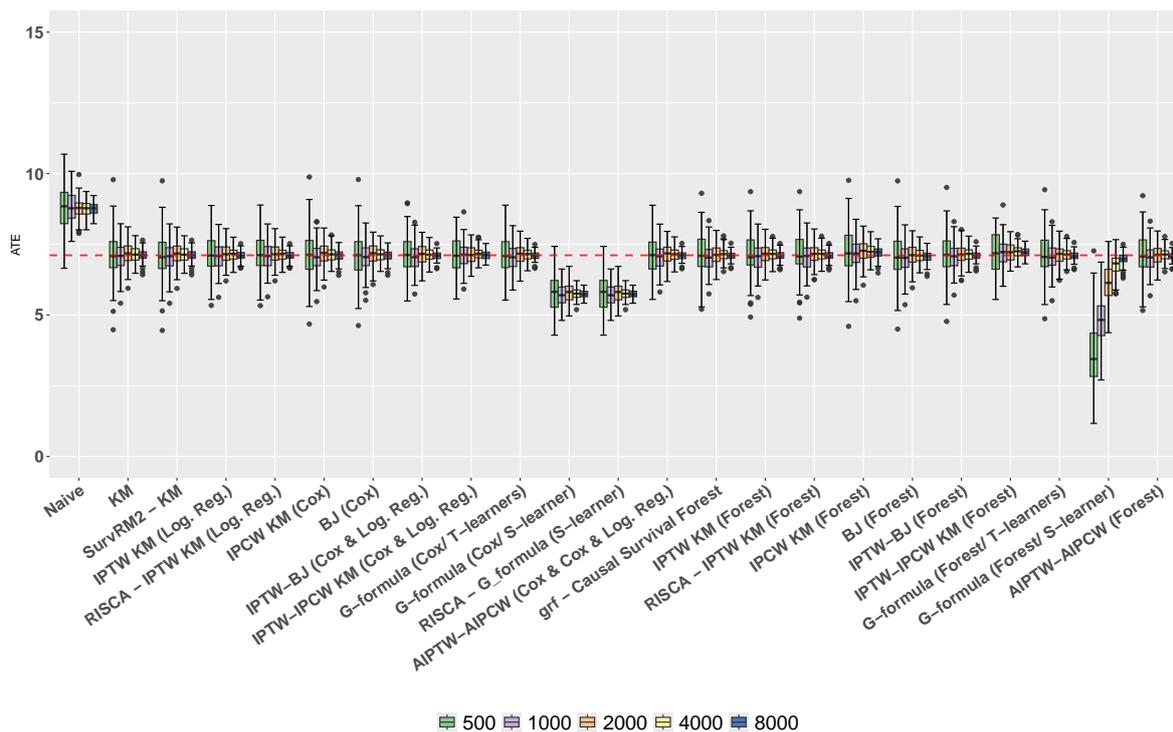


Figure 8: Results of the ATE for the simulation of a RCT with independent censoring.

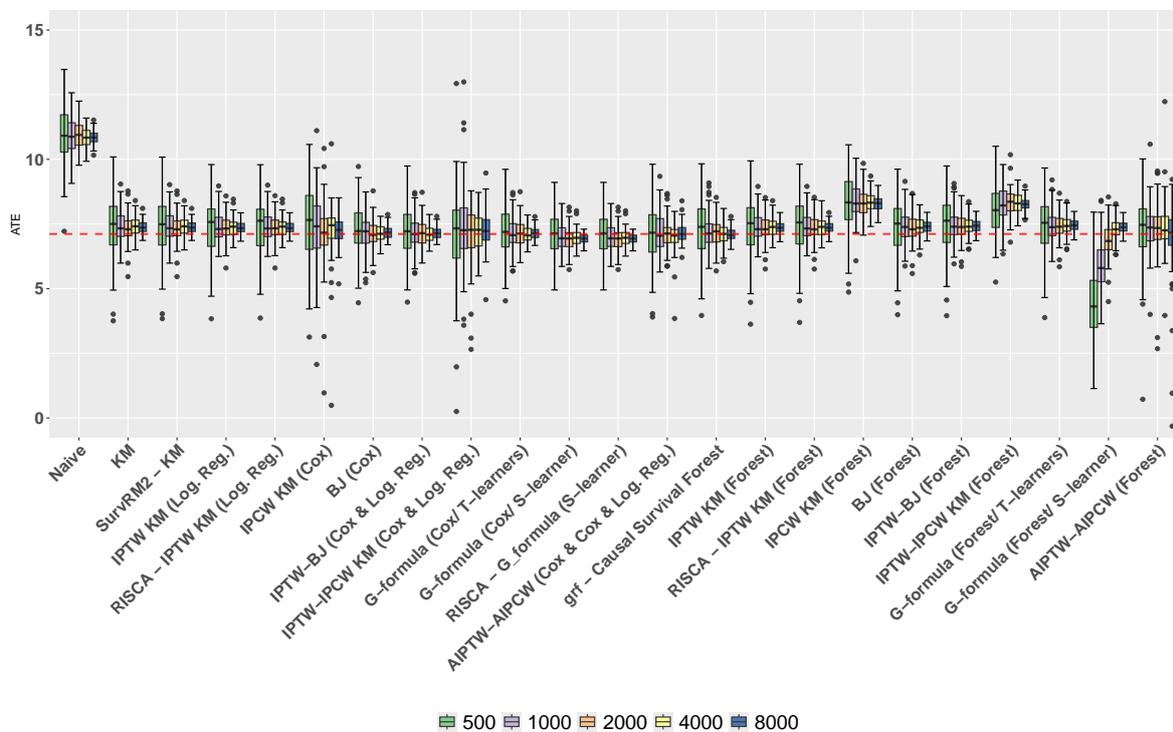


Figure 9: Estimation results of the ATE for the simulation of a RCT with dependent censoring.

James estimator BJ (Cox) is unbiased even with as few as 500 observations. The BJ estimator also demonstrates smaller variance than IPCW methods. G-formula (Cox/ T-learners) and AIPCW-AIPTW (Cox & Cox & Log.Reg.) estimators seem to perform well, even in small samples. The forest versions of these estimators seem more biased, except Causal Survival Forest and the AIPTW-AIPCW (Forest). Notably, all estimators exhibit higher variability compared to Scenario 1.

5.2 Observational data

Data Generating Process

As for Scenarii 1 and 2, we carry out two simulations of an observational study with both independent and conditional independent censoring. The only difference lies in the simulation of the propensity score, which is no longer constant. For the simulation, an iid n -sample $(X_i, A_i, C, T_i(0), T_i(1))_{i \in [n]}$ is generated as in Section 5.1, except for the treatment allocation process that is given by:

$$\text{logit}(e(X)) = \beta_A^\top X \quad \text{where} \quad \beta_A = (-1, -1, -2.5, -1),$$

where we recall that $\text{logit}(p) = \log(p/(1-p))$. The descriptive statistics for the two observational data with independent (Obs1) and conditionally independent censoring (Obs2) are displayed in Appendix (Section 9.2). Note that we did not modify the survival distribution, the target difference in RMST is thus the same.

```
[1] "The ground truth for Obs scenario 1 at time 25 is 7.1"
```

```
[1] "The ground truth for Obs scenario 2 at #time 25 is 7.1"
```

Estimation of the RMST

Figure 10 below shows the distribution of the estimators of θ_{RMST} for the observational study with independent censoring.

In the simulation of an observational study with independent censoring, confounding bias is introduced, setting it apart from RCT simulations. As expected, estimators that fail to adjust for this bias, such as unadjusted Kaplan-Meier (KM), IPCW KM (Cox), and their equivalents, are biased. However, estimators like IPTW KM (Log.Reg.), IPTW-IPCW KM (Cox & Log. Reg.) are unbiased, even if the latter estimate unnecessary nuisance components. Results with IPTW BJ (Cox & Log.Reg) are extremely variable.

The top-performing estimators in this scenario are G-formula (Cox/ T-learners) and AIPCW-AIPTW (Cox & Cox & Log.Reg.), which are unbiased even with 500 observations. The former has the lowest variance. All estimators that use forests to estimate nuisance parameters are biased across sample sizes from 500 to 8000. Although Causal Survival Forest and AIPW-AIPCW (Forest) are expected to eventually converge, they remain extremely demanding in terms of sample size. This setting thus highlights that one should either have an a priori knowledge on the specification of the models or large sample size.

Figure 11 below shows the distribution of the θ_{RMST} estimates for the observational study with conditionally independent censoring. The red dashed line represents the true θ_{RMST} for $\tau = 25$.

In the simulation of an observational study with conditionally independent censoring, estimators that do not account for both censoring and confounding bias, such as KM, IPCW KM, IPTW KM, and their package equivalents, are biased. The top-performing estimators in this scenario are G-formula (Cox/ T-learners) and AIPCW-AIPTW (Cox & Cox & Log.Reg.), which are unbiased even with 500 observations. The former has the lowest variance as expected, see Section 2.2.1. Surprisingly, the

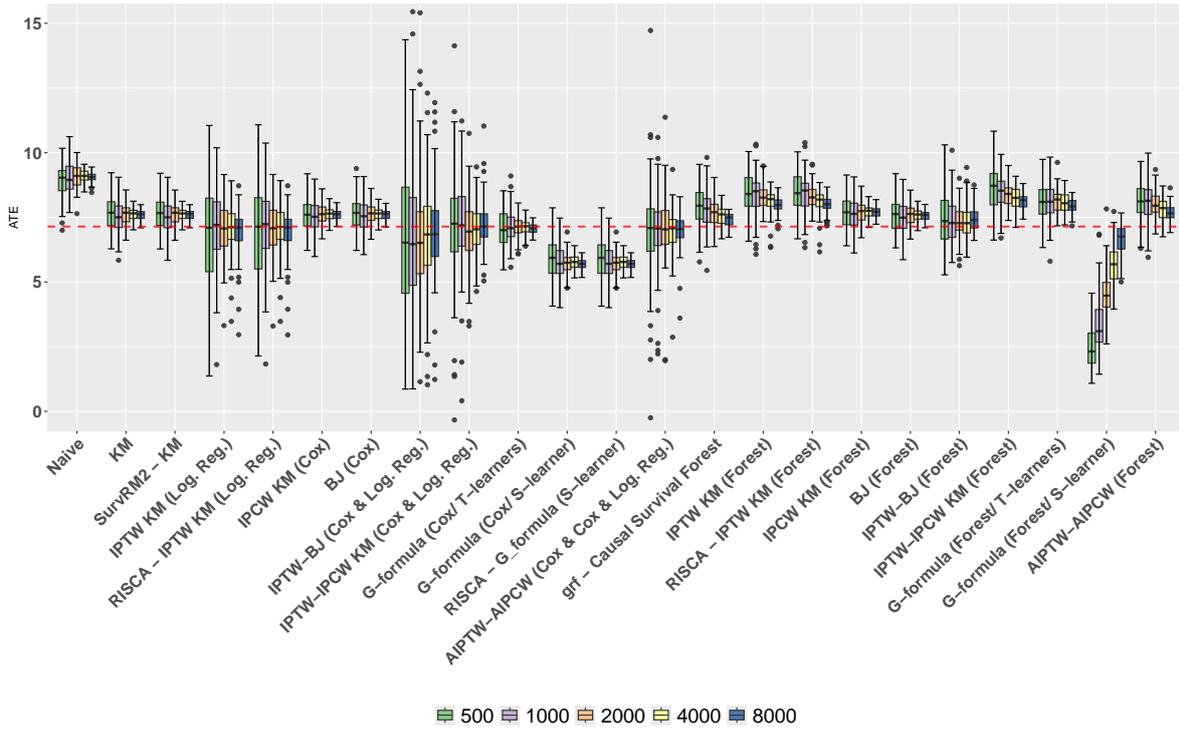


Figure 10: Estimation results of the ATE for the simulation of an observational study with independent censoring.

