



HAL
open science

Causal survival analysis, Estimation of the Average Treatment Effect (ATE): Practical Recommendations

Charlotte Voinot, Clément Berenfeld, Imke Mayer, Bernard Sebastien, Julie Josse

► **To cite this version:**

Charlotte Voinot, Clément Berenfeld, Imke Mayer, Bernard Sebastien, Julie Josse. Causal survival analysis, Estimation of the Average Treatment Effect (ATE): Practical Recommendations. 2025. hal-04852134

HAL Id: hal-04852134

<https://hal.science/hal-04852134v1>

Preprint submitted on 8 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Causal survival analysis

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

Charlotte Voinot Clément Berenfeld Imke Mayer
Bernard Sebastien Julie Josse

2024-12-20

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.

Table of 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	6
1.4	Notations	6
2	Causal survival analysis in Randomized Controlled Trials	7
2.1	Independent censoring: the Kaplan-Meier estimator	8
2.2	Conditionally independent censoring	11
2.2.1	The G-formula and the Cox Model	12

2.2.2	Censoring unbiased transformations	15
2.2.3	Augmented corrections	20
3	Causal survival analysis in observational studies	20
3.1	Independent censoring	21
3.2	Conditional independent censoring	23
3.2.1	IPTW-IPCW transformations	24
3.2.2	IPTW-BJ transformations	25
3.2.3	Double augmented corrections	26
4	Implementation	27
4.1	Summary of the estimators	27
4.2	Implementation of the estimators	27
4.3	Available packages	30
5	Simulations	31
5.1	RCT	31
5.2	Observational data	35
5.3	Mispecification of nuisance components	38
6	Conclusion	39
7	Disclosure	43
	References	43
A	Proofs	46
A.1	Proofs of Section 2.1	46
A.2	Proofs of Section 2.2	48
A.3	Proofs of Section 3.1	49
A.4	Proofs of Section 3.2	50
B	Descriptive statistics	52
B.1	RCT	52
B.2	Observational study with linear relationship	55
B.3	Observational study with interaction	57

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)

$$T = AT(1) + (1 - A)T(0). \tag{1}$$

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 $\Delta = \mathbb{I}\{T \leq C\}$, 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	Covariates			Treatment	Censoring	Status	Potential		True	Observed
							out-	comes	out-	come
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 $G(t) := \mathbb{P}(C \geq t)$ 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.

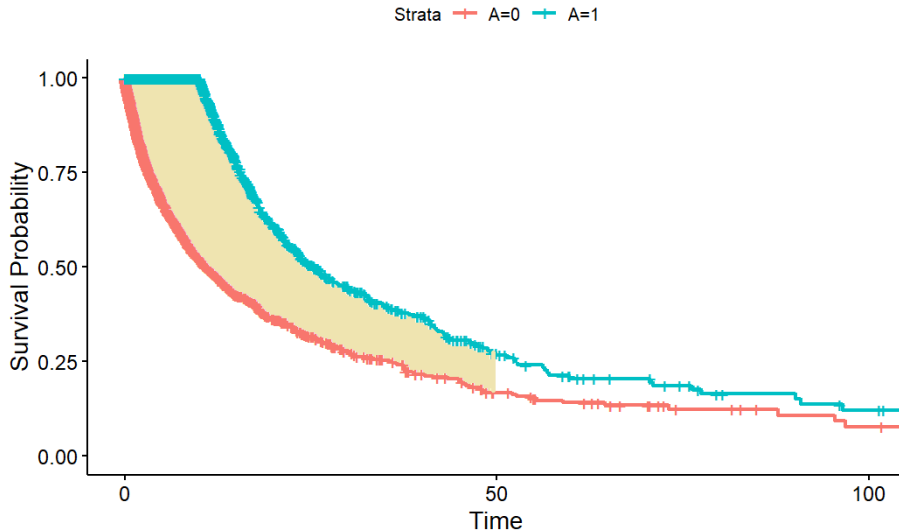


Figure 1: Plot of the estimated survival curves on synthetic toy-data. The θ_{RMST} at $\tau = 50$ corresponds to the yellow shaded area between the two survival curves. The curves have been estimated using Kaplan-Meier estimator, see Section 2.1.

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

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

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

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

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 $\lambda^{(a)}(t) = \lim_{h \rightarrow 0^+} \mathbb{P}(T(a) \in [t, t+h] T(a) \geq t) / h$ 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 $G(t) = \mathbb{P}(C \geq t)$ of the censoring time
\tilde{T}	Observed time ($T \wedge C$)
Δ	Censoring status $\mathbb{I}\{T \leq C\}$
Δ^τ	Censoring status of the restricted time $\Delta^\tau = \max\{\Delta, \mathbb{I}\{\tilde{T} \geq \tau\}\}$

Symbol	Description
$\{t_1, t_2, \dots, t_K\}$	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
$Q_S(t x, a), a \in \{0, 1\}$	$\mathbb{P}(C \geq t X = x, A = a)$
$Q_S(t x, a), a \in \{0, 1\}$	$\mathbb{E}[T \wedge \tau X = x, A = a, T \wedge \tau > t]$

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 $\mathbb{E}[T \wedge \tau | A = a]$ 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)

$$C \perp\!\!\!\perp T(0), T(1), X, A. \tag{8}$$

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

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

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 $\{t_k\}_{k \geq 1}$ be a set of positive and increasing times and assume that $T \in \{t_k\}_{k \geq 1}$ 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

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

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

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

The associated RMST estimator is then simply defined as

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

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 $\hat{S}_{\text{KM}}(t|A=a)$ of $S^{(a)}(t)$ is strongly consistent and admits the following bounds for its bias:*

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

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

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 A.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 $\sqrt{n} \left(\hat{S}_{\text{KM}}(t|A=a) - S^{(a)}(t) \right)$ 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 A.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):

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

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.