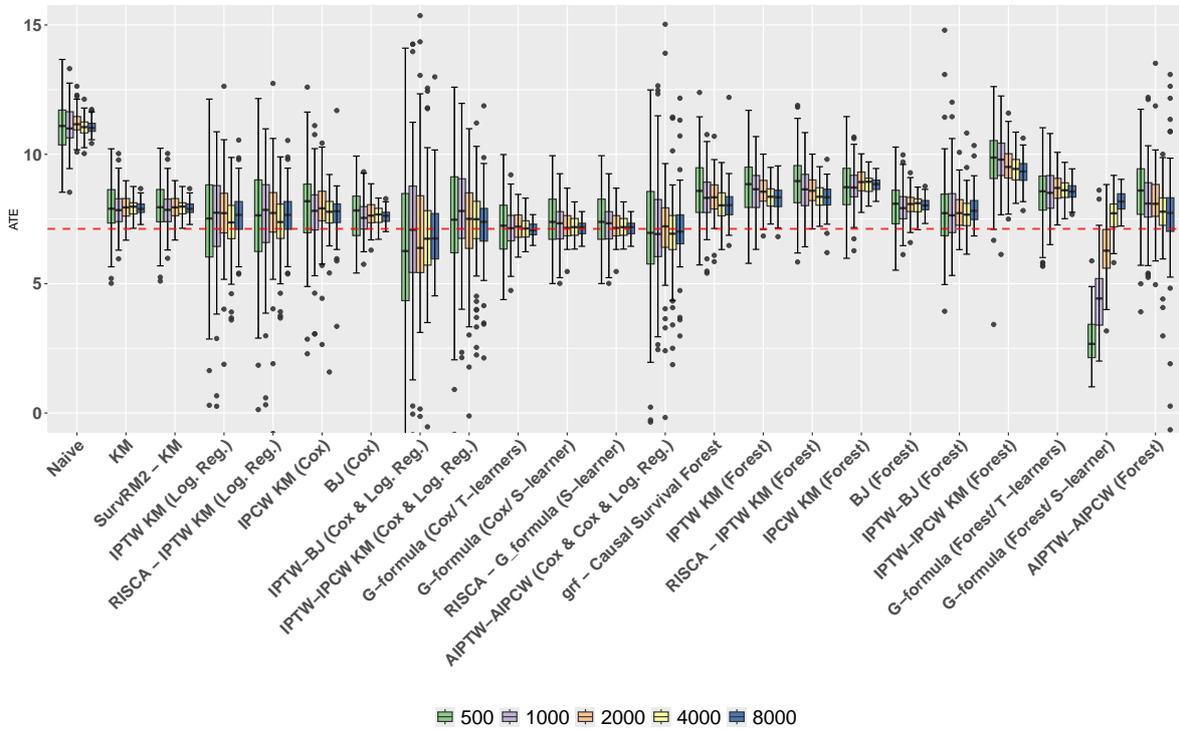


Figure 11: Estimation results of the ATE for the simulation of an observational study with dependent censoring.

G-formula (Cox/S-learner) and its equivalent from the RISCA package perform quite competitively, showing only a slight bias despite the violation of the proportional hazards assumption. All estimators that use forests to estimate nuisance parameters are biased across sample sizes from 500 to 8000. Although Causal Survival Forest and AIPTW-AIPCW (Forest) are expected to eventually converge, they remain extremely demanding in terms of sample size.

5.3 Misspecification of nuisance components

Data Generating Process

We generate an observational study with covariate interactions and conditionally independent censoring. The objective is to assess the impact of misspecifying nuisance components; specifically, we will use models that omit interactions to estimate these components. This approach enables us to evaluate the robustness properties of various estimators. In addition, in this setting forest based methods are expected to behave better.

We generate n samples $(X_i, A_i, C, T_i(0), T_i(1))$ as follows:

- $X \sim \mathcal{N}(\mu, \Sigma)$ and $\mu = (0.5, 0.5, 0.7, 0.5)$, $\Sigma = \text{Id}_4$.

- The hazard function of $T(0)$ is given by

$$\lambda^{(0)}(t|X) = \exp\{\beta_0^\top Y\} \quad \text{where} \quad \beta_0 = (0.2, 0.3, 0.1, 0.1, 1, 0, 0, 0, 0, 1),$$

$$\text{and} \quad Y = (X_1^2, X_2^2, X_3^2, X_4^2, X_1X_2, X_1X_3, X_1X_4, X_2X_3, X_2X_4, X_3X_4).$$

- The distribution of $T(1)$ is the one of $T(0)$ but shifted: $T(1) = T(0) + 1$.
- The hazard function of C is given by

$$\lambda_C(t|X) = \exp\{\beta_C^\top Y\} \quad \text{where} \quad \beta_C = (0.05, 0.05, -0.1, 0.1, 0, 1, 0, -1, 0, 0).$$

- The propensity score is

$$\text{logit}(e(x)) = \beta_A^\top Y \quad \text{where} \quad \beta_A = (0.05, -0.1, 0.5, -0.1, 1, 0, 1, 0, 0, 0).$$

When the model is well-specified, the full vector (X, Y) is given as an input of the nuisance parameter models. When it is not, only X and the first half of Y corresponding to $(X_1^2, X_2^2, X_3^2, X_4^2)$ is given as an input.

The descriptive statistics are given in Appendix (Section 9.3).

[1] "The ground truth for mis scenario at time 0.45 is 0.26"

Estimation of the RMST

First, we estimate θ_{RMST} without any model misspecification to confirm the consistency of the estimators under correctly specified nuisance models. More specifically, it means that for parametric propensity score models, semi-parametric censoring and survival models, we use models including interactions and squared assuming knowledge on the data generating process.

Next, we introduce misspecification individually for the treatment model, censoring model, and outcome model (Figure 13), i.e., we use models without interaction to estimate parametric and semi-parametric nuisance components while the data are generated with interactions.

We further examine combined misspecifications for pairs of models: treatment and censoring, treatment and outcome, and outcome and censoring. Finally, we assess the impact of misspecifying all nuisance models simultaneously (Figure 14).

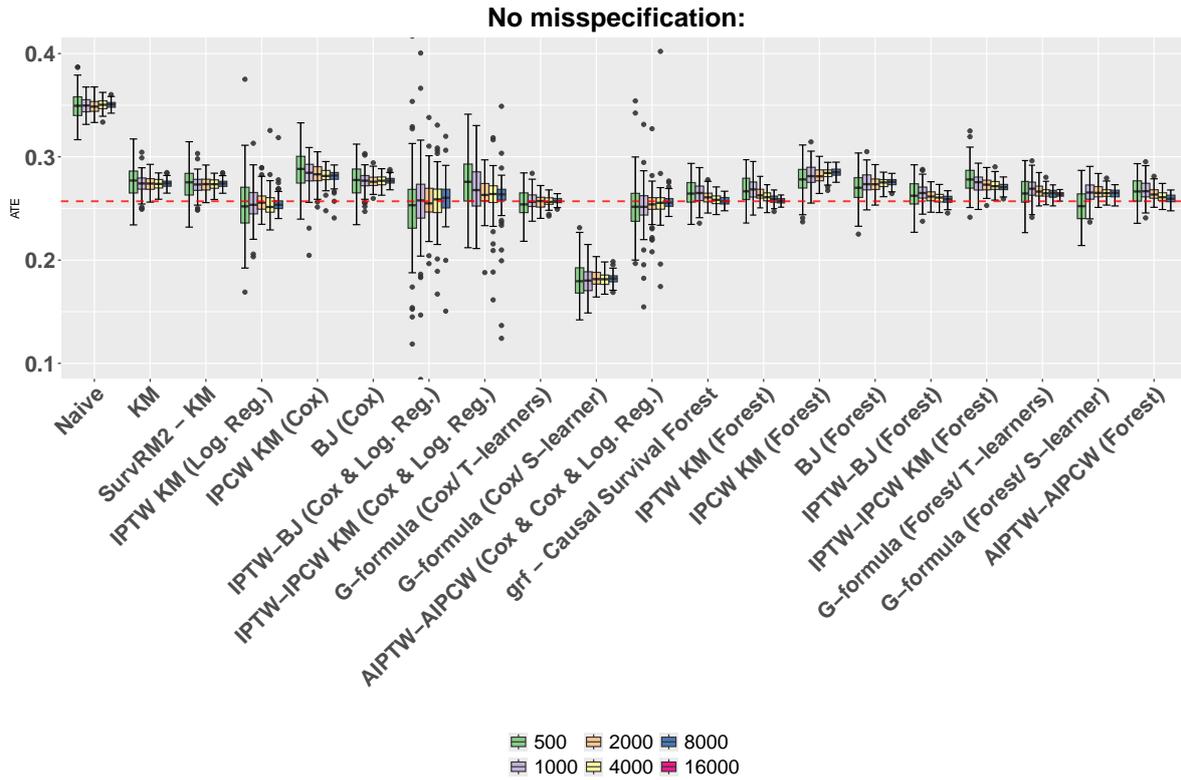


Figure 12: Estimation results of the ATE for the simulation of an observational study with dependent censoring and non linear relationships.

When there is no misspecification in Figure 12, as expected, IPTW-BJ (Cox & Log.Reg), G-formula (Cox/ T-learners) and AIPTW-AIPCW (Cox & Cox & Reg.Log) are unbiased. IPTW-IPCW KM (Cox & Log.Reg) exhibits a bias but seems to converge at larger sample size. Regarding forest-based methods, IPTW-BJ (Forest), AIPTW-AIPCW (Forest) and Causal Survival Forest estimate accurately the difference in RMST. Surprisingly, G-formula (Forest/ T-learners), G-formula (Forest/ S-learner) and IPTW-IPCW KM (Forest) exhibit small bias but are expected to eventually converge at large sample size.

Figure 13 shows that AIPTW-AIPCW (Cox & Cox & Reg.Log) is convergent when there is one nuisance misspecification. In contrary, the other estimators are biased when one of its nuisance parameter is misspecified.

Figure 14 shows that, as expected, when all nuisance models are misspecified, all estimators exhibit bias. AIPTW-AIPCW (Cox & Cox & Reg.Log) seems to converge in case where either the outcome and censoring models, or the treatment and censoring models are misspecified which deviates from initial expectations. It was anticipated that AIPTW-AIPCW would converge solely when both the censoring and treatment models were misspecified.

6 Conclusion

Based on the simulations and theoretical results, it might be advisable to stay away from the IPCW and IPTW-IPCW estimators, as they often exhibit excessive variability. Instead, we recommend implementing BJ which seems like a more stable transformation as IPCW, as well as Causal Survival Forest, G-formula (T-learners), IPTW-BJ, and AIPTW-AIPCW in both their Cox, Logistic Regression and forest versions. By qualitatively combining the results from these more robust estimators, we

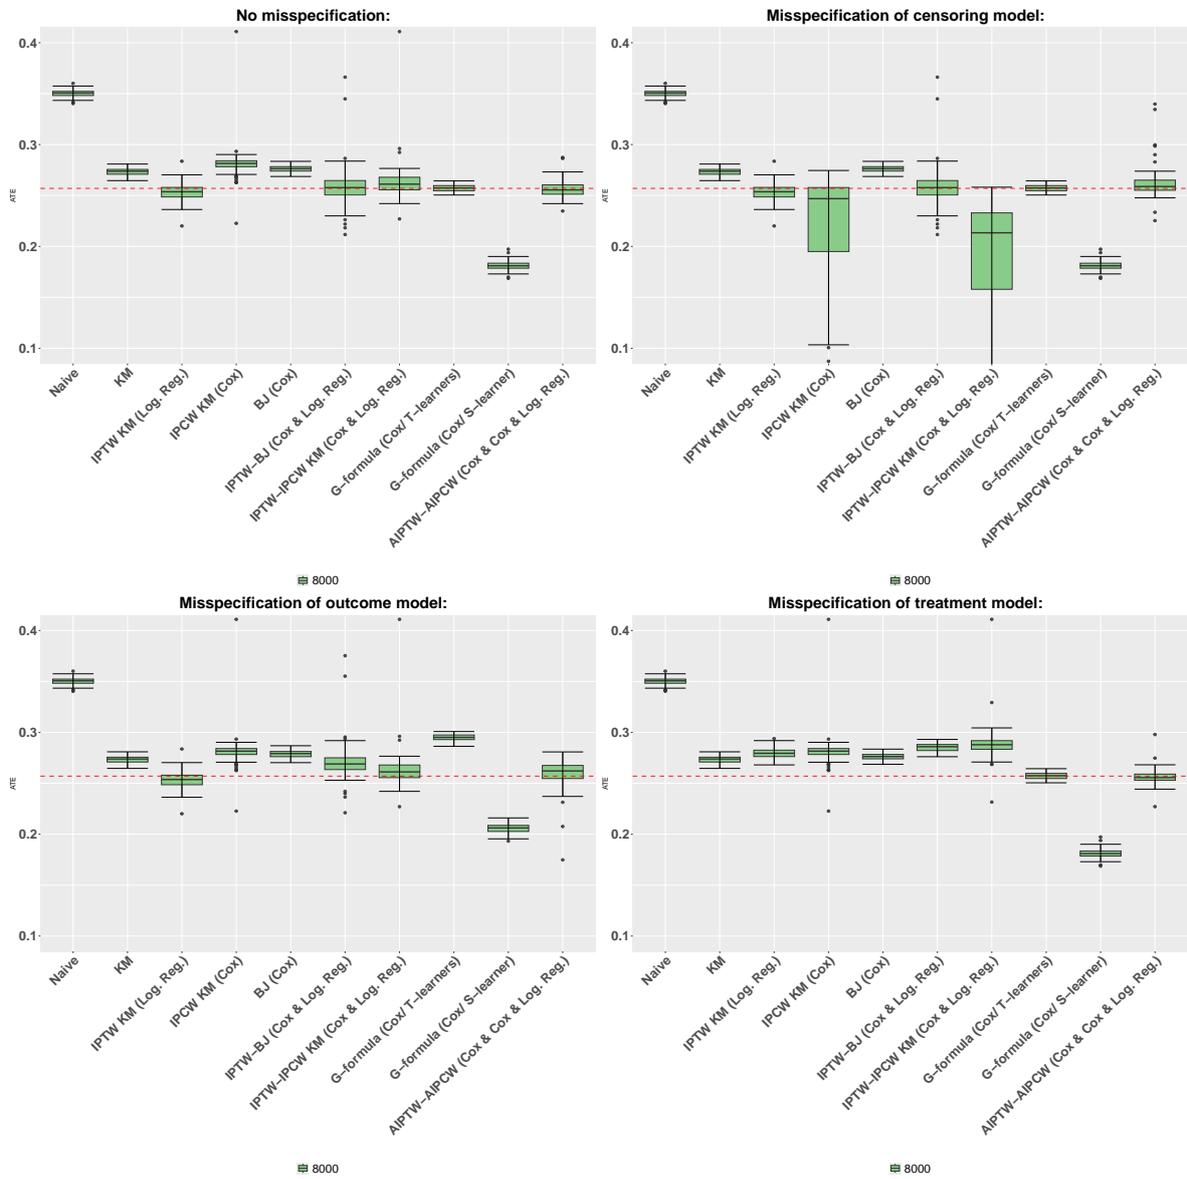


Figure 13: Estimation results of the ATE for an observational study with dependent censoring in case of a single misspecification.

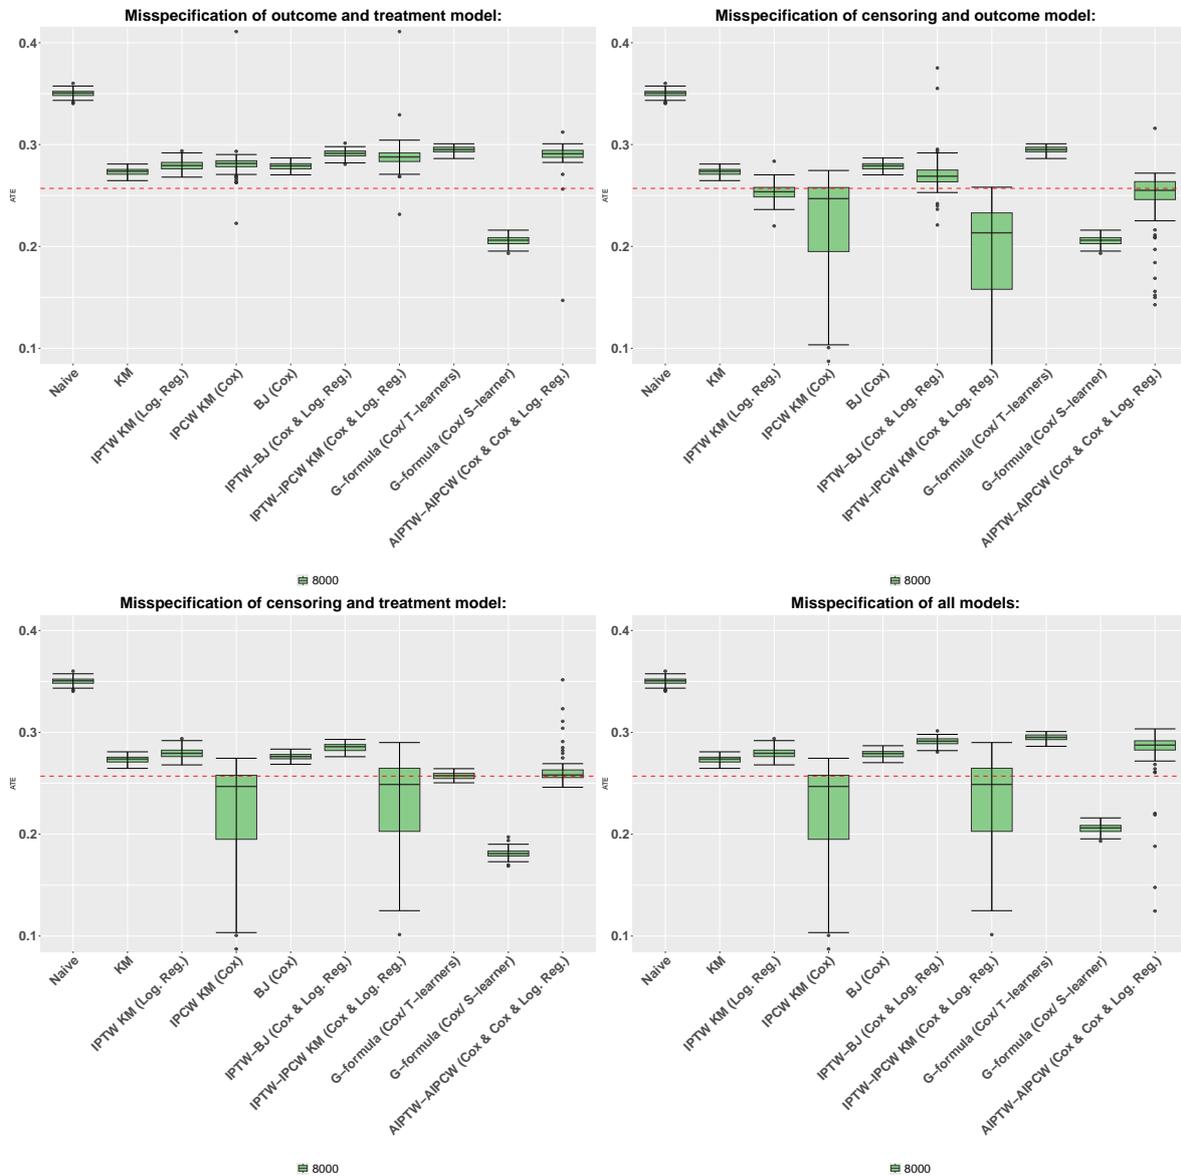


Figure 14: Estimation results of the ATE for an observational study with dependent censoring in case of a two or more misspecifications.

can expect to gain a fairly accurate understanding of the treatment effect.