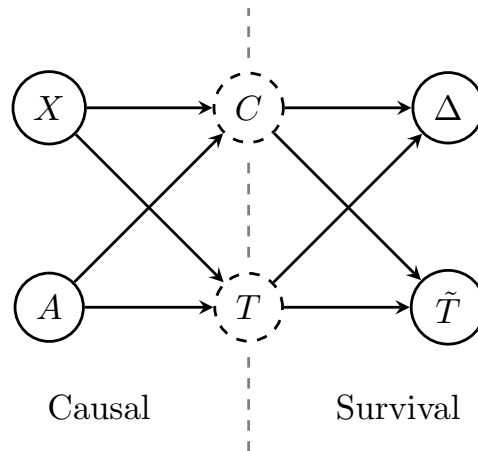


Figure 3: Causal graph in RCT survival data with dependent censoring.

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).

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 :

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

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):

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

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

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

where $\widehat{\beta}^{(a)}$ is a partial likelihood maximizer. One can show that $(\widehat{\beta}^{(a)}, \widehat{\Lambda}_0^{(a)})$ is the MLE of the true likelihood, when $\widehat{\Lambda}_0^{(a)}$ 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 $\widehat{\beta}^{(a)}$ and $\widehat{\Lambda}_0^{(a)}$ 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 $(\widehat{\beta}^{(a)}, \widehat{\Lambda}_0^{(a)})$ 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

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

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).

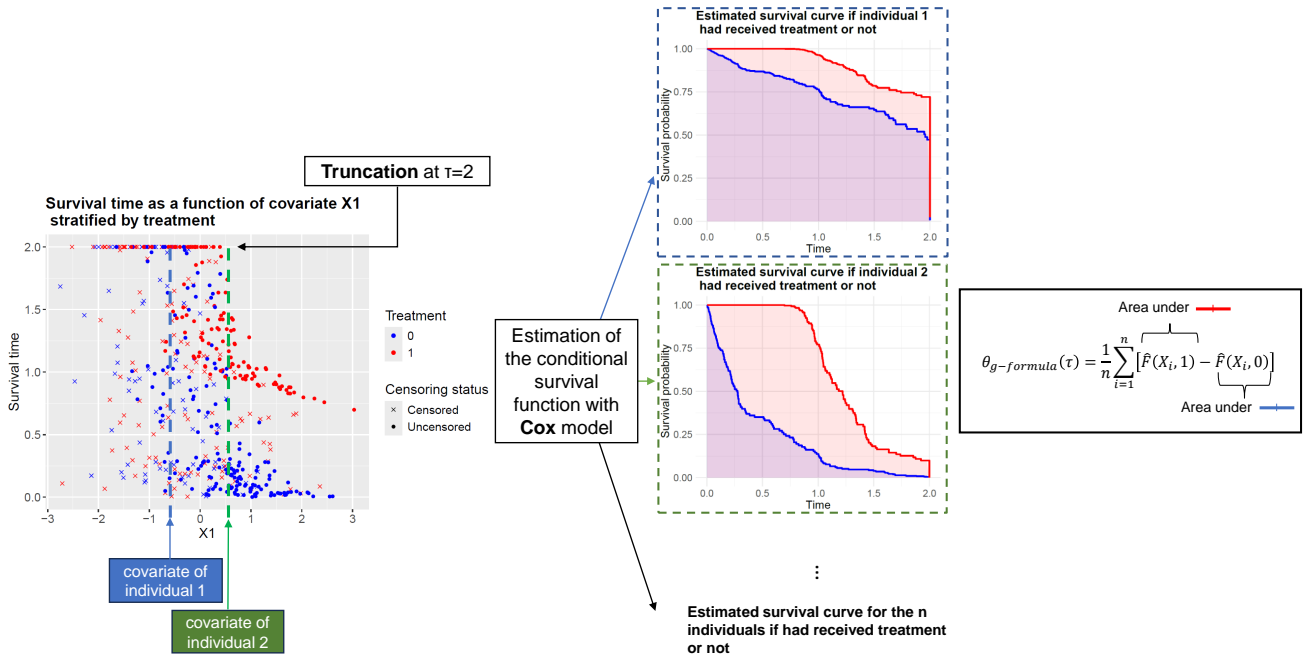


Figure 4: Illustration of the G-formula for estimating θ_{RMST} in an RCT when only one covariate X_1 influences the outcome.

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

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

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 computed 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 $\pi = \mathbb{P}(A = 1) \in (0, 1)$ 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 $n_a = \sum_i \mathbb{I}\{A = a\}$. 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.

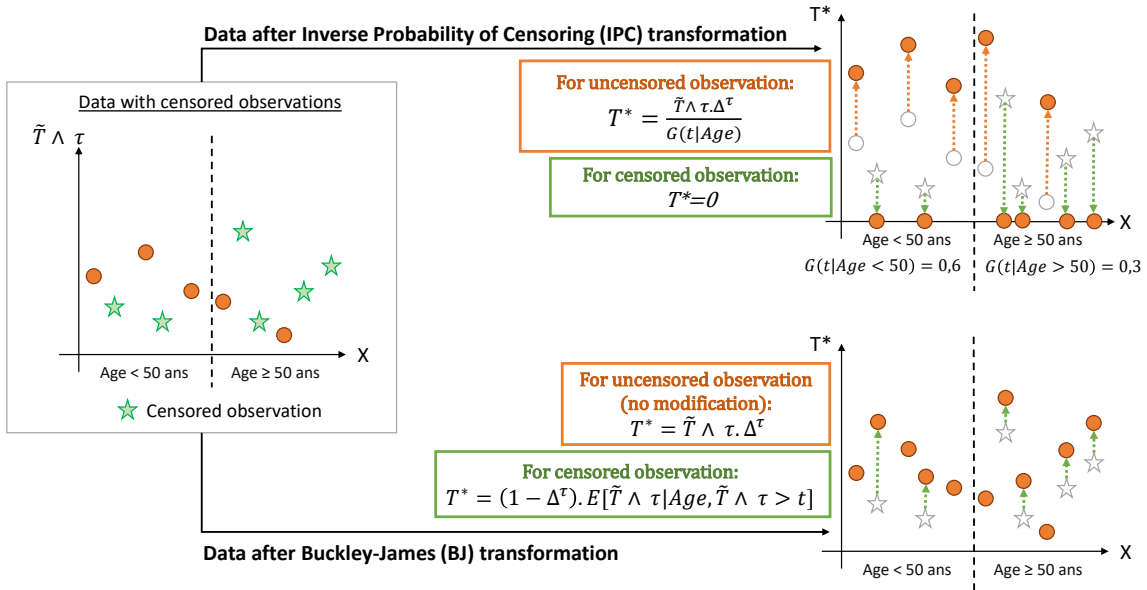


Figure 5: Illustration of Inverse-Probability-of-Censoring and Buckley-James transformations.