It is important to note that our simulations utilize large sample sizes with relatively simple relationships, which may not fully capture the complexity of real-world scenarios. In practice, most survival analysis datasets tend to be smaller and more intricate, meaning the stability of certain estimators observed in our simulations may not generalize to real-world applications. Testing these methods on real-world datasets would provide a more comprehensive evaluation of their performance in practical settings.

An interesting direction for future work would be to focus on variable selection. Indeed, there is no reason to assume that the variables related to censoring should be the same as those linked to survival or treatment allocation. We could explore differentiating these sets of variables and study the impact on the estimators' variance. Similarly to causal inference settings without survival data, we might expect, for instance, that adding precision variables—those solely related to the outcome—could reduce the variance of the estimators.

Additionally, the estimators of the Restricted Mean Survival Time (RMST) provide a valuable alternative to the Hazard Ratio. The analysis and code provided with this article enables further exploration of the advantages of the estimators of RMST for causal analysis in survival studies. This could lead to a deeper understanding of how these estimators can offer more stable and interpretable estimates of treatment effects, particularly in complex real-world datasets.

7 Disclosure

This study was funded by Sanofi. Charlotte Voinot and Bernard Sebastien are Sanofi employees and may hold shares and/or stock options in the company. Clément Berenfeld, Imke Mayer and Julie Josse have nothing to disclose.

References

- Aalen, O., O. Borgan, and H. Gjessing. 2008. *Survival and Event History Analysis: A Process Point of View*. Statistics for Biology and Health. Springer New York. <https://books.google.de/books?id=wEi26X-VuCIC>.
- Breslow, N. 1974. "Covariance Analysis of Censored Survival Data." *Biometrics* 30 (1): 89–99. <http://www.jstor.org/stable/2529620>.
- Breslow, N., and J. Crowley. 1974. "A Large Sample Study of the Life Table and Product Limit Estimates Under Random Censorship." *The Annals of Statistics* 2 (3): 437–53. <https://doi.org/10.1214/aos/1176342705>.
- Buckley, and James. 1979. "Linear regression with censored data." *Biometrika* 66 (3): 429–36. <https://doi.org/10.1093/biomet/66.3.429>.
- Chen, Pei-Yun, and Anastasios A. Tsiatis. 2001. "Causal Inference on the Difference of the Restricted Mean Lifetime Between Two Groups." *Biometrics* 57 (4): 1030–38. <https://doi.org/10.1111/j.0006-341X.2001.01030.x>.
- Colnet, Bénédicte, Julie Josse, Gaël Varoquaux, and Erwan Scornet. 2022. "Reweighting the RCT for Generalization: Finite Sample Error and Variable Selection." *arXiv Preprint arXiv:2208.07614*.
- Cox, D. R. 1972. "Regression Models and Life-Tables." *Journal of the Royal Statistical Society: Series B (Methodological)* 34 (2): 187–202. <https://doi.org/10.1111/j.2517-6161.1972.tb00899.x>.
- Cui, Yifan, Michael R Kosorok, Erik Sverdrup, Stefan Wager, and Ruoqing Zhu. 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. <https://doi.org/10.1093/jrsssb/qqac001>.

- Ebrahimi, Nader, Daniel Molefe, and Zhiliang Ying. 2003. "Identifiability and Censored Data." *Biometrika* 90 (3): 724–27.
- European Medicines Agency, EMA. 2024. "Reflection Paper on Use of Real-World Data in Non-Interventional Studies to Generate Real-World Evidence." <https://www.ema.europa.eu/en/reflection-paper-use-real-world-data-non-interventional-studies-generate-real-world-evidence-scientific-guideline>.
- Fan, Jianqing, Yang Feng, and Yichao Wu. 2010. "High-Dimensional Variable Selection for Cox's Proportional Hazards Model." In *Borrowing Strength: Theory Powering Applications—a Festschrift for Lawrence d. Brown*, 6:70–87. Institute of Mathematical Statistics.
- Fan, Jianqing, and Irène Gijbels. 1994. "Censored Regression: Local Linear Approximations and Their Applications." *Journal of the American Statistical Association* 89 (426): 560–70. <https://doi.org/10.1080/01621459.1994.10476781>.
- Foster, Jared C., Jeremy M. G. Taylor, and Stephen J. Ruberg. 2011. "Subgroup Identification from Randomized Clinical Trial Data." *Statistics in Medicine* 30 (24): 2867–80. <https://doi.org/10.1002/sim.4322>.
- Foucher, Yohann, Florent Le Borgne, and Arthur Chatton. 2019. "RISCA: Causal Inference and Prediction in Cohort-Based Analyses." *CRAN: Contributed Packages*. The R Foundation. <https://doi.org/10.32614/cran.package.risca>.
- Gill, Richard. 1983. "Large Sample Behaviour of the Product-Limit Estimator on the Whole Line." *The Annals of Statistics*, 49–58.
- Hajime, Uno, Tian Lu, Horiguchi Miki, Cronin Angel, Battoui Chakib, and Bell James. 2015. "survRM2: Comparing Restricted Mean Survival Time." *CRAN: Contributed Packages*. The R Foundation. <https://doi.org/10.32614/cran.package.survrm2>.
- Hernán, Miguel A, and James M Robins. 2010. "Causal Inference." CRC Boca Raton, FL.
- Hirano, Keisuke, Guido W Imbens, and Geert Ridder. 2003. "Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score." *Econometrica* 71 (4): 1161–89.
- Howe, Chanelle J, Stephen R. Cole, Bryan Lau, Sonia Napravnik, and Joseph J. Eron. 2016. "Selection Bias Due to Loss to Follow up in Cohort Studies." *Epidemiology* 27 1: 91–97. <https://api.semanticscholar.org/CorpusID:21993915>.
- Huitfeldt, A., M. Stensrud, and E. Suzuki. 2019. "On the Collapsibility of Measures of Effect in the Counterfactual Causal Framework." *Emerging Themes in Epidemiology* 16: 1–12.
- Ishwaran, Hemant, Udaya B. Kogalur, Eugene H. Blackstone, and Michael S. Lauer. 2008. "Random survival forests." *The Annals of Applied Statistics* 2 (3): 841–60. <https://doi.org/10.1214/08-AOAS169>.
- Kalbfleisch, John D, and Ross L Prentice. 2002. *The Statistical Analysis of Failure Time Data*. John Wiley & Sons.
- Kaplan, E. L., and Paul Meier. 1958. "Nonparametric Estimation from Incomplete Observations." *Journal of the American Statistical Association* 53 (282): 457–81. <https://doi.org/10.1080/01621459.1958.10501452>.
- Koul, H., V. Susarla, and J. Van Ryzin. 1981. "Regression Analysis with Randomly Right-Censored Data." *The Annals of Statistics* 9 (6): 1276–88. <https://doi.org/10.1214/aos/1176345644>.
- Künzel, Sören R., Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. 2019. "Metalearners for Estimating Heterogeneous Treatment Effects Using Machine Learning." *Proceedings of the National Academy of Sciences* 116 (10): 4156–65. <https://doi.org/10.1073/pnas.1804597116>.
- Laan, Mark J. van der, and James M. Robins. 2003. "Introduction." In *Unified Methods for Censored Longitudinal Data and Causality*, 8–101. New York, NY: Springer New York. https://doi.org/10.1007/978-0-387-21700-0_1.
- Martinussen, Torben, Stijn Vansteelandt, and Per Andersen. 2020. "Subtleties in the Interpretation of Hazard Contrasts." *Lifetime Data Analysis* 26 (October). <https://doi.org/10.1007/s10985-020-09501-5>.
- Ozenne, Brice Maxime Hugues, Thomas Harder Scheike, Laila Stærk, and Thomas Alexander Gerds.

2020. “On the Estimation of Average Treatment Effects with Right-Censored Time to Event Outcome and Competing Risks.” *Biometrical Journal* 62 (3): 751–63. <https://doi.org/10.1002/bimj.201800298>.
- Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao. 1994. “Estimation of Regression Coefficients When Some Regressors Are Not Always Observed.” *Journal of the American Statistical Association* 89 (427): 846–66.
- Rosenbaum, Paul R., and Donald B. Rubin. 1983. “The central role of the propensity score in observational studies for causal effects.” *Biometrika* 70 (1): 41–55. <https://doi.org/10.1093/biomet/70.1.41>.
- Royston, Patrick, and Mahesh KB Parmar. 2013. “Restricted Mean Survival Time: An Alternative to the Hazard Ratio for the Design and Analysis of Randomized Trials with a Time-to-Event Outcome.” *BMC Medical Research Methodology* 13: 1–15.
- Rubin, Daniel, and Mark J. van der Laan. 2007. *The International Journal of Biostatistics* 3 (1). <https://doi.org/10.2202/1557-4679.1052>.
- Rubin, Donald B. 1974. “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies.” *Journal of Educational Psychology* 66 (5): 688–701. <https://doi.org/10.1037/h0037350>.
- Schaubel, Douglas E., and Guanghui Wei. 2011. “Double Inverse-Weighted Estimation of Cumulative Treatment Effects under Nonproportional Hazards and Dependent Censoring.” *Biometrics* 67 (1): 29–38. <https://doi.org/10.1111/j.1541-0420.2010.01449.x>.
- Therneau, Terry M. 2001. “Survival: Survival Analysis.” *CRAN: Contributed Packages*. The R Foundation. <https://doi.org/10.32614/cran.package.survival>.
- Tibshirani, Julie, Susan Athey, Erik Sverdrup, and Stefan Wager. 2017. “Grf: Generalized Random Forests.” *CRAN: Contributed Packages*. The R Foundation. <https://doi.org/10.32614/cran.package.grf>.
- Turkson, Anthony, Francis Ayiah-Mensah, and Vivian Nimoh. 2021. “Handling Censoring and Censored Data in Survival Analysis: A Standalone Systematic Literature Review.” *International Journal of Mathematics and Mathematical Sciences* 2021 (September): 1–16. <https://doi.org/10.1155/2021/9307475>.
- Xie, Jun, and Chaofeng Liu. 2005. “Adjusted Kaplan–Meier Estimator and Log-Rank Test with Inverse Probability of Treatment Weighting for Survival Data.” *Statistics in Medicine* 24 (20): 3089–3110. <https://doi.org/10.1002/sim.2174>.
- Zhang, Zhongheng. 2016. “Parametric Regression Model for Survival Data: Weibull Regression Model as an Example.” *Annals of Translational Medicine* 4 (24).
- Zhou, M. 1988. “Two-Sided Bias Bound of the Kaplan-Meier Estimator.” *Probability Theory and Related Fields* 79 (2): 165–73.