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

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

where we recall that $G(t|X, A) := \mathbb{P}(C \geq t|X, A)$ 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 A.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

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

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

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:

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

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

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

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 $\widehat{G} = G$ is a strongly consistent and asymptotically normal estimator of $S^{(a)}(t)$.*

The proof of Proposition 2.6 can be found in Section A.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 $\widehat{G}(\cdot|A, X)$ that we provide is consistent.

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

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 A.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 A.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

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

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

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

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

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

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

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

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 A.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 A.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 $\hat{G}(\cdot|X, A)$ and $\hat{e}(\cdot)$ are uniformly weakly (resp. strongly) consistent estimators, then estimator 40 defined with \hat{e} and \hat{G} is a weakly (resp. strongly) consistent estimator of the difference in RMST.*

The proof of Proposition 3.3 can be found in Section A.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 $\hat{G}(\cdot|X, A)$ and $\hat{e}(\cdot)$ be estimators of the conditional censoring survival function and of the propensity score. We introduce

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

and

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

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

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

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

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

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 $\hat{G}(\cdot|A, X) = G(\cdot|A, X)$ and $\hat{e} = e$ is a strongly consistent and asymptotically normal estimator of $S^{(a)}(t)$.*

The proof of Proposition 3.4 can be found in Section A.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 $\hat{\theta}_{\text{IPTW-IPCW-KM}}$.

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

The proof of Proposition 3.6 can be found in Section A.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 ν 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 $\Delta_{\text{QR}}^* = \Delta_{\text{QR}}^*(F, R, p, \nu)$ satisfies, for $a \in 0, 1$,*

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

This result is similar to Ozenne et al. (2020), Thm 1, and its proof can be found in Section A.4. Based on estimators $(\widehat{G}, \widehat{S}, \widehat{\mu}, \widehat{e})$ 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):

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

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>	✓		✓					
IPCW-KM	✓		✓	✓		<i>G</i>		
BJ	✓		✓	✓	<i>S</i>			
<i>IPTW-KM</i>	✓	✓	✓				<i>e</i>	
IPCW-IPTW-KM	✓	✓	✓	✓		<i>G</i>	<i>e</i>	
IPTW-BJ	✓	✓	✓	✓	<i>S</i>		<i>e</i>	
<i>G-formula</i>	✓	✓	✓	✓	μ			
AIPTW-AIPCW	✓	✓	✓	✓	<i>S, μ</i>	<i>G</i>	<i>e</i>	✓ (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 `utilitary.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

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


- 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

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

- 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 B.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,




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

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

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-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 B.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 $\tau = 25$.




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.

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 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 B.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).

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

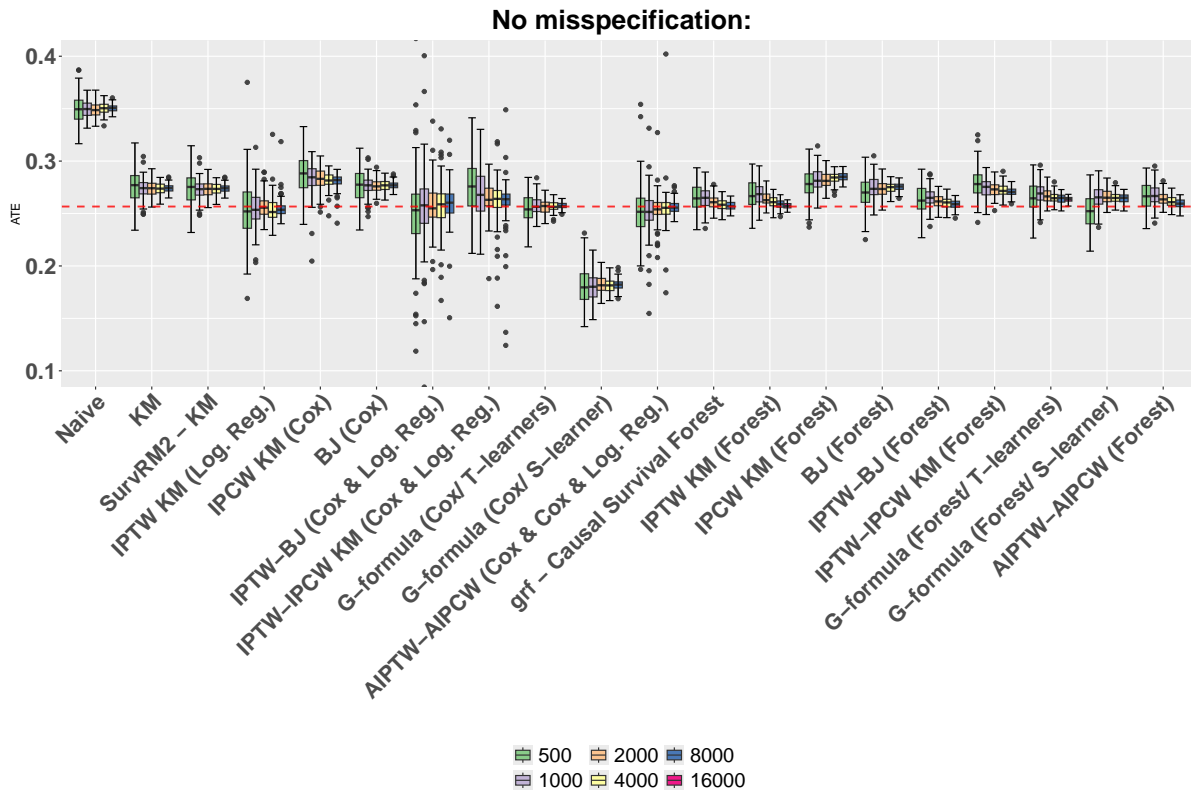


Figure 12: Estimation results of the ATE for the simulation of an observational study with dependent censoring and non linear relationships.

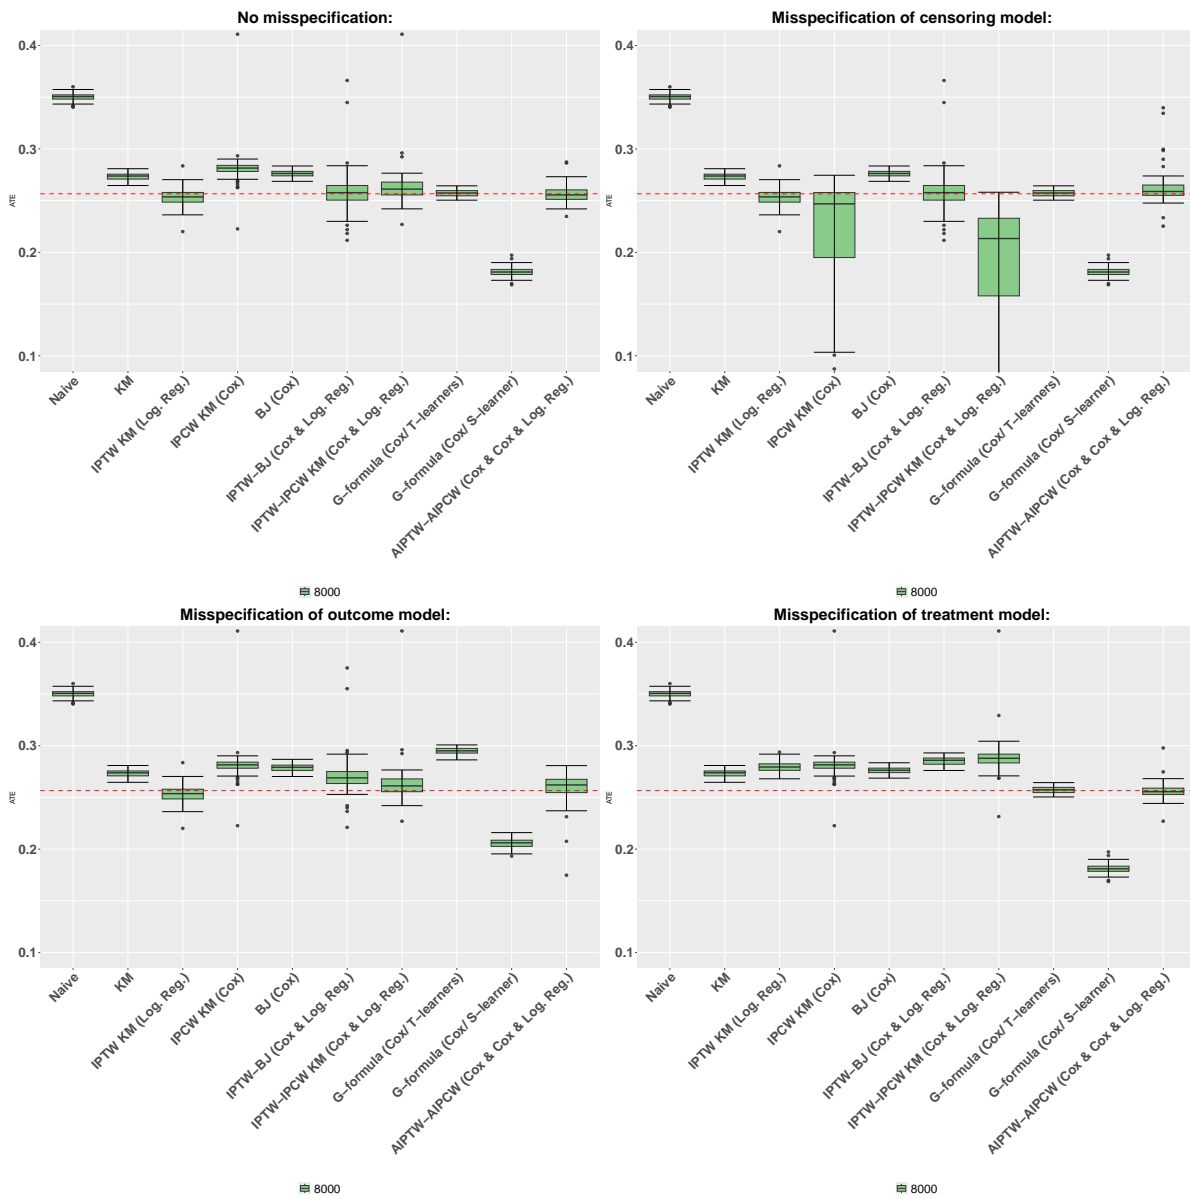


Figure 13: Estimation results of the ATE for an observational study with dependent censoring in case of a single misspecification.