8 Appendix A: Proofs

8.1 Proofs of Section 2.1

Proof. (Proposition 2.1). Consistency is a trivial consequence of the law of large number and the identity 11. To show that \hat{S}_{KM} is unbiased, let us introduce \mathcal{F}_k be the filtration generated by the set of variables

$$\{A_i, \mathbb{I}\{\tilde{T}_i = t_j\}, \mathbb{I}\{\tilde{T}_i = t_j, \Delta_i = 1\} \mid j \in [k], i \in [n]\}.$$

which corresponds to the known information up to time t_k , so that $D_k(a)$ is \mathcal{F}_k -measurable but $N_k(a)$ is \mathcal{F}_{k-1} -measurable. One can write that, for $k \geq 2$

$$\begin{aligned} \mathbb{E}[\mathbb{I}\{\tilde{T}_i = t_k, \Delta_i = 1, A_i = a\} | \mathcal{F}_{k-1}] &= \mathbb{E}[\mathbb{I}\{\tilde{T}_i = t_k, \Delta_i = 1, A_i = a\} | \mathbb{I}\{\tilde{T}_i = t_k, A_i\}] \\ &= \mathbb{I}\{A_i = a\} \mathbb{E}[\mathbb{I}\{T_i = t_k, C_i = t_k\} | \mathbb{I}\{T_i = t_k, C_i = t_k, A_i\}] \\ &= \mathbb{I}\{A_i = a\} \mathbb{I}\{C_i = t_k\} \mathbb{E}[\mathbb{I}\{T_i = t_k\} | \mathbb{I}\{T_i = t_k, A_i\}] \\ &= \mathbb{I}\{\tilde{T}_i = t_k, A_i = a\} \left(1 - \frac{S^{(a)}(t_k)}{S^{(a)}(t_{k-1})} \right), \end{aligned}$$

where we used that $T_i(a)$ is independent from A_i by Assumption 5. We then easily derive from this that

$$\mathbb{E} \left[\left(1 - \frac{D_k(a)}{N_k(a)} \right) \mathbb{I}\{N_k(a) > 0\} \middle| \mathcal{F}_{k-1} \right] = \frac{S^{(a)}(t_k)}{S^{(a)}(t_{k-1})} \mathbb{I}\{N_k(a) > 0\},$$

and then that

$$\mathbb{E} \left[\hat{S}_{\text{KM}}(t_k | A = a) \middle| \mathcal{F}_{k-1} \right] = \frac{S^{(a)}(t_k)}{S^{(a)}(t_{k-1})} \hat{S}_{\text{KM}}(t_{k-1} | A = a) + O(\mathbb{I}\{N_k(a) = 0\}),$$

By induction, we easily find that

$$\mathbb{E}[\hat{S}_{\text{KM}}(t | A = a)] = \prod_{t_j \leq t} \frac{S^{(a)}(t_j)}{S^{(a)}(t_{j-1})} + O \left(\sum_{t_j \leq t} \mathbb{P}(N_j(a) = 0) \right) = S^{(a)}(t) + O(\mathbb{P}(N_k(a) = 0))$$

where t_k is the greatest time such that $t_k \leq t$. □

Proof. (Proposition 2.2). The asymptotic normality is a mere consequence of the joint asymptotic normality of $(N_k(a), D_k(a))_{t_k}$ with an application of the δ -method. To access the asymptotic variance, notice that, using a similar reasoning as in the previous proof:

$$\begin{aligned} \mathbb{E}[(1 - D_k(a)/N_k(a))^2 | \mathcal{F}_{k-1}] &= \mathbb{E}[1 - D_k(a)/N_k(a) | \mathcal{F}_{k-1}(a)]^2 + \frac{1}{N_k(a)^2} \text{Var}(D_k(a) | \mathcal{F}_{k-1}) \\ &= s_k^2(a) + \frac{s_k(a)(1 - s_k(a))}{N_k(a)} \mathbb{I}\{N_k(a) > 0\} + O(\mathbb{I}\{N_k(a) = 0\}). \end{aligned}$$

Now we know that $N_k(a) = nr_k(a) + \sqrt{n}O_{\mathbb{P}}(1)$, with the $O_{\mathbb{P}}(1)$ having uniformly bounded moments. So that we deduce that

$$\mathbb{E}[(1 - D_k(a)/N_k(a))^2 | \mathcal{F}_{k-1}] = s_k^2(a) + \frac{s_k(a)(1 - s_k(a))}{nr_k(a)} + \frac{1}{n^{3/2}} O_{\mathbb{P}}(1),$$

where $O_{\mathbb{P}}(1)$ has again bounded moments. Using this identity, we find that

$$\begin{aligned} n \text{Var} \hat{S}_{\text{KM}}(t | A = a) &= n \left(\mathbb{E} S_{\text{KM}}(t | A = a)^2 - S^{(a)}(t)^2 \right) \\ &= n S^{(a)}(t)^2 \left(\mathbb{E} \left[\prod_{t_k \leq t} \left(1 + \frac{1 - s_k(a)}{n s_k(a) r_k(a)} + \frac{1}{n^{3/2}} O_{\mathbb{P}}(1) \right) \right] - 1 \right). \end{aligned}$$

Expanding the product and using that the $O_{\mathbb{P}}(1)$'s have bounded moments, we finally deduce that

$$\mathbb{E} \left[\prod_{t_k \leq t} \left(1 + \frac{1 - s_k(a)}{n s_k(a) r_k(a)} + \frac{1}{n^{3/2}} O_{\mathbb{P}}(1) \right) \right] - 1 = \frac{1}{n} \sum_{t_k \leq t} \frac{1 - s_k(a)}{s_k(a) r_k(a)} + \frac{1}{n^{3/2}} O(1),$$

$$n \text{Var} \hat{S}_{\text{KM}}(t | A = a) = V_{\text{KM}}(t | A = a) + O(n^{-1/2}),$$

which is what we wanted to show. □

8.2 Proofs of Section 2.2

Proof. (Proposition 2.5). Assumption 17 allows the transformation to be well-defined. Furthermore, it holds

$$\begin{aligned}
E[T_{\text{PCW}}^*|A = a, X] &= E\left[\frac{\Delta^\tau \times \tilde{T} \wedge \tau}{G(\tilde{T} \wedge \tau|A, X)} \Big| A = a, X\right] \\
&= E\left[\frac{\Delta^\tau \times T(a) \wedge \tau}{G(T(a) \wedge \tau|A, X)} \Big| A = a, X\right] \\
&= E\left[E\left[\frac{\mathbb{I}\{T(a) \wedge \tau > C\} \times T(a) \wedge \tau}{G(T(a) \wedge \tau|A, X)} \Big| A, X, T(1)\right] \Big| A = a, X\right] \\
&= E[T(a) \wedge \tau|A = a, X] \\
&= E[T(a) \wedge \tau|X].
\end{aligned}$$

We used in the second equality that on the event $\{\Delta^\tau = 1, A = a\}$, it holds $\tilde{T} \wedge \tau = T \wedge \tau = T(a) \wedge \tau$. We used in the fourth equality that $G(T(a) \wedge \tau|A, X) = E[\mathbb{I}\{T(a) \wedge \tau > C\}|X, T(a), A]$ thanks to Assumption 16, and in the last one that A is independent from X and $T(a)$ thanks to Assumption 5. \square

Proof. (Proposition 2.6). Similarly to the computations done in the proof of Proposition 2.5, it is easy to show that

$$E\left[\frac{\Delta_i^\tau}{G(\tilde{T} \wedge \tau|X, A)} \mathbb{I}(\tilde{T}_i = t_k, A = a)\right] = \mathbb{P}(A = a)\mathbb{P}(T(a) = t_k),$$

and likewise that

$$E\left[\frac{\Delta_i^\tau}{G(\tilde{T} \wedge \tau|X, A)} \mathbb{I}(\tilde{T}_i < t_k, A = a)\right] = \mathbb{P}(A = a)\mathbb{P}(T(a) < t_k),$$

so that $\hat{S}_{\text{PCW}}(t)$ converges almost surely towards the product limit

$$\prod_{t_k \leq t} \left(1 - \frac{\mathbb{P}(T(a) = t_k)}{\mathbb{P}(T(a) < t_k)}\right) = S^{(a)}(t),$$

yielding strong consistency. Asymptotic normality is straightforward. \square

Proof. (Proposition 2.7). There holds

$$\begin{aligned}
E[T_{\text{BJ}}^*|X, A = a] &= E\left[\Delta^\tau T(a) \wedge \tau + (1 - \Delta^\tau) \frac{E[T \wedge \tau \times \mathbb{I}\{T \wedge \tau > C\}|C, A, X]}{\mathbb{P}(T > C|C, A, X)} \Big| X, A = a\right] \\
&= E[\Delta^\tau T(a) \wedge \tau|X] + E\left[\underbrace{\mathbb{I}\{T \wedge \tau > C\} \frac{E[T \wedge \tau \times \mathbb{I}\{T \wedge \tau > C\}|C, A, X]}{E[\mathbb{I}\{T \wedge \tau > C\}|C, A, X]}}_{(\star)} \Big| X, A = a\right].
\end{aligned}$$

Now we easily see that conditioning wrt X in the second term yields

$$\begin{aligned}
(\star) &= E\left[E[T \wedge \tau \times \mathbb{I}\{T \wedge \tau > C\}|C, A, X] \Big| X, A = a\right] \\
&= E[(1 - \Delta^\tau) T \wedge \tau|X, A = a] \\
&= E[(1 - \Delta^\tau) T(a) \wedge \tau|X],
\end{aligned}$$

ending the proof. \square

Proof. (Theorem 2.1). We let $T^* = \Delta^\tau \phi_1 + (1 - \Delta^\tau) \phi_0$ be a transformation of the form Equation 20. There holds

$$\mathbb{E}[(T^* - T \wedge \tau)^2] = \mathbb{E}[\Delta^\tau (\phi_1 - T \wedge \tau)^2] + \mathbb{E}[(1 - \Delta^\tau) (\phi_0 - T \wedge \tau)^2].$$

The first term is non negative and is zero for the BJ transformation. Since ϕ_0 is a function of (\tilde{T}, X, A) and that $\tilde{T} = C$ on $\{\Delta^\tau = 0\}$, the second term can be rewritten in the following way. We let R be a generic quantity that does not depend on ϕ_0 .