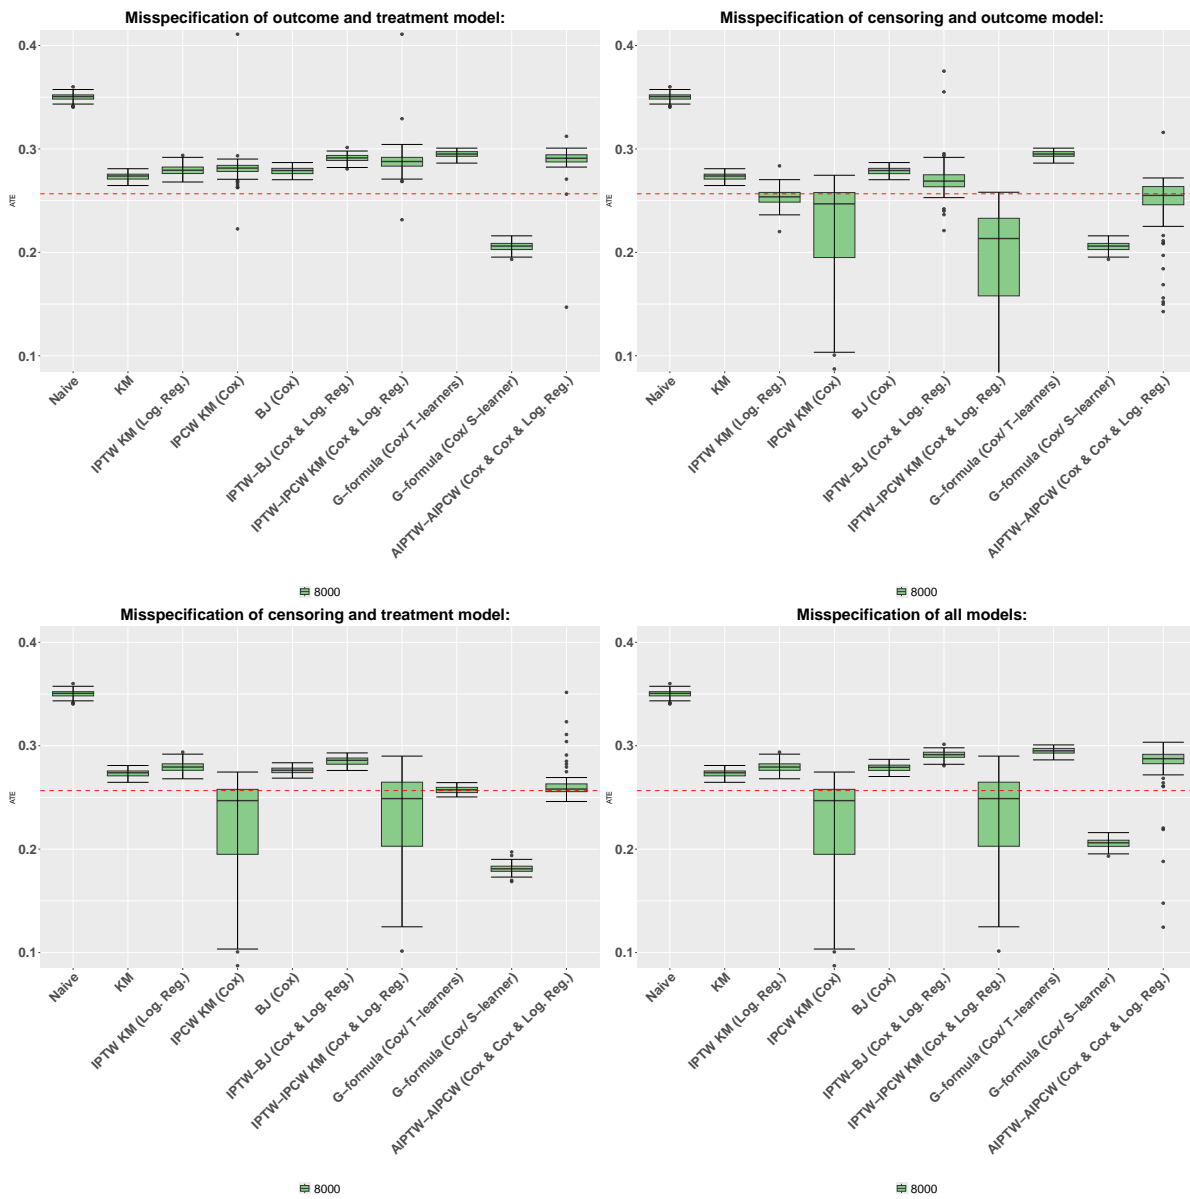


Figure 14: Estimation results of the ATE for an observational study with dependent censoring in case of a two or more misspecifications.

their Cox, Logistic Regression and forest versions. By qualitatively combining the results from these more robust estimators, we 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/qkac001>.
- 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, Battioui 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.

A Proofs

A.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 $k \geq 2$

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

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

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

and then that

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

By induction, we easily find that

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

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

Proof. (Proposition 2.2). The asymptotic normality is a mere consequence of the joint asymptotic normality of $(N_k(a), D_k(a))_{t_k \leq t}$ 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

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

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

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

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

which is what we wanted to show. \square

A.2 Proofs of Section 2.2

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

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

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 $G(T(a) \wedge \tau | A, X) = E[\mathbb{I}\{T(a) \wedge \tau \leq C\} | X, T(a), A]$ 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

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

and likewise that

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

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

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

yielding strong consistency. Asymptotic normality is straightforward. \square

Proof. (Proposition 2.7). There holds

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

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

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

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

A.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 $t_k \leq \tau$

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

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

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

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

A.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,

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

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

Proof. (Proposition 3.3). By consistency of $\hat{G}(\cdot | X, A)$ and \hat{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

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

\square

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 $t_k \leq \tau$,

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

and likewise that

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

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

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. \square

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

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

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

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

Focusing on term (A), we easily find that

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

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

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

and if $\nu(X, 1) = \mu(X, 1)$, 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

B Descriptive statistics

B.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: 961"

X1	X2	X3	X4
Min. : -2.4960	Min. : -2.9526	Min. : -4.4724	Min. : -1.9546
1st Qu.: 0.3032	1st Qu.: 0.3447	1st Qu.: -1.7380	1st Qu.: 0.3532
Median : 0.8963	Median : 0.9851	Median : -1.0741	Median : 0.9564
Mean : 0.9423	Mean : 0.9980	Mean : -1.0465	Mean : 0.9613
3rd Qu.: 1.6394	3rd Qu.: 1.6757	3rd Qu.: -0.3965	3rd Qu.: 1.5836
Max. : 4.2113	Max. : 4.0543	Max. : 1.7716	Max. : 3.9881
C	T1	T0	status
Min. : 0.0258	Min. : 10.00	Min. : 0.0003	Min. : 0.0000
1st Qu.: 9.1498	1st Qu.: 12.92	1st Qu.: 2.9198	1st Qu.: 0.0000
Median : 21.8475	Median : 19.01	Median : 9.0074	Median : 1.0000
Mean : 32.6733	Mean : 34.16	Mean : 24.1588	Mean : 0.6712
3rd Qu.: 45.7986	3rd Qu.: 34.64	3rd Qu.: 24.6408	3rd Qu.: 1.0000
Max. : 245.1850	Max. : 611.65	Max. : 601.6491	Max. : 1.0000
T_tild			
Min. : 0.00032			
1st Qu.: 2.33634			
Median : 6.06026			
Mean : 11.17432			
3rd Qu.: 14.55178			
Max. : 121.50045			

[1] "Descriptive statistics for group A=1: 1039"

X1	X2	X3	X4
Min. : -1.8689	Min. : -2.4225	Min. : -4.1827	Min. : -2.1940
1st Qu.: 0.3197	1st Qu.: 0.3776	1st Qu.: -1.6456	1st Qu.: 0.3426
Median : 0.9432	Median : 0.9893	Median : -1.0582	Median : 0.9900
Mean : 0.9653	Mean : 1.0140	Mean : -1.0098	Mean : 0.9782
3rd Qu.: 1.6105	3rd Qu.: 1.6313	3rd Qu.: -0.2902	3rd Qu.: 1.6278
Max. : 4.1308	Max. : 4.1762	Max. : 2.1513	Max. : 3.7267

C	T1	T0	status
Min. : 0.00721	Min. : 10.01	Min. : 0.0108	Min. : 0.0000
1st Qu.: 9.40216	1st Qu.: 13.01	1st Qu.: 3.0117	1st Qu.: 0.0000
Median : 23.30619	Median : 17.97	Median : 7.9681	Median : 1.0000
Mean : 33.18284	Mean : 34.27	Mean : 24.2703	Mean : 0.5024
3rd Qu.: 44.63882	3rd Qu.: 33.05	3rd Qu.: 23.0469	3rd Qu.: 1.0000
Max. : 308.12079	Max. : 1126.22	Max. : 1116.2166	Max. : 1.0000

T_tild
Min. : 0.00721
1st Qu.: 9.40216
Median : 13.05491
Mean : 16.46654
3rd Qu.: 19.77933
Max. : 172.92122

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: 1029"

X1	X2	X3	X4
Min. : -2.0253	Min. : -2.4160	Min. : -4.0326	Min. : -2.2743
1st Qu.: 0.2919	1st Qu.: 0.3481	1st Qu.: -1.6398	1st Qu.: 0.2620
Median : 0.9594	Median : 1.0196	Median : -1.0286	Median : 0.9822
Mean : 0.9946	Mean : 1.0114	Mean : -0.9952	Mean : 0.9724
3rd Qu.: 1.6849	3rd Qu.: 1.6727	3rd Qu.: -0.3055	3rd Qu.: 1.6715
Max. : 3.8826	Max. : 4.1515	Max. : 2.0179	Max. : 3.9276

C	T1	T0	status
---	----	----	--------

Min. :	0.0012	Min. :	10.03	Min. :	0.0253	Min. :	0.0000
1st Qu.:	2.4730	1st Qu.:	13.26	1st Qu.:	3.2645	1st Qu.:	0.0000
Median :	6.5279	Median :	19.07	Median :	9.0691	Median :	0.0000
Mean :	13.2357	Mean :	31.72	Mean :	21.7172	Mean :	0.4101
3rd Qu.:	15.2260	3rd Qu.:	33.59	3rd Qu.:	23.5902	3rd Qu.:	1.0000
Max. :	332.1482	Max. :	489.85	Max. :	479.8539	Max. :	1.0000
	status_tau		T_tild				
Min. :	0.0000	Min. :	0.00118				
1st Qu.:	0.0000	1st Qu.:	1.32910				
Median :	0.0000	Median :	3.59519				
Mean :	0.4402	Mean :	7.55911				
3rd Qu.:	1.0000	3rd Qu.:	8.83735				
Max. :	1.0000	Max. :	107.04441				

[1] "Descriptive statistics for group A=1: 971"

	X1	X2	X3	X4			
Min. :	-2.1394	Min. : -1.9388	Min. : -4.1306	Min. : -1.9867			
1st Qu.:	0.4003	1st Qu.:	0.3611	1st Qu.:	0.2856		
Median :	1.0362	Median :	1.0574	Median :	-0.9457	Median :	0.9863
Mean :	1.0374	Mean :	1.0611	Mean :	-0.9601	Mean :	0.9879
3rd Qu.:	1.6908	3rd Qu.:	1.7591	3rd Qu.:	-0.3385	3rd Qu.:	1.7159
Max. :	4.0722	Max. :	4.2386	Max. :	2.3858	Max. :	4.3266
	C	T1	T0	status			
Min. :	0.00532	Min. : 10.01	Min. : 0.0052	Min. : 0.0000			
1st Qu.:	2.92617	1st Qu.:	12.71	1st Qu.:	2.7145	1st Qu.:	0.0000
Median :	8.37240	Median :	18.85	Median :	8.8522	Median :	0.0000
Mean :	16.47726	Mean :	32.19	Mean :	22.1908	Mean :	0.2029
3rd Qu.:	20.25093	3rd Qu.:	33.88	3rd Qu.:	23.8824	3rd Qu.:	0.0000
Max. :	207.92277	Max. :	527.89	Max. :	517.8923	Max. :	1.0000
	status_tau		T_tild				
Min. :	0.0000	Min. :	0.00532				
1st Qu.:	0.0000	1st Qu.:	2.92617				
Median :	0.0000	Median :	8.37240				
Mean :	0.2678	Mean :	11.65063				
3rd Qu.:	1.0000	3rd Qu.:	14.77252				
Max. :	1.0000	Max. :	121.13642				

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.