$$\begin{aligned} \mathbb{E}[(1 - \Delta^\tau) (\phi_0 - T)^2] &= \mathbb{E} \left[\mathbb{I}\{T \wedge \tau > C\} \phi_0^2 - 2\mathbb{I}\{T \wedge \tau > C\} \phi_0 T \wedge \tau \right] + R \\ &= \mathbb{E} \left[\mathbb{P}(T \wedge \tau > C | C, A, X) \phi_0^2 - 2\mathbb{E}[T \wedge \tau \mathbb{I}\{T \wedge \tau > C\} | C, A, X] \phi_0 \right] + R \\ &= \mathbb{E} \left[\mathbb{P}(T \wedge \tau > C | C, A, X) \left(\phi_0 - \frac{\mathbb{E}[T \wedge \tau \mathbb{I}\{T \wedge \tau > C\} | C, A, X]}{\mathbb{P}(T \wedge \tau > C | C, A, X)} \right)^2 \right] + R. \end{aligned}$$

Now the first term in the right hand side is always non-negative, and is zero for the BJ transformation. \square

8.3 Proofs of Section 3.1

Proof. (Proposition 3.1). The fact that it is strongly consistent and asymptotically normal is again a simple application of the law of large number and of the δ -method. We indeed find that, for t_k

$$\begin{aligned} \mathbb{E} \left[\frac{1}{e(X_i)} \mathbb{1}\{\tilde{T}_i = t_k, \Delta_i = 1, A_i = 1\} \right] &= \mathbb{E} \left[\frac{A_i}{e(X_i)} \mathbb{1}\{T_i = t_k, C_i = t_k\} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\frac{A_i}{e(X_i)} \mathbb{1}\{T_i = t_k, C_i = t_k\} \middle| X_i \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\frac{A_i}{e(X_i)} \middle| X_i \right] \mathbb{P}(T_i = t_k | X_i) \mathbb{P}(C_i = t_k) \right] \\ &= \mathbb{P}(T_i = t_k) \mathbb{P}(C_i = t_k), \end{aligned}$$

where we used that A is independent from T conditionnaly on X , and that C is independent from everything. Likewise, one would get that

$$\mathbb{E} \left[\frac{1}{e(X_i)} \mathbb{1}\{\tilde{T}_i = t_k, A_i = 1\} \right] = \mathbb{P}(T_i = t_k) \mathbb{P}(C_i = t_k).$$

Similar computations hold for $A = 0$, ending the proof. \square

8.4 Proofs of Section 3.2

Proof. (Proposition 3.2). On the event $\{\Delta^\tau = 1, A = 1\}$, there holds $\tilde{T} \wedge \tau = T \wedge \tau = T(1) \wedge \tau$, whence we find that,

$$\begin{aligned} \mathbb{E}[T_{\text{IPCW}}^* | X, A = 1] &= \mathbb{E} \left[\frac{A}{e(X)} \frac{\mathbb{I}\{T(1) \wedge \tau > C\}}{G(T(1) \wedge \tau | X, A)} T(1) \wedge \tau \middle| X \right] \\ &= \mathbb{E} \left[\frac{A}{e(X)} \mathbb{E} \left[\frac{\mathbb{I}\{T(1) \wedge \tau > C\}}{G(T(1) \wedge \tau | X, A)} \middle| X, A, T(1) \right] T(1) \wedge \tau \middle| X \right] \\ &= \mathbb{E} \left[\frac{A}{e(X)} T(1) \wedge \tau \middle| X \right] \\ &= \mathbb{E} [T(1) \wedge \tau | X], \end{aligned}$$

and the same holds on the event $A = 0$. \square

Proof. (Proposition 3.3). By consistency of $\widehat{G}(\cdot|X, A)$ and \widehat{e} and by continuity, it suffices to look at the asymptotic behavior of the oracle estimator

$$\theta_{\text{IPTW-IPCW}}^* = \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i}{e(X_i)} - \frac{1-A_i}{1-e(X_i)} \right) \frac{\Delta_i^\tau}{G(\tilde{T}_i \wedge \tau | A_i, X_i)} \tilde{T}_i \wedge \tau.$$

The later is converging almost towards its mean, which, following similar computations as in the previous proof, write

$$\begin{aligned} \mathbb{E} \left[\left(\frac{A}{e(X)} - \frac{1-A}{1-e(X)} \right) \frac{\Delta^\tau}{G(\tilde{T} \wedge \tau | A, X)} \tilde{T} \wedge \tau \right] &= \mathbb{E} \left[\left(\frac{A}{e(X)} - \frac{1-A}{1-e(X)} \right) T \wedge \tau \right] \\ &= \mathbb{E} [T(1) \wedge \tau] - \mathbb{E} [T(0) \wedge \tau]. \end{aligned}$$

□

Proof. (Proposition 3.4). Asymptotic normality comes from a mere application of the δ -method, while strong consistency follows from the law of large number and the following computations. Like for the proof of Proposition 2.6, one find, by first conditioning wrt $X, A, T(a)$, that, for t_k, τ ,

$$\mathbb{E} \left[\left(\frac{A}{e(X)} + \frac{1-A}{1-e(X)} \right) \frac{\Delta^\tau}{G(\tilde{T} \wedge \tau | A, X)} \mathbb{I}\{\tilde{T} = t_k, A = a\} \right] = \mathbb{P}(T(a) = t_k)$$

and likewise that

$$\mathbb{E} \left[\left(\frac{A}{e(X)} + \frac{1-A}{1-e(X)} \right) \frac{\Delta^\tau}{G(\tilde{T} \wedge \tau | A, X)} \mathbb{I}\{\tilde{T} < t_k, A = a\} \right] = \mathbb{P}(T(a) < t_k)$$

so that indeed $S_{\text{IPTW-IPCW}}^*(t|A = a)$ converges almost surely towards $S^{(a)}(t)$. □

Proof. (Proposition 3.5). We write

$$\mathbb{E}[T_{\text{IPTW-BJ}}^* | X, A = 1] = \mathbb{E} \left[\frac{A}{e(X)} \Delta^\tau \times \tilde{T} \wedge \tau \middle| X \right] + \mathbb{E} \left[\frac{A}{e(X)} (1 - \Delta^\tau) Q_S(\tilde{T} \wedge \tau | A, X) \middle| X \right].$$

On the event $\{\Delta^\tau = 1, A = 1\}$, there holds $\tilde{T} \wedge \tau = T \wedge \tau = T(1) \wedge \tau$, whence we find that the first term on the the right hand side is equal to

$$\begin{aligned} \mathbb{E} \left[\frac{A}{e(X)} \Delta^\tau \times \tilde{T} \wedge \tau \middle| X \right] &= \mathbb{E} \left[\frac{A}{e(X)} \Delta^\tau \times T(1) \wedge \tau \middle| X \right] \\ &= \mathbb{E} [\Delta^\tau \times T(1) \wedge \tau | X]. \end{aligned}$$

For the second term in the right hand side, notice that on the event $\{\Delta^\tau = 0, A = 1\}$, there holds $\tilde{T} = C < T(1) \wedge \tau$, so that

$$\begin{aligned} &\mathbb{E} \left[\frac{A}{e(X)} \mathbb{I}\{C < T(1) \wedge \tau\} \frac{\mathbb{E}[T(1) \wedge \tau \times \mathbb{I}\{C < T(1) \wedge \tau\} | X, A, C]}{\mathbb{P}(C < T(1) \wedge \tau | C, X, A)} \middle| X \right] \\ &= \mathbb{E} \left[\frac{A}{e(X)} \mathbb{E}[T(1) \wedge \tau \times \mathbb{I}\{C < T(1) \wedge \tau\} | X, A, C] \middle| X \right] \\ &= \mathbb{E} [T(1) \wedge \tau \times \mathbb{I}\{C < T(1) \wedge \tau\} | X] \\ &= \mathbb{E} [(1 - \Delta^\tau) T(1) \wedge \tau | X], \end{aligned}$$

and the same holds on the event $\{A = 0\}$, which ends the proof. □

Proof. (Proposition 3.6). By consistency of $\widehat{G}(\cdot|X, A)$ and \widehat{e} and by continuity, it suffices to look at the asymptotic behavior of the oracle estimator

$$\theta_{\text{IPTW-BJ}}^* = \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i}{e(X_i)} - \frac{1 - A_i}{1 - e(X_i)} \right) (\Delta_i^\tau \times \tilde{T}_i \wedge \tau + (1 - \Delta_i^\tau) Q_S(\tilde{T}_i \wedge \tau | A_i, X_i)).$$

The later is converging almost towards its mean, which, following similar computations as in the previous proof, is simply equal to the difference in RMST. \square

Proof. (Proposition 3.7). We can write that

$$\Delta_{\text{QR}}^* = \underbrace{\frac{A}{p(X)} (T_{\text{DR}}^*(F, R) - v(X, 1)) + v(X, 1)}_{(A)} - \underbrace{\left(\frac{1 - A}{1 - p(X)} (T_{\text{DR}}^*(F, R) - v(X, 0)) + v(X, 0) \right)}_{(B)}.$$

Focusing on term (A), we easily find that

$$\begin{aligned} \mathbb{E}[(A)|X] &= \mathbb{E} \left[\frac{A}{p(X)} (T_{\text{DR}}^*(F, R) - v(X, 1)) + v(X, 1) \middle| X \right] \\ &= \frac{e(X)}{p(X)} (\mu(X, 1) - v(X, 1)) + v(X, 1). \end{aligned}$$

Where we used that $T_{\text{DR}}^*(F, R)$ is a censoring unbiased transform when $F = G$ or $R = S$. Now we see that if $p(X) = e(X)$, then

$$\mathbb{E}[(A)|X] = \mu(X, 1) - v(X, 1) + v(X, 1) = \mu(X, 1),$$

and if $v(X, 1) = \mu(X, 1)$, then

$$\mathbb{E}[(A)|X] = \frac{e(X)}{p(X)} \times 0 + \mu(X, 1) = \mu(X, 1).$$

Likewise, we would show that $\mathbb{E}[(B)|X] = \mu(X, 0)$ under either alternative, ending the proof. \square

9 Appendix B: Descriptive statistics

9.1 RCT

The summary by group of treatment of the generated (observed and unobserved) RCT with independent censoring is displayed below:

[1] "Descriptive statistics for group A=0: 1053"