B.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: 1097"

X1		X2		X3		X4	
Min.	:-2.1641	Min.	:-1.8109	Min.	:-2.8668	Min.	:-1.6736
1st Qu.:	0.5008	1st Qu.:	0.5243	1st Qu.:	-1.0737	1st Qu.:	0.5518
Median :	1.1753	Median :	1.1879	Median :	-0.5283	Median :	1.2237
Mean :	1.1725	Mean :	1.1822	Mean :	-0.4929	Mean :	1.2263
3rd Qu.:	1.8411	3rd Qu.:	1.8649	3rd Qu.:	0.0493	3rd Qu.:	1.9119
Max. :	4.1548	Max. :	4.5749	Max. :	2.2581	Max. :	4.5120
C		T1		T0		status	
Min. :	0.04185	Min. :	10.00	Min. :	0.0005	Min. :	0.0000
1st Qu.:	9.95416	1st Qu.:	12.53	1st Qu.:	2.5266	1st Qu.:	0.0000
Median :	22.61758	Median :	17.82	Median :	7.8249	Median :	1.0000
Mean :	32.62643	Mean :	32.25	Mean :	22.2488	Mean :	0.6764
3rd Qu.:	46.33098	3rd Qu.:	33.80	3rd Qu.:	23.7966	3rd Qu.:	1.0000
Max. :	251.27060	Max. :	527.79	Max. :	517.7890	Max. :	1.0000
T_tild							
Min. :	0.00049						
1st Qu.:	1.81799						
Median :	5.38692						
Mean :	10.00168						
3rd Qu.:	13.04893						
Max. :	117.13428						

[1] "Descriptive statistics for group A=1: 903"

X1		X2		X3		X4	
Min.	:-2.42956	Min.	:-2.55388	Min.	:-4.2621	Min.	:-2.3497
1st Qu.:	0.02135	1st Qu.:	0.03554	1st Qu.:	-2.1193	1st Qu.:	0.1442
Median :	0.70026	Median :	0.68627	Median :	-1.6093	Median :	0.7740
Mean :	0.70584	Mean :	0.69164	Mean :	-1.6077	Mean :	0.7734
3rd Qu.:	1.37383	3rd Qu.:	1.33898	3rd Qu.:	-1.0284	3rd Qu.:	1.4302
Max. :	4.28548	Max. :	3.50823	Max. :	0.9609	Max. :	3.5146
C		T1		T0		status	
Min. :	0.01247	Min. :	10.01	Min. :	0.0113	Min. :	0.0000

1st Qu.:	9.89830	1st Qu.:	13.51	1st Qu.:	3.5132	1st Qu.:	0.0000
Median :	22.04596	Median :	20.37	Median :	10.3705	Median :	0.0000
Mean :	32.37769	Mean :	35.97	Mean :	25.9707	Mean :	0.4839
3rd Qu.:	44.29837	3rd Qu.:	36.94	3rd Qu.:	26.9384	3rd Qu.:	1.0000
Max. :	283.66431	Max. :	622.86	Max. :	612.8589	Max. :	1.0000

T_tild

Min. :	0.01247
1st Qu.:	9.89830
Median :	13.39646
Mean :	17.43497
3rd Qu.:	20.82123
Max. :	176.02564

The covariates between the two groups of treatment are unbalanced because of dependent treatment assignment. The mean of X_1 , X_2 , X_3 and X_4 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: 1099"

X1	X2	X3	X4
Min. : -2.1278	Min. : -1.626	Min. : -3.13844	Min. : -2.1011
1st Qu.: 0.4903	1st Qu.: 0.480	1st Qu.: -1.03980	1st Qu.: 0.5833
Median : 1.2076	Median : 1.186	Median : -0.49548	Median : 1.2356
Mean : 1.1778	Mean : 1.172	Mean : -0.46597	Mean : 1.2288
3rd Qu.: 1.8542	3rd Qu.: 1.875	3rd Qu.: 0.09934	3rd Qu.: 1.8959
Max. : 4.3029	Max. : 4.357	Max. : 2.47981	Max. : 4.2822

C	T1	T0	status
Min. : 0.02418	Min. : 10.00	Min. : 0.0027	Min. : 0.0000
1st Qu.: 2.63552	1st Qu.: 12.81	1st Qu.: 2.8065	1st Qu.: 0.0000
Median : 7.11481	Median : 18.53	Median : 8.5258	Median : 0.0000
Mean : 14.75487	Mean : 30.73	Mean : 20.7272	Mean : 0.4622
3rd Qu.: 16.93227	3rd Qu.: 31.44	3rd Qu.: 21.4443	3rd Qu.: 1.0000
Max. : 269.87017	Max. : 521.90	Max. : 511.8959	Max. : 1.0000

status_tau	T_obs	e
Min. : 0.0000	Min. : 0.00269	Min. : 0.0000141
1st Qu.: 0.0000	1st Qu.: 1.42605	1st Qu.: 0.0217159
Median : 0.0000	Median : 4.02644	Median : 0.0939441

```

Mean      :0.4959   Mean      : 7.92379   Mean      :0.1977888
3rd Qu.  :1.0000   3rd Qu.  : 9.20963   3rd Qu.  :0.2939396
Max.     :1.0000   Max.     :157.60939   Max.     :0.9783493

```

```
[1] "Descriptive statistics for group A=1: 901"
```

```

      X1              X2              X3              X4
Min.   :-2.21674   Min.   :-2.0999   Min.   :-4.640   Min.   :-2.75769
1st Qu.: 0.08678   1st Qu.: 0.1222   1st Qu.: -2.244   1st Qu.: 0.05356
Median : 0.74634   Median : 0.7908   Median : -1.684   Median : 0.71404
Mean   : 0.73616   Mean   : 0.7853   Mean   : -1.671   Mean   : 0.71517
3rd Qu.: 1.38391   3rd Qu.: 1.4753   3rd Qu.: -1.100   3rd Qu.: 1.35296
Max.   : 4.07993   Max.   : 3.2937   Max.   : 1.026   Max.   : 4.13187

      C              T1              T0              status
Min.   : 0.00318   Min.   : 10.01   Min.   : 0.0138   Min.   :0.000
1st Qu.: 2.67961   1st Qu.: 12.94   1st Qu.: 2.9443   1st Qu.:0.000
Median : 7.69254   Median : 19.58   Median : 9.5772   Median :0.000
Mean   : 15.58449   Mean   : 32.69   Mean   : 22.6936   Mean   :0.182
3rd Qu.: 19.14077   3rd Qu.: 34.69   3rd Qu.: 24.6859   3rd Qu.:0.000
Max.   :300.46870   Max.   :737.24   Max.   :727.2356   Max.   :1.000

      status_tau      T_obs              e
Min.   :0.0000   Min.   : 0.00318   Min.   :0.01787
1st Qu.:0.0000   1st Qu.: 2.67961   1st Qu.:0.61911
Median :0.0000   Median : 7.69254   Median :0.85517
Mean   :0.2364   Mean   : 11.13241   Mean   :0.75406
3rd Qu.:0.0000   3rd Qu.: 14.36652   3rd Qu.:0.96434
Max.   :1.0000   Max.   :103.32168   Max.   :0.99997

```

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.