X1	X2	X3	X4
Min. : -1.7349	Min. : -2.5663	Min. : -3.8925	Min. : -2.0360
1st Qu.: 0.3874	1st Qu.: 0.3788	1st Qu.: -1.7245	1st Qu.: 0.4379
Median : 1.0238	Median : 1.0206	Median : -1.0558	Median : 1.0744
Mean : 1.0227	Mean : 1.0285	Mean : -1.0472	Mean : 1.0664
3rd Qu.: 1.7022	3rd Qu.: 1.6939	3rd Qu.: -0.3646	3rd Qu.: 1.7108
Max. : 4.3746	Max. : 4.4499	Max. : 2.5493	Max. : 4.1453
C	T1	T0	status
Min. : 0.01707	Min. : 10.00	Min. : 0.0038	Min. : 0.000
1st Qu.: 10.15906	1st Qu.: 12.42	1st Qu.: 2.4243	1st Qu.: 0.000

```

Median : 24.58672  Median : 17.69  Median : 7.6852  Median :1.000
Mean   : 35.10085  Mean   : 30.75  Mean   : 20.7507  Mean   :0.698
3rd Qu.: 49.45058  3rd Qu.: 30.54  3rd Qu.: 20.5401  3rd Qu.:1.000
Max.   :283.29895  Max.   :965.02  Max.   :955.0185  Max.   :1.000

```

```

T_tild
Min.   : 0.0038
1st Qu.: 1.9573
Median : 5.3226
Mean   :10.1376
3rd Qu.:13.1953
Max.   :94.1313

```

```
[1] "Descriptive statistics for group A=1: 947"
```

```

      X1          X2          X3          X4
Min.  :-2.8647  Min.  :-2.1411  Min.  :-4.0319  Min.  :-2.5553
1st Qu.: 0.2581  1st Qu.: 0.3464  1st Qu.:-1.5864  1st Qu.: 0.3386
Median : 0.9561  Median : 1.0099  Median :-0.9133  Median : 1.0125
Mean   : 0.9466  Mean   : 0.9991  Mean   :-0.9312  Mean   : 1.0299
3rd Qu.: 1.6559  3rd Qu.: 1.6977  3rd Qu.:-0.2499  3rd Qu.: 1.7116
Max.   : 4.3198  Max.   : 3.7245  Max.   : 2.5541  Max.   : 3.9869

```

```

      C          T1          T0          status
Min.   : 0.03081  Min.   : 10.02  Min.   : 0.0189  Min.   :0.0000
1st Qu.: 9.45254  1st Qu.: 13.17  1st Qu.: 3.1690  1st Qu.:0.0000
Median : 25.31415  Median : 19.68  Median : 9.6783  Median :1.0000
Mean   : 34.66363  Mean   : 31.99  Mean   : 21.9885  Mean   :0.5048
3rd Qu.: 46.72254  3rd Qu.: 33.99  3rd Qu.: 23.9901  3rd Qu.:1.0000
Max.   :217.99849  Max.   :365.43  Max.   :355.4287  Max.   :1.0000

```

```

T_tild
Min.   : 0.03081
1st Qu.: 9.45254
Median : 13.22399
Mean   : 17.39243
3rd Qu.: 21.17930
Max.   :149.40454

```

Covariates are balanced between groups, and censoring times are the same (independent censoring). However, there are more censored observations in the treated group ($A = 1$) than in the control group ($A = 0$). This is due to the higher instantaneous hazard of the event in the treated group (with $T_1 = T_0 + 10$) compared to the constant hazard of censoring.

The summary of the generated (observed and unobserved) RCT with conditionally independent censoring stratified by treatment is displayed below.

```
[1] "Descriptive statistics for group A=0: 1006"
```

```

      X1          X2          X3          X4
Min.  :-1.8983  Min.  :-2.6673  Min.  :-4.0236  Min.  :-2.2433
1st Qu.: 0.4118  1st Qu.: 0.3122  1st Qu.:-1.7435  1st Qu.: 0.3677
Median : 1.0311  Median : 0.9990  Median :-0.9930  Median : 1.0121
Mean   : 1.0264  Mean   : 1.0227  Mean   :-1.0016  Mean   : 1.0204
3rd Qu.: 1.7015  3rd Qu.: 1.7126  3rd Qu.:-0.2952  3rd Qu.: 1.6647
Max.   : 4.4370  Max.   : 4.4759  Max.   : 2.0828  Max.   : 4.4475

```

C	T1	T0	status
Min. : 0.00146	Min. : 10.03	Min. : 0.0281	Min. : 0.0000
1st Qu.: 2.58160	1st Qu.: 12.60	1st Qu.: 2.6012	1st Qu.: 0.0000
Median : 6.71686	Median : 17.26	Median : 7.2615	Median : 0.0000
Mean : 13.42545	Mean : 29.69	Mean : 19.6897	Mean : 0.4473
3rd Qu.: 16.28881	3rd Qu.: 30.70	3rd Qu.: 20.7009	3rd Qu.: 1.0000
Max. : 209.00395	Max. : 397.44	Max. : 387.4447	Max. : 1.0000

status_tau	T_tild
Min. : 0.0000	Min. : 0.00146
1st Qu.: 0.0000	1st Qu.: 1.27297
Median : 0.0000	Median : 3.37826
Mean : 0.4881	Mean : 7.25646
3rd Qu.: 1.0000	3rd Qu.: 7.87051
Max. : 1.0000	Max. : 127.65479

[1] "Descriptive statistics for group A=1: 994"

X1	X2	X3	X4
Min. : -1.9653	Min. : -3.0018	Min. : -3.8764	Min. : -2.2017
1st Qu.: 0.2779	1st Qu.: 0.3366	1st Qu.: -1.6976	1st Qu.: 0.3591
Median : 0.9193	Median : 0.9757	Median : -1.0526	Median : 1.0136
Mean : 0.9680	Mean : 0.9870	Mean : -1.0353	Mean : 1.0163
3rd Qu.: 1.6599	3rd Qu.: 1.6134	3rd Qu.: -0.3524	3rd Qu.: 1.6847
Max. : 3.8066	Max. : 4.1911	Max. : 2.0895	Max. : 4.3685

C	T1	T0	status
Min. : 0.0097	Min. : 10.01	Min. : 0.0147	Min. : 0.0000
1st Qu.: 3.1894	1st Qu.: 13.07	1st Qu.: 3.0665	1st Qu.: 0.0000
Median : 8.6746	Median : 18.91	Median : 8.9100	Median : 0.0000
Mean : 17.7391	Mean : 33.09	Mean : 23.0947	Mean : 0.2314
3rd Qu.: 20.4410	3rd Qu.: 34.28	3rd Qu.: 24.2818	3rd Qu.: 0.0000
Max. : 532.6940	Max. : 633.48	Max. : 623.4836	Max. : 1.0000

status_tau	T_tild
Min. : 0.0000	Min. : 0.00974
1st Qu.: 0.0000	1st Qu.: 3.18939
Median : 0.0000	Median : 8.67462
Mean : 0.2907	Mean : 11.97123
3rd Qu.: 1.0000	3rd Qu.: 14.69202
Max. : 1.0000	Max. : 134.65165

Covariates are balanced between the two groups. However, censoring times differ between groups due to conditionally independent censoring based on covariates and treatment group. Indeed, the distribution of C is different between the treatment group.

9.2 Observational study with linear relationship

The summary of the generated (observed and unobserved) data set observational study with independent censoring stratified by treatment is displayed below to enhance the difference with the other scenario.

[1] "Descriptive statistics for group A=0: 1109"

X1	X2	X3	X4
Min. : -2.3330	Min. : -1.8751	Min. : -3.26193	Min. : -1.5008

1st Qu.: 0.5218	1st Qu.: 0.5882	1st Qu.: -1.05047	1st Qu.: 0.4846
Median : 1.2327	Median : 1.2913	Median : -0.53051	Median : 1.1734
Mean : 1.1866	Mean : 1.2613	Mean : -0.52258	Mean : 1.1842
3rd Qu.: 1.8763	3rd Qu.: 1.8910	3rd Qu.: 0.05191	3rd Qu.: 1.8448
Max. : 4.7463	Max. : 4.2293	Max. : 2.18757	Max. : 4.3594
C	T1	T0	status
Min. : 0.00037	Min. : 10.01	Min. : 0.0064	Min. : 0.0000
1st Qu.: 9.08596	1st Qu.: 12.63	1st Qu.: 2.6296	1st Qu.: 0.0000
Median : 23.44589	Median : 17.53	Median : 7.5317	Median : 1.0000
Mean : 34.00533	Mean : 30.04	Mean : 20.0356	Mean : 0.6943
3rd Qu.: 48.09966	3rd Qu.: 31.93	3rd Qu.: 21.9266	3rd Qu.: 1.0000
Max. : 251.41596	Max. : 491.67	Max. : 481.6699	Max. : 1.0000
T_tild			
Min. : 0.00037			
1st Qu.: 1.96940			
Median : 5.39172			
Mean : 10.32146			
3rd Qu.: 12.79874			
Max. : 168.69106			

[1] "Descriptive statistics for group A=1: 891"

X1	X2	X3	X4
Min. : -2.37512	Min. : -2.1902	Min. : -4.5059	Min. : -2.01609
1st Qu.: 0.08916	1st Qu.: 0.1560	1st Qu.: -2.1669	1st Qu.: 0.05362
Median : 0.75965	Median : 0.8066	Median : -1.5974	Median : 0.71811
Mean : 0.75126	Mean : 0.7846	Mean : -1.6236	Mean : 0.71452
3rd Qu.: 1.39138	3rd Qu.: 1.4246	3rd Qu.: -1.0577	3rd Qu.: 1.35626
Max. : 3.87355	Max. : 3.6591	Max. : 0.6806	Max. : 3.97256
C	T1	T0	status
Min. : 0.0766	Min. : 10.01	Min. : 0.0071	Min. : 0.0000
1st Qu.: 8.9967	1st Qu.: 12.92	1st Qu.: 2.9225	1st Qu.: 0.0000
Median : 22.5362	Median : 18.37	Median : 8.3709	Median : 0.0000
Mean : 31.6058	Mean : 31.71	Mean : 21.7056	Mean : 0.4882
3rd Qu.: 44.9464	3rd Qu.: 33.31	3rd Qu.: 23.3142	3rd Qu.: 1.0000
Max. : 201.0258	Max. : 382.17	Max. : 372.1723	Max. : 1.0000
T_tild			
Min. : 0.0766			
1st Qu.: 8.9967			
Median : 12.8082			
Mean : 16.0926			
3rd Qu.: 19.9619			
Max. : 136.1725			

The covariates between the two groups of treatment are unbalanced because of dependent treatment assignment. The mean of X1, X2, X3 and X4 is bigger in the control group than in the treated group. The censoring times have the same distribution (independent censoring). There are more censored observation in the treated group (A=1) than in the control group (A=0) for the same reason than in the RCT scenario.

The summary of the generated (observed and unobserved) data set observational study with conditionally independent censoring stratified by treatment is displayed below.

[1] "Descriptive statistics for group A=0: 1135"

X1		X2		X3		X4	
Min.	:-2.1274	Min.	:-1.6571	Min.	:-3.09281	Min.	:-1.8423
1st Qu.:	0.5145	1st Qu.:	0.5386	1st Qu.:	-1.11137	1st Qu.:	0.5463
Median :	1.1402	Median :	1.1625	Median :	-0.57357	Median :	1.2030
Mean :	1.1748	Mean :	1.2001	Mean :	-0.52394	Mean :	1.2382
3rd Qu.:	1.8621	3rd Qu.:	1.8824	3rd Qu.:	0.02643	3rd Qu.:	1.9129
Max. :	4.1307	Max. :	4.3984	Max. :	3.41547	Max. :	3.9189

C		T1		T0		status	
Min. :	0.00205	Min. :	10.00	Min. :	0.0004	Min. :	0.0000
1st Qu.:	2.55604	1st Qu.:	12.50	1st Qu.:	2.5008	1st Qu.:	0.0000
Median :	6.96852	Median :	17.72	Median :	7.7205	Median :	0.0000
Mean :	14.12909	Mean :	30.46	Mean :	20.4648	Mean :	0.4793
3rd Qu.:	16.38653	3rd Qu.:	29.98	3rd Qu.:	19.9838	3rd Qu.:	1.0000
Max. :	199.00909	Max. :	501.09	Max. :	491.0921	Max. :	1.0000

status_tau		T_obs		e	
Min. :	0.0000	Min. :	0.00039	Min. :	0.0000033
1st Qu.:	0.0000	1st Qu.:	1.34819	1st Qu.:	0.0263812
Median :	1.0000	Median :	3.48376	Median :	0.1059645
Mean :	0.5119	Mean :	7.46724	Mean :	0.2042602
3rd Qu.:	1.0000	3rd Qu.:	8.51456	3rd Qu.:	0.3137048
Max. :	1.0000	Max. :	135.08089	Max. :	0.9788361

[1] "Descriptive statistics for group A=1: 865"

X1		X2		X3		X4	
Min.	:-2.5798	Min.	:-2.64701	Min.	:-4.394	Min.	:-2.34523
1st Qu.:	0.1759	1st Qu.:	0.07471	1st Qu.:	-2.189	1st Qu.:	0.07272
Median :	0.7802	Median :	0.76465	Median :	-1.644	Median :	0.72962
Mean :	0.7726	Mean :	0.73404	Mean :	-1.648	Mean :	0.69508
3rd Qu.:	1.3668	3rd Qu.:	1.37752	3rd Qu.:	-1.134	3rd Qu.:	1.34778
Max. :	3.4179	Max. :	4.47830	Max. :	1.138	Max. :	3.53681

C		T1		T0		status	
Min. :	0.0182	Min. :	10.01	Min. :	0.0105	Min. :	0.000
1st Qu.:	2.7635	1st Qu.:	13.18	1st Qu.:	3.1771	1st Qu.:	0.000
Median :	7.4038	Median :	18.88	Median :	8.8809	Median :	0.000
Mean :	14.5227	Mean :	35.30	Mean :	25.2969	Mean :	0.178
3rd Qu.:	17.4710	3rd Qu.:	33.06	3rd Qu.:	23.0577	3rd Qu.:	0.000
Max. :	380.7475	Max. :	863.19	Max. :	853.1887	Max. :	1.000

status_tau		T_obs		e	
Min. :	0.0000	Min. :	0.01817	Min. :	0.001564
1st Qu.:	0.0000	1st Qu.:	2.76351	1st Qu.:	0.603723
Median :	0.0000	Median :	7.40385	Median :	0.857623
Mean :	0.2324	Mean :	10.90832	Mean :	0.752468
3rd Qu.:	0.0000	3rd Qu.:	13.54555	3rd Qu.:	0.964409
Max. :	1.0000	Max. :	140.04319	Max. :	0.999832

The covariates between the two groups are unbalanced. The censoring time is dependent on the covariates also, as the covariates are unbalanced between the two groups, the censoring time is also unbalanced. In particular, the mean of X1, X2, X3 and X4 is bigger in the control group than in the treated group. Also, the number of events is bigger in the control than treated group.

9.3 Observational study with interaction

X1		X2		X3		X4	
Min.	:-2.6724	Min.	:-2.8608	Min.	:-2.92627	Min.	:-2.7028
1st Qu.:	-0.1854	1st Qu.:	-0.1614	1st Qu.:	-0.01601	1st Qu.:	-0.2035
Median :	0.5083	Median :	0.4781	Median :	0.69895	Median :	0.5130
Mean :	0.5221	Mean :	0.4889	Mean :	0.69275	Mean :	0.5140
3rd Qu.:	1.1920	3rd Qu.:	1.1382	3rd Qu.:	1.38013	3rd Qu.:	1.2303
Max. :	3.8850	Max. :	3.4889	Max. :	3.79283	Max. :	4.3792

tau		A		T0		T1	
Min.	:0.5	Min.	:0.0000	Min.	: 0.00000	Min.	: 1.000
1st Qu.:	0.5	1st Qu.:	0.0000	1st Qu.:	0.03164	1st Qu.:	1.032
Median :	0.5	Median :	1.0000	Median :	0.18757	Median :	1.188
Mean :	0.5	Mean :	:0.5865	Mean :	0.77911	Mean :	1.779
3rd Qu.:	0.5	3rd Qu.:	:1.0000	3rd Qu.:	0.69832	3rd Qu.:	1.698
Max. :	0.5	Max. :	:1.0000	Max. :	:68.87148	Max. :	:69.871

C		T_obs		T_obs_tau		status	
Min.	: 0.0000	Min.	: 0.00000	Min.	:0.00000	Min.	:0.000
1st Qu.:	0.1453	1st Qu.:	0.05287	1st Qu.:	0.05287	1st Qu.:	0.000
Median :	0.5355	Median :	0.24785	Median :	0.24785	Median :	0.000
Mean :	5.8443	Mean :	0.54939	Mean :	0.26867	Mean :	0.437
3rd Qu.:	1.7237	3rd Qu.:	0.99074	3rd Qu.:	0.50000	3rd Qu.:	1.000
Max. :	:2220.3576	Max. :	:18.47672	Max. :	:0.50000	Max. :	:1.000

status_tau		censor.status		e	
Min.	:0.0000	Min.	:0.000	Min.	:0.0000764
1st Qu.:	0.0000	1st Qu.:	0.000	1st Qu.:	0.4171371
Median :	1.0000	Median :	1.000	Median :	0.6044427
Mean :	0.6045	Mean :	0.563	Mean :	0.5928267
3rd Qu.:	1.0000	3rd Qu.:	1.000	3rd Qu.:	0.8245093
Max. :	1.0000	Max. :	1.000	Max. :	0.9999938

The summary of the generated (observed and unobserved) data set complex observational study (conditionally independent censoring) stratified by treatment is displayed below.

[1] "Descriptive statistics for group A=0: 827"

X1		X2		X3		C	
Min.	:-2.4572	Min.	:-2.86076	Min.	:-1.9492	Min.	: 0.0001
1st Qu.:	-0.2401	1st Qu.:	-0.08007	1st Qu.:	-0.1292	1st Qu.:	0.1601
Median :	0.4972	Median :	0.63144	Median :	0.4602	Median :	0.5402
Mean :	0.4792	Mean :	0.65492	Mean :	0.4339	Mean :	3.8330
3rd Qu.:	1.1795	3rd Qu.:	1.42315	3rd Qu.:	1.0360	3rd Qu.:	1.7511
Max. :	3.8850	Max. :	3.42233	Max. :	2.8309	Max. :	:603.7879

T1		T0		status		T_obs	
Min.	: 1.000	Min.	: 0.0000	Min.	:0.0000	Min.	:0.00000
1st Qu.:	1.027	1st Qu.:	0.0266	1st Qu.:	0.0000	1st Qu.:	0.01867
Median :	1.162	Median :	0.1619	Median :	1.0000	Median :	0.09428
Mean :	1.698	Mean :	0.6980	Mean :	0.6542	Mean :	0.28128
3rd Qu.:	1.692	3rd Qu.:	0.6915	3rd Qu.:	1.0000	3rd Qu.:	0.31235
Max. :	:41.723	Max. :	:40.7229	Max. :	:1.0000	Max. :	:5.88375

status_tau		e	
Min.	:0.0000	Min.	:0.0000764

```

1st Qu.:0.0000  1st Qu.:0.2367490
Median :1.0000  Median :0.4479076
Mean  :0.7243  Mean  :0.4265800
3rd Qu.:1.0000  3rd Qu.:0.5988384
Max.   :1.0000  Max.   :0.9707041

```

[1] "Descriptive statistics for group A=1: 1173"

```

      X1          X2          X3          C
Min.   :-2.6724  Min.   :-2.4416  Min.   :-2.9263  Min.   :  0.0000
1st Qu.:-0.1461  1st Qu.:-0.2560  1st Qu.: 0.1952  1st Qu.:  0.1315
Median : 0.5185  Median : 0.4078  Median : 0.9093  Median :  0.5335
Mean   : 0.5524  Mean   : 0.3718  Mean   : 0.8753  Mean   :  7.2623
3rd Qu.: 1.2016  3rd Qu.: 0.9739  3rd Qu.: 1.6755  3rd Qu.:  1.7117
Max.   : 3.7301  Max.   : 3.4889  Max.   : 3.7928  Max.   :2220.3576

      T1          T0          status      T_obs
Min.   : 1.000    Min.   : 0.00000  Min.   :0.0000  Min.   : 0.000041
1st Qu.: 1.035    1st Qu.: 0.03514  1st Qu.:0.0000  1st Qu.: 0.131473
Median : 1.197    Median : 0.19658  Median :0.0000  Median : 0.533500
Mean   : 1.836    Mean   : 0.83630  Mean   :0.2839  Mean   : 0.738415
3rd Qu.: 1.700    3rd Qu.: 0.69997  3rd Qu.:1.0000  3rd Qu.: 1.064682
Max.   :69.871    Max.   :68.87148  Max.   :1.0000  Max.   :18.476721

      status_tau      e
Min.   :0.00    Min.   :0.05582
1st Qu.:0.00    1st Qu.:0.55353
Median :1.00    Median :0.73499
Mean   :0.52    Mean   :0.71004
3rd Qu.:1.00    3rd Qu.:0.90282
Max.   :1.00    Max.   :0.99999

```

The observations are the same than the previous scenario: The covariates and the censoring time between the two groups are unbalanced. To be able to evaluate the estimators, we need to know the true θ_{RMST} at time τ .