B.3 Observational study with interaction

```

      X1              X2              X3              X4
Min.   :-3.1050   Min.   :-2.8416   Min.   :-3.05920   Min.   :-3.2757
1st Qu.: -0.1533   1st Qu.: -0.1816   1st Qu.: -0.01324   1st Qu.: -0.1819
Median : 0.4849   Median : 0.4709   Median : 0.67314   Median : 0.4667
Mean   : 0.4936   Mean   : 0.5001   Mean   : 0.69411   Mean   : 0.4590

```

3rd Qu.: 1.1758	3rd Qu.: 1.1659	3rd Qu.: 1.38927	3rd Qu.: 1.0941
Max. : 3.7534	Max. : 4.0851	Max. : 4.16151	Max. : 3.3850
tau	A	T0	T1
Min. :0.5	Min. :0.0000	Min. : 0.0000	Min. : 1.000
1st Qu.:0.5	1st Qu.:0.0000	1st Qu.: 0.0364	1st Qu.: 1.036
Median :0.5	Median :1.0000	Median : 0.2033	Median : 1.203
Mean :0.5	Mean :0.5665	Mean : 0.9848	Mean : 1.985
3rd Qu.:0.5	3rd Qu.:1.0000	3rd Qu.: 0.7503	3rd Qu.: 1.750
Max. :0.5	Max. :1.0000	Max. :343.7007	Max. :344.701
C	T_obs	T_obs_tau	status
Min. : 0.000	Min. : 0.00000	Min. :0.000002	Min. :0.000
1st Qu.: 0.146	1st Qu.: 0.05307	1st Qu.:0.053073	1st Qu.:0.000
Median : 0.555	Median : 0.23531	Median :0.235310	Median :0.000
Mean : 25.188	Mean : 0.58820	Mean :0.264448	Mean :0.463
3rd Qu.: 1.931	3rd Qu.: 0.83602	3rd Qu.:0.500000	3rd Qu.:1.000
Max. :30982.747	Max. :151.25974	Max. :0.500000	Max. :1.000
status_tau	censor.status	e	
Min. :0.000	Min. :0.000	Min. :0.0000003	
1st Qu.:0.000	1st Qu.:0.000	1st Qu.:0.4297210	
Median :1.000	Median :1.000	Median :0.5971153	
Mean :0.609	Mean :0.537	Mean :0.5884060	
3rd Qu.:1.000	3rd Qu.:1.000	3rd Qu.:0.7977037	
Max. :1.000	Max. :1.000	Max. :0.9999122	

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: 867"

X1	X2	X3	C
Min. :-3.10500	Min. :-2.841601	Min. :-1.8996	Min. : 0.000
1st Qu.: -0.07291	1st Qu.: -0.002758	1st Qu.: -0.1253	1st Qu.: 0.174
Median : 0.52177	Median : 0.708172	Median : 0.4257	Median : 0.633
Mean : 0.53635	Mean : 0.716400	Mean : 0.4548	Mean : 48.170
3rd Qu.: 1.23299	3rd Qu.: 1.425439	3rd Qu.: 1.0251	3rd Qu.: 2.098
Max. : 3.75343	Max. : 3.862628	Max. : 2.5052	Max. :30982.747
T1	T0	status	T_obs
Min. : 1.000	Min. : 0.00000	Min. :0.0000	Min. :0.000002
1st Qu.: 1.027	1st Qu.: 0.02734	1st Qu.:0.0000	1st Qu.:0.018567
Median : 1.158	Median : 0.15835	Median :1.0000	Median :0.088354
Mean : 1.616	Mean : 0.61591	Mean :0.7001	Mean :0.245368
3rd Qu.: 1.586	3rd Qu.: 0.58625	3rd Qu.:1.0000	3rd Qu.:0.307923

Max. :65.745	Max. :64.74474	Max. :1.0000	Max. :6.465842
status_tau	e		
Min. :0.0000	Min. :0.0000003		
1st Qu.:0.0000	1st Qu.:0.2470120		
Median :1.0000	Median :0.4678179		
Mean :0.7463	Mean :0.4317838		
3rd Qu.:1.0000	3rd Qu.:0.5951117		
Max. :1.0000	Max. :0.9871857		

[1] "Descriptive statistics for group A=1: 1133"

X1	X2	X3	C
Min. :-2.9456	Min. :-2.2904	Min. :-3.0592	Min. : 0.000
1st Qu.: -0.1914	1st Qu.: -0.3197	1st Qu.: 0.1134	1st Qu.: 0.134
Median : 0.4645	Median : 0.3072	Median : 0.9322	Median : 0.508
Mean : 0.4609	Mean : 0.3346	Mean : 0.8772	Mean : 7.601
3rd Qu.: 1.1406	3rd Qu.: 0.9584	3rd Qu.: 1.6611	3rd Qu.: 1.805
Max. : 3.6716	Max. : 4.0851	Max. : 4.1615	Max. :3333.917
T1	T0	status	T_obs
Min. : 1.000	Min. : 0.0000	Min. :0.0000	Min. : 0.00002
1st Qu.: 1.048	1st Qu.: 0.0475	1st Qu.:0.0000	1st Qu.: 0.13409
Median : 1.240	Median : 0.2405	Median :0.0000	Median : 0.50780
Mean : 2.267	Mean : 1.2671	Mean :0.2816	Mean : 0.85055
3rd Qu.: 1.885	3rd Qu.: 0.8852	3rd Qu.:1.0000	3rd Qu.: 1.05814
Max. :344.701	Max. :343.7007	Max. :1.0000	Max. :151.25974
status_tau	e		
Min. :0.000	Min. :0.004657		
1st Qu.:0.000	1st Qu.:0.551534		
Median :1.000	Median :0.725028		
Mean :0.504	Mean :0.708257		
3rd Qu.:1.000	3rd Qu.:0.898320		
Max. :1.000	Max. :0.999912		

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 .